
GEOSX Documentation

Randolph Settgast

Sep 11, 2021

Contents

1	Quick Start Guide	3
1.1	Frequently Asked Questions	3
1.2	Repository Organization	4
1.3	Username and Authentication	5
1.4	Download	5
1.5	Configuration	7
1.6	Compilation	8
1.7	Running	9
1.8	Testing	9
2	Tutorials	11
2.1	Tutorial 1: First Steps	11
2.2	Tutorial 2: External Meshes	24
2.3	Tutorial 3: Regions and Property Specifications	33
2.4	Tutorial 4: Boundary Conditions and Time-Dependent Functions	44
3	Basic Examples	49
3.1	Multiphase Flow	49
3.2	Multiphase Flow with Wells	58
3.3	CO ₂ Injection	71
3.4	Poromechanics	80
3.5	Hydraulic Fracturing	85
4	Advanced Examples	99
4.1	Validation and Verification Studies	99
4.2	Performance Benchmarks	128
4.3	Application Studies	128
5	User Guide	129
5.1	Input Files	129
5.2	Meshes	132
5.3	Physics Solvers	139
5.4	Constitutive Models	168
5.5	Event Management	210
5.6	Tasks Manager	215
5.7	Functions	216
5.8	Linear Solvers	221

5.9	Numerical Methods	225
5.10	Parallel Partitioning	227
5.11	Outputs	228
5.12	Advanced XML Features (geosx_xml_tools)	231
5.13	pygeosx — GEOSX in Python	235
5.14	Indices and tables	237
6	Developer Guide	239
6.1	Contributing	239
6.2	Code Components	270
7	Doxygen	327
8	Build Guide	329
8.1	System prerequisites	329
8.2	Third-party dependencies	330
8.3	Building GEOSX	333
8.4	Spack and Uberenv	334
8.5	Continuous Integration process	337
9	Datastructure Index	339
9.1	Input Schema Definitions	339
9.2	Element: AcousticSEM	340
9.3	Element: Benchmarks	341
9.4	Element: BiotPorosity	341
9.5	Element: BlackOilFluid	342
9.6	Element: Blueprint	343
9.7	Element: BoundedPlane	343
9.8	Element: Box	343
9.9	Element: BrooksCoreyBakerRelativePermeability	344
9.10	Element: BrooksCoreyCapillaryPressure	345
9.11	Element: BrooksCoreyRelativePermeability	345
9.12	Element: CO2BrineFluid	345
9.13	Element: CarmanKozenyPermeability	346
9.14	Element: CellElementRegion	346
9.15	Element: ChomboIO	346
9.16	Element: CompositeFunction	346
9.17	Element: CompositionalMultiphaseFVM	348
9.18	Element: CompositionalMultiphaseFluid	349
9.19	Element: CompositionalMultiphaseHybridFVM	351
9.20	Element: CompositionalMultiphaseReservoir	352
9.21	Element: CompositionalMultiphaseWell	353
9.22	Element: CompressibleSinglePhaseFluid	354
9.23	Element: CompressibleSolidCarmanKozenyPermeability	354
9.24	Element: CompressibleSolidConstantPermeability	354
9.25	Element: CompressibleSolidParallelPlatesPermeability	355
9.26	Element: ConstantPermeability	355
9.27	Element: Constitutive	355
9.28	Element: Contact	356
9.29	Element: Coulomb	357
9.30	Element: Cylinder	357
9.31	Element: DamageElasticIsotropic	358
9.32	Element: DamageSpectralElasticIsotropic	358
9.33	Element: DamageVolDevElasticIsotropic	358
9.34	Element: DeadOilFluid	359

9.35	Element: Dirichlet	360
9.36	Element: DruckerPrager	360
9.37	Element: ElasticIsotropic	360
9.38	Element: ElasticOrthotropic	361
9.39	Element: ElasticTransverseIsotropic	361
9.40	Element: ElementRegions	361
9.41	Element: EmbeddedSurfaceGenerator	362
9.42	Element: Events	362
9.43	Element: ExtendedDruckerPrager	363
9.44	Element: FieldSpecification	363
9.45	Element: FieldSpecifications	363
9.46	Element: File	364
9.47	Element: FiniteElementSpace	364
9.48	Element: FiniteElements	364
9.49	Element: FiniteVolume	364
9.50	Element: FlowProppantTransport	365
9.51	Element: Functions	365
9.52	Element: Geometry	366
9.53	Element: HaltEvent	366
9.54	Element: HybridMimeticDiscretization	366
9.55	Element: Hydrofracture	368
9.56	Element: Included	369
9.57	Element: InternalMesh	369
9.58	Element: InternalWell	370
9.59	Element: InternalWellbore	371
9.60	Element: LagrangianContact	372
9.61	Element: LaplaceFEM	373
9.62	Element: LaplaceVEM	374
9.63	Element: LinearSolverParameters	376
9.64	Element: Mesh	377
9.65	Element: MultiphasePoromechanics	377
9.66	Element: NonlinearSolverParameters	379
9.67	Element: NullModel	380
9.68	Element: NumericalMethods	380
9.69	Element: Outputs	380
9.70	Element: PAMELAMeshGenerator	380
9.71	Element: PackCollection	381
9.72	Element: ParallelPlatesPermeability	381
9.73	Element: Parameter	381
9.74	Element: Parameters	381
9.75	Element: ParticleFluid	382
9.76	Element: Perforation	382
9.77	Element: PeriodicEvent	383
9.78	Element: PermeabilityBase	383
9.79	Element: PhaseFieldDamageFEM	384
9.80	Element: PhaseFieldFracture	386
9.81	Element: PorousDruckerPrager	387
9.82	Element: PorousElasticIsotropic	387
9.83	Element: PorousElasticOrthotropic	387
9.84	Element: PorousElasticTransverseIsotropic	387
9.85	Element: PorousExtendedDruckerPrager	387
9.86	Element: PressurePorosity	388
9.87	Element: Problem	388
9.88	Element: ProppantSlurryFluid	388

9.89	Element: ProppantTransport	390
9.90	Element: Python	391
9.91	Element: Restart	391
9.92	Element: Run	391
9.93	Element: Silo	392
9.94	Element: SinglePhaseFVM	393
9.95	Element: SinglePhaseHybridFVM	394
9.96	Element: SinglePhasePoromechanics	395
9.97	Element: SinglePhasePoromechanicsEmbeddedFractures	396
9.98	Element: SinglePhaseProppantFVM	397
9.99	Element: SinglePhaseReservoir	398
9.100	Element: SinglePhaseWell	399
9.101	Element: SolidMechanicsEmbeddedFractures	400
9.102	Element: SolidMechanicsLagrangianSSLE	402
9.103	Element: SolidMechanics_LagrangianFEM	404
9.104	Element: SoloEvent	405
9.105	Element: Solvers	406
9.106	Element: SourceFlux	407
9.107	Element: StrainDependentPermeability	407
9.108	Element: SurfaceElementRegion	407
9.109	Element: SurfaceGenerator	408
9.110	Element: SymbolicFunction	409
9.111	Element: TableFunction	409
9.112	Element: TableRelativePermeability	410
9.113	Element: Tasks	410
9.114	Element: ThickPlane	410
9.115	Element: TimeHistory	411
9.116	Element: Traction	412
9.117	Element: TriaxialDriver	413
9.118	Element: TwoPointFluxApproximation	413
9.119	Element: VTK	413
9.120	Element: VanGenuchtenBakerRelativePermeability	414
9.121	Element: VanGenuchtenCapillaryPressure	415
9.122	Element: WellControls	417
9.123	Element: WellElementRegion	418
9.124	Element: lassen	418
9.125	Element: quartz	418
9.126	Datastructure: AcousticSEM	419
9.127	Datastructure: Benchmarks	419
9.128	Datastructure: BiotPorosity	420
9.129	Datastructure: BlackOilFluid	420
9.130	Datastructure: Blueprint	421
9.131	Datastructure: BoundedPlane	421
9.132	Datastructure: Box	421
9.133	Datastructure: BrooksCoreyBakerRelativePermeability	421
9.134	Datastructure: BrooksCoreyCapillaryPressure	421
9.135	Datastructure: BrooksCoreyRelativePermeability	422
9.136	Datastructure: CO2BrineFluid	424
9.137	Datastructure: CarmanKozenyPermeability	425
9.138	Datastructure: CellElementRegion	425
9.139	Datastructure: ChomboIO	425
9.140	Datastructure: CompositeFunction	425
9.141	Datastructure: CompositionalMultiphaseFVM	426
9.142	Datastructure: CompositionalMultiphaseFluid	428

9.143	Datastructure: CompositionalMultiphaseHybridFVM	429
9.144	Datastructure: CompositionalMultiphaseReservoir	429
9.145	Datastructure: CompositionalMultiphaseWell	429
9.146	Datastructure: CompressibleSinglePhaseFluid	430
9.147	Datastructure: CompressibleSolidCarmanKozenyPermeability	430
9.148	Datastructure: CompressibleSolidConstantPermeability	430
9.149	Datastructure: CompressibleSolidParallelPlatesPermeability	430
9.150	Datastructure: ConstantPermeability	430
9.151	Datastructure: Constitutive	430
9.152	Datastructure: ConstitutiveModels	431
9.153	Datastructure: Contact	431
9.154	Datastructure: Coulomb	432
9.155	Datastructure: Cylinder	432
9.156	Datastructure: DamageElasticIsotropic	432
9.157	Datastructure: DamageSpectralElasticIsotropic	432
9.158	Datastructure: DamageVolDevElasticIsotropic	432
9.159	Datastructure: DeadOilFluid	433
9.160	Datastructure: Dirichlet	433
9.161	Datastructure: DruckerPrager	434
9.162	Datastructure: ElasticIsotropic	434
9.163	Datastructure: ElasticOrthotropic	434
9.164	Datastructure: ElasticTransverseIsotropic	435
9.165	Datastructure: ElementRegions	435
9.166	Datastructure: EmbeddedSurfaceGenerator	436
9.167	Datastructure: Events	436
9.168	Datastructure: ExtendedDruckerPrager	437
9.169	Datastructure: FaceManager	437
9.170	Datastructure: FieldSpecification	439
9.171	Datastructure: FieldSpecifications	439
9.172	Datastructure: File	439
9.173	Datastructure: FiniteElementSpace	439
9.174	Datastructure: FiniteElements	439
9.175	Datastructure: FiniteVolume	439
9.176	Datastructure: FlowProppantTransport	440
9.177	Datastructure: Functions	440
9.178	Datastructure: Geometry	440
9.179	Datastructure: HaltEvent	440
9.180	Datastructure: HybridMimeticDiscretization	441
9.181	Datastructure: Hydrofracture	441
9.182	Datastructure: Included	441
9.183	Datastructure: InternalMesh	441
9.184	Datastructure: InternalWell	441
9.185	Datastructure: InternalWellbore	442
9.186	Datastructure: LagrangianContact	442
9.187	Datastructure: LaplaceFEM	442
9.188	Datastructure: LaplaceVEM	442
9.189	Datastructure: Level0	443
9.190	Datastructure: LinearSolverParameters	443
9.191	Datastructure: Mesh	443
9.192	Datastructure: MeshBodies	443
9.193	Datastructure: MultiphasePoromechanics	443
9.194	Datastructure: NonlinearSolverParameters	444
9.195	Datastructure: NullModel	444
9.196	Datastructure: NumericalMethods	444

9.197	Datastructure: Outputs	444
9.198	Datastructure: PAMELAMeshGenerator	444
9.199	Datastructure: PackCollection	444
9.200	Datastructure: ParallelPlatesPermeability	445
9.201	Datastructure: Parameter	445
9.202	Datastructure: Parameters	445
9.203	Datastructure: ParticleFluid	445
9.204	Datastructure: Perforation	445
9.205	Datastructure: PeriodicEvent	446
9.206	Datastructure: PermeabilityBase	446
9.207	Datastructure: PhaseFieldDamageFEM	446
9.208	Datastructure: PhaseFieldFracture	446
9.209	Datastructure: PorousDruckerPrager	446
9.210	Datastructure: PorousElasticIsotropic	447
9.211	Datastructure: PorousElasticOrthotropic	447
9.212	Datastructure: PorousElasticTransverseIsotropic	447
9.213	Datastructure: PorousExtendedDruckerPrager	447
9.214	Datastructure: PressurePorosity	447
9.215	Datastructure: Problem	448
9.216	Datastructure: ProppantSlurryFluid	448
9.217	Datastructure: ProppantTransport	449
9.218	Datastructure: Python	449
9.219	Datastructure: Restart	449
9.220	Datastructure: Run	449
9.221	Datastructure: Silo	449
9.222	Datastructure: SinglePhaseFVM	449
9.223	Datastructure: SinglePhaseHybridFVM	450
9.224	Datastructure: SinglePhasePoromechanics	450
9.225	Datastructure: SinglePhasePoromechanicsEmbeddedFractures	450
9.226	Datastructure: SinglePhaseProppantFVM	451
9.227	Datastructure: SinglePhaseReservoir	451
9.228	Datastructure: SinglePhaseWell	451
9.229	Datastructure: SolidMechanicsEmbeddedFractures	452
9.230	Datastructure: SolidMechanicsLagrangianSSLE	453
9.231	Datastructure: SolidMechanics_LagrangianFEM	454
9.232	Datastructure: SoloEvent	455
9.233	Datastructure: Solvers	455
9.234	Datastructure: SourceFlux	456
9.235	Datastructure: StrainDependentPermeability	456
9.236	Datastructure: SurfaceElementRegion	456
9.237	Datastructure: SurfaceGenerator	458
9.238	Datastructure: SymbolicFunction	459
9.239	Datastructure: TableFunction	459
9.240	Datastructure: TableRelativePermeability	459
9.241	Datastructure: Tasks	459
9.242	Datastructure: ThickPlane	459
9.243	Datastructure: TimeHistory	460
9.244	Datastructure: Traction	460
9.245	Datastructure: TriaxialDriver	460
9.246	Datastructure: TwoPointFluxApproximation	460
9.247	Datastructure: VTK	460
9.248	Datastructure: VanGenuchtenBakerRelativePermeability	461
9.249	Datastructure: VanGenuchtenCapillaryPressure	461
9.250	Datastructure: WellControls	461

9.251 Datastructure: WellElementRegion	462
9.252 Datastructure: WellElementRegionuniqueSubRegion	463
9.253 Datastructure: cellBlocks	463
9.254 Datastructure: cellManager	464
9.255 Datastructure: commandLine	464
9.256 Datastructure: domain	465
9.257 Datastructure: edgeManager	466
9.258 Datastructure: elementRegionsGroup	467
9.259 Datastructure: elementSubRegions	467
9.260 Datastructure: embeddedSurfacesEdgeManager	467
9.261 Datastructure: embeddedSurfacesNodeManager	468
9.262 Datastructure: finiteVolumeStencils	468
9.263 Datastructure: lassen	468
9.264 Datastructure: neighborData	469
9.265 Datastructure: nodeManager	469
9.266 Datastructure: quartz	470
9.267 Datastructure: sets	470
9.268 Datastructure: wellElementSubRegion	470
10 Contributors	471
11 Publications	473
11.1 Preprints and Early-Views	475
11.2 2021	476
11.3 2020	477
11.4 2019	478
12 Acknowledgements	479
13 Indices and tables	481
Index	483

Welcome to our documentation pages!

GEOSX is a code framework focused on enabling streamlined development of physics simulations on high performance computing platforms. Our documentation is organized into several separate guides, given that different users will have different needs.

We recommend all users begin with the Quick Start guide, which will walk you through downloading and compiling the code. Application focused users may then want to explore our Tutorials, which provide an introduction to the basic capabilities of the code. More detailed descriptions of these capabilities can then be found in the User Guide.

For those interested in developing new capabilities in GEOSX, we provide a Developer Guide. The code itself is also documented inline using doxygen. The Build Guide contains more detailed information about third-party dependencies, the build system, and the continuous integration system. Finally, GEOSX has a self-documenting data structure. The Datastructure Index is an automatically generated list of all available input parameters and data structures in the code. This is a comprehensive resource, but probably not the place to start.

High quality documentation is a critical component of a successful code. If you have suggestions for improving the guides below, please post an issue on our [issue tracker](#).

The goal of this page is to get you started as quickly as possible using GEOSX. We will walk you through downloading the source, compiling the code, and testing the installation.

Before jumping to the installation process, we want to first address some frequently asked questions we get from new users. If you are itching to get started, feel free to jump ahead to the relevant sections.

1.1 Frequently Asked Questions

1.1.1 Does GEOSX have a graphical user interface?:

Given the focus on rapid development and HPC environments, GEOSX does not have a graphical user interface. This is consistent with many other high performance computing packages, but we recognize it can be a deal-breaker for certain users. For those who can get past this failing, we promise we still have a lot to offer. In a typical workflow, you will prepare an XML-based input file describing your problem. You may also prepare a mesh file containing geometric and property information describing, say, a reservoir you would like to simulate. There is no shortage of GUI tools that can help you in this model building stage. The resulting input deck is then consumed by GEOSX to run the simulation and produce results. This may be done in a terminal of your local machine or by submitting a job to a remote server. The resulting output files can then be visualized by any number of graphical visualization programs (typically [VisIt](#) or [paraview](#)). Thus, while GEOSX is GUI free, the typical workflow is not.

1.1.2 Do I need to be a code developer to use GEOSX?:

For the moment, most users will need to download and compile the code from source, which we readily admit this requires a certain level of development expertise. We try to make this process as easy as possible, and we are working on additional deployment options to make this process easier. Once installed, however, our goal is to make GEOSX accessible to developers and non-developers alike. Our target audience includes engineers and scientists who want to solve tough application problems, but could care less about the insides of the tool. For those of you who *are* interested in scientific computing, however, GEOSX is an open source project and we welcome external contributions.

1.1.3 What are the system requirements?:

GEOSX is primarily written in C++, with a focus on standards compliance and platform-to-platform portability. It is designed to run on everything from commodity laptops to the world's most powerful supercomputers. We regularly test the code across a variety of operating systems and compilers. Most of these operating systems are Linux/UNIX based (e.g. Ubuntu, CentOS, Mac OSX). We do have developers working in Windows environments, but they use a Virtual Machine or work within a docker image rather than directly in the Windows environment. In the instructions below, we assume you have access to fairly standard development tools. Using advanced features of GEOSX, like GPU-acceleration, will of course introduce additional hardware and software requirements.

1.1.4 Help, I get errors while trying to download/compile/run!:

Unfortunately, no set of instructions is foolproof. It is simply impossible to anticipate every system configuration or user. If you run into problems during the installation, we recommend the following five-step process:

1. Take a moment to relax, and then re-read the instructions carefully. Perhaps you overlooked a key step? Re-read the error message(s) closely. Modern compilation tools are often quite helpful in reporting exactly why things fail.
2. Type a few keywords from your error into a search engine. It is possible someone else out there has encountered your problem before, and a well-chosen keyword can often produce an instant solution. Note that when a compilation fails, you may get pages and pages of errors. Try to identify the *first* one to occur and fix that. One error will often trigger subsequent errors, and looking at the *last* error on the screen may not be so helpful.
3. If you encounter problems building one of the third-party libraries we depend on, check out their support pages. They may be able to help you more directly than we can.
4. Still stuck? Check out our [issues tracker](#), searching current or closed issues that may address your problem. Perhaps someone has had an identical issue, or something close. The issue tracker has a convenient search bar where you can search for relevant keywords. Remember to remove the default `is:open` keyword to search both open and closed issues.
5. If you have exhausted the options above, it is time to seek help from the developers. Post an issue on our issue tracker. Be specific, providing as much information as possible about your system setup and the error you are encountering. Please be patient in this process, as we may need to correspond a few times and ask you to run additional tests. Most of the time, users have a slightly unusual system configuration that we haven't encountered yet, such as an older version of a particular library. Other times there is a legitimate bug in GEOSX to be addressed. Take pride in the fact that you may be saving the next user from wasted time and frustration.

1.2 Repository Organization

The source for GEOSX and related tools are hosted on [Github](#). We use [Git workflows](#) to version control our code and manage the entire development process. On Github, we have a [GEOSX Organization](#) that hosts several related repositories.

You should sign up for a free Github account, particularly if you are interested in posting issues to our issue tracker and communicating with the developers. The main repository of interest is obviously GEOSX itself: [GEOSX](#)

We also rely on two types of dependencies: first-party and third-party. First-party dependencies are projects directly associated with the GEOSX effort, but kept in separate repositories because they form stand-alone tools. For example, there is a geologic mesh handling package called [PAMELA](#) and an equation-of-state package called [PVTPackage](#). These packages are handled as [Git Submodules](#), which provides a transparent way of coordinating multiple code development projects. Most users will never have to worry that these modules are in fact separate projects from GEOSX.

We also rely on several open-source Third-Party Libraries (TPLs) (see [thirdPartyLibs](#)). These are well-respected projects developed externally to GEOSX. We have found, however, that many compilation issues stem from version incompatibilities between different packages. To address this, we provide a mirror of these TPLs, with version combinations we know play nicely together. We also provide a build script that conveniently and consistently builds those dependencies.

Our build system will automatically use the mirror package versions by default. You are welcome to tune your configuration, however, to point to different versions installed on your system. If you work on an HPC platform, for example, common packages may already be available and optimized for platform hardware. For new users, however, it may be safer to begin with the TPL mirror.

Note: If you are working on an HPC platform with several other GEOSX users, we often compile the TPLs in a shared location so individual users don't have to waste their storage quota. Inquire with your institution's point-of-contact whether this option already exists. For all LLNL systems, the answer is yes.

Finally, there are also several private repositories only accessible to the core development team, which we use for behind-the-scene testing and maintenance of the code.

1.3 Username and Authentication

New users should sign up for a free [Github account](#).

If you intend to develop in the GEOSX codebase, you may benefit from setting up your git credentials (see [Git Workflow](#)).

1.4 Download

It is possible to directly download the source code as a zip file. We strongly suggest, however, that users don't rely on this option. Instead, most users should use Git to either *clone* or *fork* the repository. This makes it much easier to stay up to date with the latest releases and bug fixes. If you are not familiar with the basics of Git, [here is a helpful resource](#) to get you started.

The tutorial here assumes you will use a https clone with no specific credentials. Using an ssh connection pattern requires a very slight modification. See the **Additional Notes** at the end of this section for details.

If you do not already have Git installed on your system, you will need to install it. We recommend using a relatively recent version of Git, as there have been some notable improvements over the past few years. You can check if Git is already available by opening a terminal and typing

```
git --version
```

You'll also need the `git-lfs` large file extension.

The first task is to clone the GEOSX and `thirdPartyLibs` repositories. If you do not tell it otherwise, the build system will expect the GEOSX and `thirdPartyLibs` to be parallel to each other in the directory structure. For example,

```
codes/
├── GEOSX/
└── thirdPartyLibs/
```

where the toplevel `codes` directory can be re-named and located wherever you like. It is possible to customize the build system to expect a different structure, but for now let us assume you take the simplest approach.

First, using a terminal, create the `codes` directory wherever you like.

```
cd /insert/your/desired/path/
mkdir codes
cd codes
```

Inside this directory, we can clone the GEOSX repository. We will also use some Git commands to initialize and download the submodules (e.g. PAMELA). Note that most users will not have access to our integrated tests repository, and so we “deinit” (deactivate) this submodule. Developers who will be working with the integratedTests repository should skip this line.

```
git clone https://github.com/GEOSX/GEOSX.git
cd GEOSX
git lfs install
git submodule init
git submodule deinit integratedTests
git submodule update
cd ..
```

If all goes well, you should have a complete copy of the GEOSX source at this point. The most common errors people encounter here have to do with Github not recognizing their authentication settings and/or repository permissions. See the previous section for tips on ensuring your SSH is working properly.

Note: The integratedTests submodule is not publicly available, with access limited to the core development team. This may cause the `git submodule update` command to fail if you forget the `git submodule deinit integratedTests` step above. This submodule is not required for building GEOSX. If you see an error message here, however, you may need to initialize and update the submodules manually:

```
cd GEOSX
git submodule update --init src/coreComponents/LvArray
git submodule update --init src/coreComponents/fileIO/coupling/hdf5_interface
git submodule update --init src/externalComponents/PAMELA
git submodule update --init src/externalComponents/PVTPackage
cd ..
```

Once we have grabbed GEOSX, we do the same for the thirdPartyLibs repository. From the `codes` directory, type

```
git clone https://github.com/GEOSX/thirdPartyLibs.git
cd thirdPartyLibs
git lfs install
git pull
git submodule init
git submodule update
cd ..
```

Again, if all goes well you should now have a copy of all necessary TPL packages.

Additional Notes:

#. `git-lfs` may not function properly (or may be very slow) if your version of `git` and `git-lfs` are not current. If you are using an older version, you may need to add `git lfs pull` after `git pull` in the above procedures.

#. You can adapt the commands if you use an `ssh` connection instead. The `clone https://github.com/GEOSX/GEOSX.git` becomes `git clone git@github.com:GEOSX/GEOSX.git`. You may also be willing to insert your credentials in the command line (less secure) `git clone https://${USER}:${TOKEN}@github.com/GEOSX/GEOSX.git`.

1.5 Configuration

At a minimum, you will need a relatively recent compiler suite installed on your system (e.g. [GCC](#), [Clang](#)) as well as [CMake](#). If you want to run jobs using MPI-based parallelism, you will also need an MPI implementation (e.g. [OpenMPI](#), [MVAPICH](#)). Note that GEOSX supports a variety of parallel computing models, depending on the hardware and software environment. Advanced users are referred to the [Build Guide](#) for a discussion of the available configuration options.

Before beginning, it is a good idea to have a clear idea of the flavor and version of the build tools you are using. If something goes wrong, the first thing the support team will ask you for is this information.

```
cpp --version
mpic++ --version
cmake --version
```

Here, you may need to replace `cpp` with the full path to the C++ compiler you would like to use, depending on how your path and any aliases are configured.

GEOSX compilations are driven by a `cmake host-config` file, which tells the build system about the compilers you are using, where various packages reside, and what options you want to enable. We have created a number of default `hostconfig` files for common systems. You should browse them to see if any are close to your needs:

```
cd GEOSX/host-configs
```

We maintain `host configs` (ending in `.cmake`) for HPC systems at various institutions, as well as ones for common personal systems. If you cannot find one that matches your needs, we suggest beginning with one of the shorter ones and modifying as needed. A typical one may look like:

```
# file: your-platform.cmake

# detect host and name the configuration file
site_name(HOST_NAME)
set(CONFIG_NAME "your-platform" CACHE PATH "")
message( "CONFIG_NAME = ${CONFIG_NAME}" )

# set paths to C, C++, and Fortran compilers. Note that while GEOSX does not contain
↳ any Fortran code,
# some of the third-party libraries do contain Fortran code. Thus a Fortran compiler
↳ must be specified.
set(CMAKE_C_COMPILER "/usr/bin/clang" CACHE PATH "")
set(CMAKE_CXX_COMPILER "/usr/bin/clang++" CACHE PATH "")
set(CMAKE_Fortran_COMPILER "/usr/local/bin/gfortran" CACHE PATH "")
set(ENABLE_Fortran OFF CACHE BOOL "" FORCE)

# enable MPI and set paths to compilers and executable.
# Note that the MPI compilers are wrappers around standard serial compilers.
# Therefore, the MPI compilers must wrap the appropriate serial compilers specified
# in CMAKE_C_COMPILER, CMAKE_CXX_COMPILER, and CMAKE_Fortran_COMPILER.
set(ENABLE_MPI ON CACHE PATH "")
set(MPI_C_COMPILER "/usr/local/bin/mpicc" CACHE PATH "")
set(MPI_CXX_COMPILER "/usr/local/bin/mpicxx" CACHE PATH "")
set(MPI_Fortran_COMPILER "/usr/local/bin/mpifort" CACHE PATH "")
set(MPIEXEC "/usr/local/bin/mpirun" CACHE PATH "")

# disable CUDA and OpenMP
set(CUDA_ENABLED "OFF" CACHE PATH "" FORCE)
set(ENABLE_OPENMP "OFF" CACHE PATH "" FORCE)
```

(continues on next page)

(continued from previous page)

```
# enable PAMELA and PVTPackage
set(ENABLE_PAMELA ON CACHE BOOL "" FORCE)
set(ENABLE_PVTPackage ON CACHE BOOL "" FORCE)

# enable tests
set(ENABLE_GTEST_DEATH_TESTS ON CACHE BOOL "" FORCE )
```

The various `set()` commands are used to set environment variables that control the build. You will see in the above example that we set the C++ compiler to `/user/bin/clang++` and so forth. We also disable CUDA and OpenMP, but enable PAMELA and PVTPackage. The final line is related to our unit test suite. See the [Build Guide](#) for more details on available options.

Note: If you develop a new `host-config` for a particular platform that may be useful for other users, please consider sharing it with the developer team.

1.6 Compilation

We will begin by compiling the TPLs, followed by the main code. If you work on an HPC system with other GEOSX developers, check with them to see if the TPLs have already been compiled in a shared directory. If this is the case, you can skip ahead to just compiling the main code. If you are working on your own machine, you will need to compile both.

We strongly suggest that GEOSX and TPLs be built with the same `hostconfig` file. Below, we assume that you keep it in, say, `GEOSX/host-configs/your-platform.cmake`, but this is up to you.

We begin with the third-party libraries, and use a `python config-build.py` script to configure and build all of the TPLs. Note that we will request a Release build type, which will enable various optimizations. The other option is a Debug build, which allows for debugging but will be much slower in production mode. The TPLS will then be built in a build directory named consistently with your `hostconfig` file.

```
cd thirdPartyLibs
python scripts/config-build.py -hc ../GEOSX/host-configs/your-platform.cmake -bt_
↪Release
cd build-your-platform-release
make
```

Note that building all of the TPLs can take quite a while, so you may want to go get a cup of coffee at this point. Also note that you should *not* use a `parallel make -j N` command to try and speed up the build time.

The next step is to compile the main code. Again, the `config-build.py` sets up `cmake` for you, so the process is very similar.

```
cd ../../GEOSX
python scripts/config-build.py -hc host-configs/your-platform.cmake -bt Release
cd build-your-platform-release
make -j4
make install
```

Here, the `parallel make -j 4` will use four processes for compilation, which can substantially speed up the build if you have a multi-processor machine. You can adjust this value to match the number of processors available on your machine. The `make install` command then installs GEOSX to a default location unless otherwise specified. If all goes well, a `geosx` executable should now be available:

```
GEOSX/install-your-platform-release/bin/geosx
```

1.7 Running

We can do a quick check that the geosx executable is working properly by calling the executable with our help flag

```
./bin/geosx --help
```

This should print out a brief summary of the available command line arguments:

```
USAGE: geosx -i input.xml [options]

Options:
-?, --help
-i, --input,          Input xml filename (required)
-r, --restart,        Target restart filename
-x, --x-partitions,   Number of partitions in the x-direction
-y, --y-partitions,   Number of partitions in the y-direction
-z, --z-partitions,   Number of partitions in the z-direction
-s, --schema,         Name of the output schema
-b, --use-nonblocking, Use non-blocking MPI communication
-n, --name,           Name of the problem, used for output
-s, --suppress-pinned Suppress usage of pinned memory for MPI communication buffers
-o, --output,         Directory to put the output files
-t, --timers,         String specifying the type of timer output.
An input xml must be specified!
```

Obviously this doesn't do much interesting, but it will at least confirm that the executable runs. In typical usage, an input XML must be provided describing the problem to be run, e.g.

```
./bin/geosx -i your-problem.xml
```

In a parallel setting, the command might look something like

```
mpirun -np 8 ./bin/geosx -i your-problem.xml -x 2 -y 2 -z 2
```

Note that we provide a series of *Tutorials* to walk you through the actual usage of the code, with several input examples. Once you are comfortable the build is working properly, we suggest new users start working through these tutorials.

1.8 Testing

It is wise to run our unit test suite as an additional check that everything is working properly. You can run them in the build folder you just created.

```
cd GEOSX/build-your-platform-release
ctest -V
```

This will run a large suite of simple tests that check various components of the code. If you have access, you may also consider running the integrated tests. Please refer to *Integrated Tests* for further information.

Note: If *all* of the unit tests fail, there is likely something wrong with your installation. Refer to the FAQs above for how best to proceed in this situation. If only a few tests fail, it is possible that your platform configuration has exposed

some issue that our existing platform tests do not catch. If you suspect this is the case, please consider posting an issue to our issue tracker (after first checking whether other users have encountered a similar issue).


The easiest way to learn to use GEOSX is through worked examples. Here, we have included tutorials showing how to run some common problems. After working through these examples, you should have a good understanding of how to set up and solve your own models.

Note that these tutorials are intended to be followed in sequence, as each step introduces a few new skills. Most of the tutorial models are also quite small, so that large computational resources are not required.

2.1 Tutorial 1: First Steps

Context

In this tutorial, we use a single-phase flow solver (see [Singlephase Flow Solver](#)) to solve for pressure propagation on a 10x10x10 cube mesh with anisotropic permeability values. The pressure source is the lowest-left corner element, and the pressure sink sits at the opposite top corner.



Objectives

At the end of this tutorial you will know:

- the basic structure of XML input files used by GEOSX,
- how to run GEOSX on a simple case requiring no external input files,
- the basic syntax of a solver block for single-phase problems,
- how to control output and visualize results.

Input file

GEOSX runs by reading user input information from one or more XML files. For this tutorial, we only need a single GEOSX input file located at:

```
src/coreComponents/physicsSolvers/fluidFlow/integratedTests/singlePhaseFlow/3D_  
↪ 10x10x10_compressible.xml
```

Running GEOSX

If our XML input file is called `my_input.xml`, GEOSX runs this file by executing:

```
/path/to/geosx -i /path/to/my_input.xml
```

The `-i` flag indicates the path to the XML input file. To get help on what other command line input flags GEOSX supports, run `geosx --help`.

Input file structure

XML files store information in a tree-like structure using nested blocks of information called *elements*. In GEOSX, the root of this tree structure is the element called `Problem`. All elements in an XML file are defined by an opening *tag* (`<ElementName>`) and end by a corresponding closing tag (`</ElementName>`). Elements can have properties defined as *attributes* with `key="value"` pairs. A typical GEOSX input file contains the following tags:

1. *Solver*
2. *Mesh*

3. *Geometry*
4. *Events*
5. *NumericalMethods*
6. *ElementRegions*
7. *Constitutive*
8. *FieldSpecifications*
9. *Outputs*

XML validation tools

If you have not already done so, please use or enable an XML validation tool (see **User Guide/Input Files/Input Validation**). Such tools will help you identify common issues that may occur when working with XML files.

Note: Common errors come from the fact that XML is case-sensitive, and all opened tags must be properly closed.

2.1.1 Single-phase solver

GEOSX is a multiphysics simulator. To find the solution to different physical problems such as diffusion or mechanical deformation, GEOSX uses one or more physics solvers. The `Solvers` tag is used to define and parameterize these solvers. Different combinations of solvers can be applied in different regions of the domain at different moments of the simulation.

In this first example, we use one type of solver in the entire domain and for the entire duration of the simulation. The solver we are specifying here is a single-phase flow solver. In GEOSX, such a solver is created using a `SinglePhaseFVM` element. This type of solver is one among several cell-centered single-phase finite volume methods.

The XML block used to define this single-phase finite volume solver is shown here:

```
<Solvers>
  <SinglePhaseFVM
    name="SinglePhaseFlow"
    logLevel="1"
    discretization="singlePhaseTPFA"
    fluidNames="{ water }"
    solidNames="{ rock }"
    permeabilityNames="{ rockPerm }"
    targetRegions="{ mainRegion }">
    <NonlinearSolverParameters
      newtonTol="1.0e-6"
      newtonMaxIter="8"/>
    <LinearSolverParameters
      solverType="gmres"
      krylovTol="1.0e-10"/>
  </SinglePhaseFVM>
</Solvers>
```

Each type of solver has a specific set of parameters that are required and some parameters that are optional. Optional values are usually set with sensible default values.

name

First, we register a solver of type `SinglePhaseFVM` with a user-chosen name, here `SinglePhaseFlow`. This unique user-defined name can be almost anything. However, some symbols are known to cause issues in names : avoid commas, slashes, curly braces. GEOSX is case-sensitive: it makes a distinction between two `SinglePhaseFVM` solvers called `mySolver` and `MySolver`. Giving elements a name is a common practice in GEOSX: users need to give unique identifiers to objects they define. That name is the handle to this instance of a solver class.

logLevel

Then, we set a solver-specific level of console logging (`logLevel` set to 1 here). Notice that the value (1) is between double-quotes. This is a general convention for all attributes: we write `key="value"` regardless of the value type (integers, strings, lists, etc.).

For `logLevel`, higher values lead to more console output or intermediate results saved to files. When debugging, higher `logLevel` values is often convenient. In production runs, you may want to suppress most console output.

discretization

For solvers of the `SinglePhaseFVM` family, one required attribute is a discretization scheme. Here, we use a Two-Point Flux Approximation (TPFA) finite volume discretization scheme called `singlePhaseTPFA`. To know the list of admissible values of an attribute, please see GEOSX's XML schema. This discretization type must know how to find permeability values that it uses internally to compute transmissibilities. The `permeabilityNames` attribute tells the solver the user-defined name (the *handle*) of the permeability values that will be defined elsewhere in the input file. Note that the order of attributes inside an element is not important.

fluidNames, solidNames, targetRegions

Here, we specify a collection of fluids, rocks, and target regions of the mesh on which the solver will apply. Curly brackets are used in GEOSX inputs to indicate collections of values (sets or lists). The curly brackets used here are necessary, even if the collection contains a single value. Commas are used to separate members of a set.

Nested elements

Finally, note that other XML elements can be nested inside the `Solvers` element. Here, we use specific XML elements to set values for numerical tolerances. The solver stops when numerical residuals are smaller than the specified tolerances (convergence is achieved) or when the maximum number of iterations allowed is exceeded (convergence not achieved).

2.1.2 Mesh

To solve this problem, we need to define a mesh for our numerical calculations. This is the role of the **Mesh** element.

There are two approaches to specifying meshes in GEOSX: internal or external.

- The external approach allows to import mesh files created outside GEOSX, such as a corner-point grid or an unstructured grid representing complex shapes and structures.
- The internal approach uses GEOSX's built-in capability to create simple meshes from a small number of parameters. It does not require any external file information. The geometric complexity of internal meshes is limited, but many practical problems can be solved on such simple grids.

In this tutorial, to keep things self-contained, we use the internal mesh generator. We parameterize it with the **InternalMesh** element.

```
<Mesh>
  <InternalMesh
    name="mesh"
    elementTypes="{ C3D8 }"
    xCoords="{ 0, 10 }"
    yCoords="{ 0, 10 }"
```

(continues on next page)

(continued from previous page)

```
zCoords="{ 0, 10 }"  
nx="{ 10 }"  
ny="{ 10 }"  
nz="{ 10 }"  
cellBlockNames="{ cellBlock }"/>  
</Mesh>
```

name

Just like for solvers, we register the `InternalMesh` element using a unique **name** attribute. Here the `InternalMesh` object is instantiated with the name `mesh`.


elementTypes

We specify the collection of elements types that this mesh contains. Tetrahedra, hexahedra, wedges, prisms are examples of element types. If a mesh contains different types of elements (a hybrid mesh), we should indicate this here by listing all unique types of elements in curly brackets. Keeping things simple, our element collection has only one type of element: a `C3D8` type representing a hexahedral element (linear 8-node brick).

A mesh can contain several geometrical types of elements. For numerical convenience, elements are aggregated by types into `cellBlocks`. Here, we only linear 8-node brick elements, so the entire domain is one object called `cellBlock`.

xCoords, yCoords, zCoords, nx, ny, nz

This specifies the spatial arrangement of the mesh elements. The mesh defined here goes from coordinate $x=0$ to $x=10$ in the x -direction, with $nx=10$ subdivisions along this segment. The same is true for the y -dimension and the z -dimension. Our mesh is a cube of $10 \times 10 \times 10 = 1,000$ elements with a bounding box defined by corner coordinates $(0,0,0)$ and $(10,10,10)$.



2.1.3 Geometry

The `Geometry` tag allows users to capture subregions of a mesh and assign them a unique name. Here, we name two `Box` elements, one for the location of the `source` and one for the `sink`. Pressure values are assigned to these named regions elsewhere in the input file.

The pressure source is the element in the (0,0,0) corner of the domain, and the sink is the element in the (10,10,10) corner.

For an element to be inside a geometric region, it must have all its vertices strictly inside that region. Consequently, we need to extend the geometry limits a small amount beyond the actual coordinates of the elements to catch all vertices. Here, we use a safety padding of 0.01.

```
<Geometry>
  <Box
    name="source"
    xMin="{ -0.01, -0.01, -0.01 }"
    xMax="{ 1.01, 1.01, 1.01 }"/>
  <Box
```

(continues on next page)


(continued from previous page)

```

name="sink"
xMin="{ 8.99, 8.99, 8.99 }"
xMax="{ 10.01, 10.01, 10.01 }"/>
</Geometry>

```

There are several methods to achieve similar conditions (Dirichlet boundary condition on faces, etc.). The `Box` defined here is one of the simplest approaches.



2.1.4 Events

In GEOSX, we call `Events` anything that happens at a set time or frequency. Events are a central element for time-stepping in GEOSX, and a dedicated section just for events is necessary to give them the treatment they deserve.

For now, we focus on three simple events: the time at which we wish the simulation to end (`maxTime`), the times at which we want the solver to perform updates, and the times we wish to have simulation output values reported.

In GEOSX, all times are specified in **seconds**, so here `maxTime=5000.0` means that the simulation will run from time 0 to time 5,000 seconds.

If we focus on the `PeriodicEvent` elements, we see :

1. A **periodic solver** application: this event is named `solverApplications`. With the attribute `forceDt=20`, it tells the solver to compute results at 20-second time intervals. We know what this event does by looking at its `target` attribute: here, from time 0 to `maxTime` and with a forced time step of 20 seconds, we instruct GEOSX to call the solver registered as `SinglePhaseFlow`. Note the hierarchical structure of the target formulation, using `'/'` to indicate a specific named instance (`SinglePhaseFlow`) of an element (`Solvers`). If the solver needs to take smaller time steps, it is allowed to do so, but it will have to compute results for every 20-second increment between time zero and `maxTime` regardless of possible intermediate time steps.
2. An **output event**: this event is used for reporting purposes and instructs GEOSX to write out results at specific frequencies. Here, we need to see results at every 100-second increment. This event triggers a full application of solvers, even if solvers were not summoned by the previous event. In other words, an output event will force an application of solvers, possibly in addition to the periodic events requested directly.

```
<Events
  maxTime="5000.0">
  <PeriodicEvent
    name="solverApplications"
    forceDt="20.0"
    target="/Solvers/SinglePhaseFlow"/>
  <PeriodicEvent
    name="outputs"
    timeFrequency="100.0"
    target="/Outputs/siloOutput"/>
</Events>
```

2.1.5 Numerical methods

GEOSX comes with several useful numerical methods. In the `Solvers` elements, for instance, we had specified to use a two-point flux approximation as discretization scheme for the finite volume single-phase solver. Now to use this scheme, we need to supply more details in the `NumericalMethods` element.

```
<NumericalMethods>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="singlePhaseTPFA"
      fieldName="pressure"
      coefficientName="permeability"
      coefficientModelNames="{ rockPerm }"/>
    </FiniteVolume>
  </NumericalMethods>
```

The `fieldName` attribute specifies which property will be used for flux computations, and also specifies that for Dirichlet boundary conditions, the pressure value at the element face is used. The `coefficientName` attribute is used for the stencil transmissibility computations.

Note that in GEOSX, there is a difference between physics solvers and numerical methods. Their parameterizations are thus independent. We can have multiple solvers using the same numerical scheme but with different tolerances, for instance.

2.1.6 Regions

In GEOSX, `ElementsRegions` are used to attach material properties to regions of elements. Here, we use only one **CellElementRegion** to represent the entire domain (user name: `mainRegion`). It contains all the blocks called `cellBlock` defined in the mesh section. We specify the materials contained in that region using a `materialList`.

Several materials coexist in `cellBlock`, and we list them using their user-defined names: `water`, `rockPorosity`, and `rockPerm`, etc. What these names mean, and the physical properties that they are attached to are defined next.

```
<ElementRegions>
  <CellElementRegion
    name="mainRegion"
    cellBlocks="{ cellBlock }"
    materialList="{ water, rock, rockPerm, rockPorosity, nullSolid }"/>
  </ElementRegions>
```

2.1.7 Constitutive models

The Constitutive element attaches physical properties to all materials contained in the domain.

The physical properties of the materials defined as `water`, `rockPorosity`, and `rockPerm` are provided here, each material being derived from a different material type: `CompressibleSinglePhaseFluid` for the water, `PressurePorosity` for the rock porosity, and `ConstantPermeability` for rock permeability. The list of attributes differs between these constitutive materials.

```
<Constitutive>
  <CompressibleSinglePhaseFluid
    name="water"
    defaultDensity="1000"
    defaultViscosity="0.001"
    referencePressure="0.0"
    compressibility="5e-10"
    viscosibility="0.0"/>

  <CompressibleSolidConstantPermeability
    name="rock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"/>

  <NullModel
    name="nullSolid"/>

  <PressurePorosity
    name="rockPorosity"
    defaultReferencePorosity="0.05"
    referencePressure="0.0"
    compressibility="1.0e-9"/>

  <ConstantPermeability
    name="rockPerm"
    permeabilityComponents="{ 1.0e-12, 1.0e-12, 1.0e-15 }"/>
</Constitutive>
```

The names `water`, `rockPorosity` and `rockPerm` are defined by the user as handles to specific instances of physical materials. GEOSX uses S.I. units throughout, not field units. Pressures, for instance, are in Pascal, not psia. The x- and y-permeability are set to $1.0\text{e-}12\text{ m}^2$ corresponding to approximately 1 Darcy.

We have used the handles `water`, `rockPorosity` and `rockPerm` in the input file in the `ElementRegions` section of the XML file, before the registration of these materials took place here, in `Constitutive` element.

Note: This highlights an important aspect of using XML in GEOSX: the order in which objects are registered and

used in the XML file is not important.

2.1.8 Defining properties

In the `FieldSpecifications` section, properties such as source and sink pressures are set. GEOSX offers a lot of flexibility to specify field values through space and time.

Spatially, in GEOSX, all field specifications are associated to a target object on which the field values are mounted. This allows for a lot of freedom in defining fields: for instance, one can have volume property values attached to a subset of volume elements of the mesh, or surface properties attached to faces of a subset of elements.


For each `FieldSpecification`, we specify a name, a `fieldName` (this name is used by solvers or numerical methods), an `objectPath`, `setNames` and a `scale`. The `ObjectPath` is important and it reflects the internal class hierarchy of the code. Here, for the `fieldName` `pressure`, we assign the value defined by `scale` (5e6 Pascal) to one of the `ElementRegions` (class) called `mainRegions` (instance). More specifically, we target the `elementSubRegions` called `cellBlock` (this contains all the C3D8 elements, effectively all the domain). The `setNames` allows to use the elements defined in `Geometry`, or use everything in the object path (using the `all`).

```
<FieldSpecifications>
  <FieldSpecification
    name="initialPressure"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/mainRegion/elementSubRegions/cellBlock"
    fieldName="pressure"
    scale="5e6"/>

  <FieldSpecification
    name="sourceTerm"
    objectPath="ElementRegions/mainRegion/elementSubRegions/cellBlock"
    fieldName="pressure"
    scale="1e7"
    setNames="{ source }"/>

  <FieldSpecification
    name="sinkTerm"
    objectPath="ElementRegions/mainRegion/elementSubRegions/cellBlock"
    fieldName="pressure"
    scale="0.0"
    setNames="{ sink }"/>
</FieldSpecifications>
```

The image below shows the pressures after the very first time step, with the domain initialized at 5 MPa, the sink at 0 MPa on the top right, and the source in the lower left corner at 10 MPa.



2.1.9 Output

In order to retrieve results from a simulation, we need to instantiate one or multiple `Outputs`.

Here, we define a single object of type `Silo`. `Silo` is a library and a format for reading and writing a wide variety of scientific data. Data in `Silo` format can be read by `VisIt`.

This `Silo` output object is called `siloOutput`. We had referred to this object already in the `Events` section: it was the target of a periodic event named `outputs`. You can verify that the `Events` section is using this object as a target. It does so by pointing to `/Outputs/siloOutput`.

```
<Outputs>
  <Silo
    name="siloOutput"/>
</Outputs>
```

GEOSX currently supports outputs that are readable by `VisIt` and Kitware's `Paraview`, as well as other visualization tools. In this example, we only request a `Silo` format compatible with `VisIt`.

All elements are now in place to run GEOSX.

2.1.10 Running GEOSX

The command to run GEOSX is

```
path/to/geosx -i path/to/this/xml_file.xml
```

Note that all paths for files included in the XML file are relative to this XML file.

While running GEOSX, it logs status information on the console output with a verbosity that is controlled at the object level, and that can be changed using the `logLevel` flag.

The first few lines appearing to the console are indicating that the XML elements are read and registered correctly:

```
Adding Solver of type SinglePhaseFVM, named SinglePhaseFlow
Adding Mesh: InternalMesh, mesh
Adding Geometric Object: Box, source
Adding Geometric Object: Box, sink
Adding Event: PeriodicEvent, solverApplications
Adding Event: PeriodicEvent, outputs
Adding Output: Silo, siloOutput
Adding Object CellElementRegion named mainRegion from ObjectManager::Catalog.
  mainRegion/cellBlock/water is allocated with 1 quadrature points.
  mainRegion/cellBlock/rock is allocated with 1 quadrature points.
  mainRegion/cellBlock/rockPerm is allocated with 1 quadrature points.
  mainRegion/cellBlock/rockPorosity is allocated with 1 quadrature points.
  mainRegion/cellBlock/nullSolid is allocated with 1 quadrature points.
```

Then, we go into the execution of the simulation itself:

```
Time: 0s, dt:20s, Cycle: 0
  Attempt: 0, NewtonIter: 0
  ( R ) = ( 5.65e+00 ) ;
  Attempt: 0, NewtonIter: 1
  ( R ) = ( 2.07e-04 ) ;
  Last LinSolve(iter,res) = ( 63, 8.96e-11 ) ;
  Attempt: 0, NewtonIter: 2
  ( R ) = ( 9.86e-11 ) ;
  Last LinSolve(iter,res) = ( 70, 4.07e-11 ) ;
```

Each time iteration at every 20s interval is logged to console, until the end of the simulation at `maxTime=5000`:

```
Time: 4980s, dt:20s, Cycle: 249
  Attempt: 0, NewtonIter: 0
  ( R ) = ( 4.74e-09 ) ;
  Attempt: 0, NewtonIter: 1
  ( R ) = ( 2.05e-14 ) ;
  Last LinSolve(iter,res) = ( 67, 5.61e-11 ) ;
SinglePhaseFlow: Newton solver converged in less than 4 iterations, time-step
→required will be doubled.
Cleaning up events
Umpire          HOST sum across ranks: 14.8 MB
Umpire          HOST          rank max: 14.8 MB
total time      5.658s
initialization time 0.147s
run time        3.289s
```


Note that by default a direct linear solver is used, so the `LinSolve` information is not so meaningful. Information on run times, initialization times, and maximum amounts of memory (high water mark) are given at the end of the simulation, if successful.

Congratulations on completing this first run!

2.1.11 Visualization

Here, we have requested results to be written in Silo, a format compatible with [VisIt](#). To visualize results, open VisIt and directly load the database of simulation output files.

After a few time step, pressure between the source and sink are in equilibrium, as shown on the representation below.



2.1.12 To go further

Feedback on this tutorial

This concludes the single-phase internal mesh tutorial. For any feedback on this tutorial, please submit a [GitHub issue](#) on the project's [GitHub](#) page.

For more details

- More on single-phase flow solvers, please see [Singlephase Flow Solver](#).
- More on meshes, please see [Meshes](#).

- More on events, please see [Event Management](#).

2.2 Tutorial 2: External Meshes

Context

In this tutorial, we use a simple single-phase flow solver (see [Singlephase Flow Solver](#)) to solve for pressure propagation on a mesh that is imported into GEOSX. The main goal of this tutorial is to learn how to work with external meshes, and to learn how easy it is to swap meshes on the same physical problem in GEOSX. This makes GEOSX a powerful tool to solve real field applications with complex geometries and perform assessments of mesh geometry and resolution effects.

Objectives

At the end of this tutorial you will know:

- the syntax and format of input meshes,
- how to input external files into a GEOSX input XML file,
- how to run the same physical problem with two different meshes,
- how to use and visualize hexahedral and tetrahedral meshes.

Input Files

This tutorial uses an XML file containing the main input for GEOSX and a separate file with all the mesh information. As we will see later, the main XML file points to the external mesh file with an `include` statement. The XML input file for this test case is located at:

```
src/coreComponents/physicsSolvers/fluidFlow/integratedTests/singlePhaseFlow/pamela_  
→test/3D_10x10x10_compressible_pamela_hex_gravity.xml
```

The mesh file format used in this tutorial is called **MSH**. This format is a standard scientific meshing format not specific to GEOSX. It is maintained as the native format of the meshing tool **Gmsh**. MSH is designed for unstructured meshes and contains a compact and complete representation of the mesh geometry and of its properties. The mesh file used here is human-readable ASCII. It contains a list of nodes with their (x,y,z) coordinates, and a list of elements that are constructed from these nodes.

2.2.1 Hexahedral elements

In the first part of the tutorial, we will run flow simulations on a mesh made of hexahedral elements. These types of elements are used in classical cartesian grids (sugar cubes) or corner-point grids or pillar grids.

Brief discussion about hexahedral meshes in GEOSX

Although closely related, the hexahedral grids that GEOSX can process are slightly different than either structured grid or corner-point grids. The differences are worth pointing out here. In GEOSX:

- **hexahedra can have irregular shapes:** no pillars are needed and vertices can be anywhere in space. This is useful for grids that turn, fold, or are heavily bent. Hexahedral blocks should nevertheless not be deprecated and have 8 distinct vertices. Some tolerance exists for deprecation to wedges or prisms in some solvers (finite element solvers), but it is best to avoid such situations and label elements according to their actual shape. Butterfly cells, flat cells, negative or zero volume cells will cause problems.


- **the mesh needs to be conformal:** in 3D, this means that neighboring grid blocks have to share exactly a complete face. Note that corner-point grids do not have this requirement and neighboring blocks can be offset. When importing grids from commonly-used ge modeling packages, this is an important consideration. This problem is solved by splitting shifted grid blocks to restore conformity. While it may seem convenient to be able to have offset grid blocks at first, the advantages of conformal grids used in GEOSX are worth the extra meshing effort: by using conformal grids, GEOSX can run finite element and finite volume simulations on the same mesh without problems, going seamlessly from one numerical method to the other. This is key to enabling multiphysics simulation.
- **there is no assumption of overall structure:** GEOSX does not need to know a number of block in the X, Y, Z direction (no NX, NY, NZ) and does not assume that the mesh is a full cartesian domain that the interesting parts of the reservoir must be carved out from. Blocks are numbered by indices that assume nothing about spatial positioning and there is no concept of (i,j,k). This approach also implies that no “masks” are needed to remove inactive or dead cells, as often done in cartesian grids to get the actual reservoir contours from a bounding box, and here we only need to specify grid blocks that are active. For performance and flexibility, this lean approach to meshes is important.

Importing an external mesh with PAMELA

In this first part of the tutorial, we use a hexahedral mesh provided with GEOSX. This hexahedral mesh is strictly identical to the grid used in the first tutorial (*Tutorial 1: First Steps*), but instead of using the internal grid generator GEOSX, we specify it with spatial node coordinates in MSH format.

The process by which grids are imported into GEOSX is worth explaining. To import external grid into GEOSX, we use an external component (submodule) called **PAMELA**. PAMELA (Parallel Meshing Library) was developed as a stand-alone utility to import grids in multiple formats and write them into memory for GEOSX. Although PAMELA is not necessary to run GEOSX (the internal grid generator of GEOSX has plenty of interesting features), you need PAMELA if you want to import external grids.

So here, our mesh consists of a simple sugar-cube stack of size 10x10x10. We inject fluid from one vertical face of a cube (the face corresponding to $x=0$), and we let the pressure equilibrate in the closed domain. The displacement is a single-phase, compressible fluid subject to gravity forces, so we expect the pressure to be constant on the injection face, and to be close to hydrostatic on the opposite plane ($x=10$). We use GEOSX to compute the pressure inside each grid block over a period of time of 100 seconds.



To see how to import such a mesh, we inspect the following XML file:

```
src/coreComponents/physicsSolvers/integratedTests/singlePhaseFlow/pamela_test/3D_
↪10x10x10_compressible_pamela_hex_gravity.xml
```

In the XML Mesh tag, instead of an InternalMesh tag, we have a PAMELAMeshGenerator tag. We see that a file called cube_10x10x10_hex.msh is imported using PAMELA, and this object is instantiated with a user-defined name value. The file here contains geometric information in **MSH** format (it can also contain properties, as we will see in the next tutorial).

```
<Mesh>
  <PAMELAMeshGenerator
    name="CubeHex"
    file="cube_10x10x10_hex.msh"/>
</Mesh>
```

Here are the first few lines of the msh file :

Listing 2.1: cube_10x10x10_hex.msh

```
$MeshFormat
2.2 0 8
$EndMeshFormat
$Nodes
1331
1      0      0      0
2      10     0      0
3      10     10     0
4      0      10     0
5      0      0      10
6      10     0      10
7      10     10     10
8      0      10     10
9      1      0      0
```

(continues on next page)

(continued from previous page)

10	2	0	0
11	3	0	0
12	4	0	0
13	5	0	0
14	6	0	0
15	7	0	0

GEOSX can run different physical solvers on different regions of the mesh at different times. Here, to keep things simple, we run one solver (single-phase flow) on the entire domain throughout the simulation. Even this is trivial, we need to define and name a region encompassing the entire domain and assign it to the single-phase flow solver. We also need to provide material properties to the regions. This is done by specifying `ElementRegions`. Here, the entire field is one region called `Domain`, and contains multiple constitutive models, including `water`, `rockPorosity`, and `rockPerm`.

```
<ElementRegions>
  <CellElementRegion
    name="Domain"
    cellBlocks="{ DEFAULT_HEX }"
    materialList="{ water, rock, rockPerm, rockPorosity, nullSolid }"/>
</ElementRegions>
```

Running GEOSX

The command to run GEOSX is

```
path/to/geosx -i ../../../../coreComponents/physicsSolvers/fluidFlow/integratedTests/
↳ singlePhaseFlow/pamela_test/3D_10x10x10_compressible_pamela_hex_gravity.xml
```

Note that all paths for files included in the XML file are relative to this XML file, not to the GEOSX executable. When running GEOSX, console messages will provide indications regarding the status of the simulation.

In our case, the first lines are:

```
Adding Solver of type SinglePhaseFlow, named SinglePhaseFlow
Adding Mesh: PAMELAMeshGenerator, CubeHex
Adding Geometric Object: Box, all
Adding Geometric Object: Box, left
Adding Event: PeriodicEvent, solverApplications
Adding Event: PeriodicEvent, outputs
Adding Event: PeriodicEvent, restarts
Adding Output: Silo, siloWellPump
Adding Output: Restart, restartOutput
Adding Object CellElementRegion named Domain from ObjectManager::Catalog.
```

This indicates initialization of GEOSX. The mesh preprocessing tool PAMELA is launched next, with console messages as follows.

```
0 >>> *****
0 >>> PAMELA Library Import tool
0 >>> *****
0 >>> GMSH FORMAT IDENTIFIED
0 >>> *** Importing Gmsh mesh format...
0 >>> Reading nodes...
0 >>> Done0
0 >>> Reading elements...
```

(continues on next page)

(continued from previous page)

```
0 >>> Number of nodes = 1331
0 >>> Number of triangles = 0
0 >>> Number of quadrilaterals = 0
0 >>> Number of tetrahedra = 0
0 >>> Number of hexahedra = 1000
0 >>> Number of pyramids = 0
0 >>> Number of prisms = 0
0 >>> *** Done
0 >>> *** Creating Polygons from Polyhedra...
0 >>> 3300 polygons have been created
0 >>> *** Done
0 >>> *** Perform partitioning...
0 >>> TRIVIAL partitioning...
0 >>> Ghost elements...
0 >>> Clean mesh...
0 >>> *** Done...
0 >>> Clean Adjacency...
0 >>> *** Done...
Writing into the GEOSX mesh data structure
Running simulation
```

Notice the specification of the number of nodes (1331), and hexahedra (1000). After the adjacency calculations, GEOSX starts the simulation itself, with the time-step increments specified in the XML file.

At the end of your simulation, you should see something like:


```
Time: 99s, dt:1s, Cycle: 99
Cleaning up events
Writing out restart file at 3D_10x10x10_compressible_pamela_hex_gravity_restart_
→000000100/rank_0000000.hdf5

init time = 0.081181s, run time = 5.4595s
```

Once this is done, GEOSX is finished and we can inspect the outcome.

Visualization of results in VisIt

All results are written in a format compatible with [VisIt](#). To load the results, point VisIt to the database file written in the Silo output folder.



We see that the face $x=0$ shown here in the back of the illustration applies a constant pressure boundary condition (colored in green), whereas the face across from it displays a pressure field under gravity effect, equilibrated and hydrostatic. These results are consistent with what we expect.

Let us now see if a tetrahedral mesh, under the same exact physical conditions, can reproduce these results.

2.2.2 Externally Generated Tetrahedral Elements

In the second part of the tutorial, we discretize the same cubic domain but with tetrahedral elements. Tetrahedral meshes are not yet common in geomodeling but offer tremendous flexibility in modeling fracture planes, faults, complex reservoir horizons and boundaries. Just like for hexahedral meshes, and for the same reasons (compatibility with finite volume and finite element methods), tetrahedral meshes in GEOSX must be conformal.

As stated previously, the problem we wish to solve here is the exact same physical problem as with hexahedral grid blocks. We apply a constant pressure condition (injection) from the $x=0$ vertical face of the domain, and we let pressure equilibrate over time. We observe the opposite side of the cube and expect to see hydrostatic pressure profiles because of the gravitational effect. The displacement is a single phase, compressible flow subject to gravity forces. We use GEOSX to compute the pressure inside each grid block.

The set-up for this problem is almost identical to the hexahedral mesh set-up. We simply point our `Mesh` tag to include a tetrahedral grid. The beauty of not relying on I,J,K indices for any property specification or well trajectory makes it **easy to try different meshes for the same physical problems with GEOSX**. Swapping out meshes without requiring other modifications to the input files makes mesh refinement studies easy to perform with GEOSX.


Like before, the XML file for this problem is the following:

```
src\CoreComponents\physicsSolvers\integratedTests\singlePhaseFlow\pamela_test\3D_
↪10x10x10_compressible_pamela_tetra_gravity.xml
```

The only difference, is that now, the `Mesh` tag points GEOSX to a different mesh file called `cube_10x10x10_tet.msh`. This file contains nodes and tetrahedral elements in `Gmsh` format, representing a different discretization of the exact same 10x10x10 cubic domain.

```
<Mesh>
  <PAMELAMeshGenerator
    name="CubeTetra"
    file="cube_10x10x10_tet.msh"/>
</Mesh>
```

The mesh now looks like this:



And the MSH file starts as follows (notice the tetrahedral point coordinates as real numbers):

Listing 2.2: cube_10x10x10_tet.msh

```
$MeshFormat
2.2 0 8
$EndMeshFormat
$Nodes
366
1 0 0 10
2 0 0 0
3 0 10 10
4 0 10 0
5 10 0 10
6 10 0 0
7 10 10 10
8 10 10 0
9 0 0 1.6666666666666662
10 0 0 3.3333333333333323
11 0 0 4.9999999999999986
12 0 0 6.6666666666666647
13 0 0 8.3333333333333321
14 0 1.6666666666666662 10
15 0 3.3333333333333323 10
```

Again, the entire field is one region called `Domain` and contains `water` and `rock` only.

```

<ElementRegions>
  <CellElementRegion
    name="Domain"
    cellBlocks="{ DEFAULT_TETRA }"
    materialList="{ water, rock, rockPerm, rockPorosity, nullSolid }"/>
  </CellElementRegion>
</ElementRegions>

```

Running GEOSX

The command to run GEOSX is

```

path/to/geosx -i ../../../../coreComponents/physicsSolvers/fluidFlow/integratedTests/
↳singlePhaseFlow/pamela_test/cube_10x10x10_tet.msh

```

Again, all paths for files included in the XML file are relative to this XML file, not to the GEOSX executable. When running GEOSX, console messages will provide indications regarding the status of the simulation. In our case, the first lines are:

```

Adding Solver of type SinglePhaseFVM, named SinglePhaseFlow
Adding Mesh: PAMELAMeshGenerator, CubeTetra
Adding Geometric Object: Box, left
Adding Event: PeriodicEvent, solverApplications
Adding Event: PeriodicEvent, outputs
Adding Event: PeriodicEvent, restarts
Adding Output: Silo, siloWellPump
Adding Output: Restart, restartOutput
Adding Object CellElementRegion named Domain from ObjectManager::Catalog.
Reading external mesh from /Users/j0529096/Documents/code/GEOSX/src/coreComponents/
↳physicsSolvers/fluidFlow/integratedTests/singlePhaseFlow/pamela_test/cube_10x10x10_
↳tet.msh

```

Followed by:

```

0 >>> *****
0 >>> PAMELA Library Import tool
0 >>> *****
0 >>> GMSH FORMAT IDENTIFIED
0 >>> *** Importing Gmsh mesh format...
0 >>> Reading nodes...
0 >>> Done0
0 >>> Reading elements...
0 >>> Number of nodes = 366
0 >>> Number of triangles = 624
0 >>> Number of quadrilaterals = 0
0 >>> Number of tetrahedra = 1153
0 >>> Number of hexahedra = 0
0 >>> Number of pyramids = 0
0 >>> Number of prisms = 0
0 >>> *** Done
0 >>> *** Creating Polygons from Polyhedra...
0 >>> 1994 polygons have been created
0 >>> *** Done
0 >>> *** Perform partitioning...
0 >>> TRIVIAL partitioning...
0 >>> Ghost elements...
0 >>> Clean mesh...

```

(continues on next page)

(continued from previous page)


```
0 >>> *** Done...
0 >>> Clean Adjacency...
0 >>> *** Done...
Writing into the GEOSX mesh data structure
  Domain/DEFAULT_TETRA/water is allocated with 1 quadrature points.
  Domain/DEFAULT_TETRA/rock is allocated with 1 quadrature points.
  Domain/DEFAULT_TETRA/rockPerm is allocated with 1 quadrature points.
  Domain/DEFAULT_TETRA/rockPorosity is allocated with 1 quadrature points.
  Domain/DEFAULT_TETRA/nullSolid is allocated with 1 quadrature points.
PAMELAMeshGenerator CubeTetra: importing field data from mesh dataset
```

We see that we have now 366 nodes and 1153 tetrahedral elements. And finally, when the simulation is successfully done we see:

```
Time: 0s, dt:1s, Cycle: 0
Time: 1s, dt:1s, Cycle: 1
Time: 2s, dt:1s, Cycle: 2
Time: 3s, dt:1s, Cycle: 3
Time: 4s, dt:1s, Cycle: 4
Time: 5s, dt:1s, Cycle: 5
...
Time: 95s, dt:1s, Cycle: 95
Time: 96s, dt:1s, Cycle: 96
Time: 97s, dt:1s, Cycle: 97
Time: 98s, dt:1s, Cycle: 98
Time: 99s, dt:1s, Cycle: 99
Cleaning up events
Umpire          HOST sum across ranks:    4.9 MB
Umpire          HOST          rank max:    4.9 MB
total time                      3.164s
initialization time             0.178s
run time                       2.659s
```

Visualization of results in VisIt

All results are written in a format compatible with [VisIt](#) by default. If we load into VisIt the *.database* file found in the Silo folder, we observe the following results:



Here, we can see that despite the different mesh sizes and shapes, we are able to recover our pressure profile without any problems, or degradation in runtime performance.

2.2.3 To go further

Feedback on this tutorial

This concludes the single-phase external mesh tutorial. For any feedback on this tutorial, please submit a [GitHub issue](#) on the project's [GitHub page](#).

For more details

- A complete description of the Internal Mesh generator is found here [Meshes](#).
- PAMELA being an external submodule has less documentation, but the same [Meshes](#) page may get you started.
- GEOSX can handle tetrahedra, hexahedra, prisms, pyramids, wedges, and any combination thereof in one mesh. For more information on how MSH formats can help you specify these mesh types, see the [Gmsh](#) website.

2.3 Tutorial 3: Regions and Property Specifications

Context

In this tutorial, we set up a simple field case for single-phase flow simulation (see [Singlephase Flow Solver](#)). We demonstrate how to run a basic flow simulation in the reservoir layer. We do not consider any coupling with wells. Injection and production will be specified by imposing a high pressure in the cells close to the injection area and a low pressure in the cells close to the production area.

Objectives

At the end of this tutorial you will know:

- how to import external mesh information and properties,
- how to run a specific solver (here, flow) in a specific region only,


- the basic method of using boxes to set up boundary conditions,
- how to use *TableFunction* to import fields varying in time and/or space,
- how to control output frequency and export results for visualization.

Input file

The XML input file for this test case is located at:


```
src/coreComponents/physicsSolvers/multiphysics/integratedTests/FieldCaseTutorial1.xml
```

We consider the following mesh as a numerical support to the simulations in this tutorial:



This mesh contains three continuous regions:

- a Top region (overburden, elementary tag = *Overburden*)
- a Middle region (reservoir layer, elementary tag = *Reservoir*)
- a Bottom region (underburden, elementary tag = *Underburden*)



The mesh is defined using the GMSH file format (see [Meshes](#) for more information on the supported mesh file format). Each tetrahedron is associated to a unique tag.

The XML file considered here follows the typical structure of the GEOSX input files:

1. *Solver*
2. *Mesh*
3. *Geometry*
4. *Events*
5. *NumericalMethods*
6. *ElementRegions*
7. *Constitutive*
8. *FieldSpecifications*
9. *Outputs*
10. *Functions*

2.3.1 Single-phase solver

Let us inspect the **Solver** XML tags.

```

<Solvers
  gravityVector="{ 0.0, 0.0, 0.0 }">
    <SinglePhaseFVM
      name="SinglePhaseFlow"
      discretization="singlePhaseTPFA"
      fluidNames="{ water }"
      solidNames="{ rock }"
      permeabilityNames="{ rockPerm }"
      targetRegions="{ Reservoir }">
        <NonlinearSolverParameters
          newtonTol="1.0e-6"
          newtonMaxIter="8"/>
        <LinearSolverParameters
          solverType="gmres"
          preconditionerType="amg"
          krylovTol="1.0e-10"/>
      </SinglePhaseFVM>
    </Solvers>

```

This node gathers all the information previously defined. We use a classical `SinglePhaseFVM` Finite Volume Method, with the two-point flux approximation as will be defined in the **NumericalMethod** tag. The `targetRegions` refers only to the Reservoir region because we only solve for flow in this region. The `fluidNames` and `solidNames` refer the materials defined in the **Constitutive** tag.

The `NonlinearSolverParameters` and `LinearSolverParameters` are used to set usual numerical solver parameters such as the linear and nonlinear tolerances, the preconditioner and solver types or the maximum number of nonlinear iterations.

2.3.2 Mesh

Here, we use the `PAMELAMeshGenerator` to load the mesh (see *Importing the Mesh*). The syntax to import external meshes is simple : in the XML file, the mesh `file` is included with its relative or absolute path to the location of the GEOSX XML file and a user-specified name label for the mesh object.

```

<Mesh>
  <PAMELAMeshGenerator
    name="SyntheticMesh"
    file="synthetic.msh"/>
  </Mesh>

```

2.3.3 Geometry

Here, we are using definition of `source` and `sink` boxes in addition to the `all` box in order to flag sets of nodes or cells which will act as injection or production.

```

<Geometry>
  <Box
    name="all"
    xMin="{ -1e9, -1e9, -1e9 }"
    xMax="{ 1e9, 1e9, 1e9 }"/>

  <Box
    name="source"
    xMin="{ -0.1, 9700, 4499.9 }"

```

(continues on next page)

(continued from previous page)


```

xMax="{ 300, 10000.01, 5500.1 }"/>

<Box
  name="sink"
  xMin="{ 19700, -0.1, 4499.9 }"
  xMax="{ 20000.1, 300, 5500.1 }"/>
</Geometry>

```

In order to define a box, the user defines `xMax` and `xMin`, two diagonally opposite nodes of the box.



2.3.4 Events

The events are used here to guide the simulation through time, and specify when outputs must be triggered.

```

<Events
  maxTime="500.0e6">
  <PeriodicEvent
    name="solverApplications"
    forceDt="2.0e6"
    target="/Solvers/SinglePhaseFlow"/>

  <PeriodicEvent
    name="outputs"
    timeFrequency="10.0e6"

```

(continues on next page)

(continued from previous page)

```
target="/Outputs/syntheticReservoirVizFile"/>
</Events>
```

The **Events** tag is associated with the `maxTime` keyword defining the maximum time. If this time is ever reached or exceeded, the simulation ends.

Two `PeriodicEvent` are defined. - The first one, `solverApplications`, is associated with the solver. The `forcedDt` keyword means that there will always be time-steps of 23 days (2 000 000 seconds). - The second, `outputs`, is associated with the output. The `timeFrequency` keyword means that it will be executed every 116 days (10 000 000 seconds). The `targetExactTimestep` is set to 1, meaning that the Event Manager will impose this event will be triggered exactly every 116 days, constraining schedule decided by application to match this date.

2.3.5 Numerical methods

Defining the numerical method used in the solver, we will provide information on how to discretize our equations. Here a classical two-point flux approximation (TPFA) scheme is used to discretize water fluxes over faces.

```
<NumericalMethods>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="singlePhaseTPFA"
      targetRegions="{ Reservoir }"
      fieldName="pressure"
      coefficientName="permeability"/>
    </FiniteVolume>
  </NumericalMethods>
```

The `TwoPointFluxApproximation` node should specify the primary field to solve for as `fieldName`. For a flow problem, this field is the pressure. Here we specified `targetRegions` as we only solve flow for reservoir. The field under `coefficientName` is used during TPFA transmissibilities construction.

2.3.6 Regions

Assuming that the overburden and the underburden are impermeable, and flow only takes place in the reservoir, we need to define regions.

There are two methods to achieve this regional solve.

- The first solution is to define a unique `CellElementRegion` corresponding to the reservoir.

```
<ElementRegion>
  <CellElementRegion name="ReservoirLayer"
    cellBlocks="{Reservoir_TETRA}"
    materialList="{ water, rock, rockPerm, rockPorosity,
↵nullSolid }">
  </CellElementRegion>
</ElementRegion>
```

- The second solution is to define all the `CellElementRegions` as they are in the GMSH file, but defining the solvers only on the reservoir layer. In this case, the **ElementRegion** tag is :

```
<ElementRegions>
  <CellElementRegion
    name="Reservoir"
    cellBlocks="{ Reservoir_TETRA }"
```

(continues on next page)

(continued from previous page)

```

materialList="{ water, rock, rockPerm, rockPorosity, nullSolid }"/>

<CellElementRegion
  name="Burden"
  cellBlocks="{ Overburden_TETRA, Underburden_TETRA }"
  materialList="{ water, rock, rockPerm, rockPorosity, nullSolid }"/>
</ElementRegions>

```

We opt for the latest as it allows to visualize over- and underburdens and to change regions handling in their tag without needing to amend the **ElementRegion** tag.

Note: The material list here was set for a single-phase flow problem. This list is subject to change if the problem is not a single-phase flow problem.

2.3.7 Constitutive models

We simulate a single-phase flow in the reservoir layer, hence with multiple types of materials, a fluid (water) and solid (rock permeability and porosity).

```

<Constitutive>
  <CompressibleSinglePhaseFluid
    name="water"
    defaultDensity="1000"
    defaultViscosity="0.001"
    referencePressure="0.0"
    compressibility="1e-9"
    viscosibility="0.0"/>

  <CompressibleSolidConstantPermeability
    name="rock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"/>

  <NullModel
    name="nullSolid"/>

  <PressurePorosity
    name="rockPorosity"
    defaultReferencePorosity="0.05"
    referencePressure="0.0"
    compressibility="1.0e-9"/>

  <ConstantPermeability
    name="rockPerm"
    permeabilityComponents="{ 1.0e-14, 1.0e-14, 1.0e-18 }"/>
</Constitutive>

```

The constitutive parameters such as the density, the viscosity, and the compressibility are specified in the International System of Units.

Note: To consider an incompressible fluid, the user has to set the compressibility to 0.

2.3.8 Defining properties

The next step is to specify fields, including:

- The initial value (here, the pressure has to be initialized)
- The static properties (here, we have to define the permeability tensor and the porosity)
- The boundary conditions (here, the injection and production pressure have to be set)

```
<FieldSpecifications>
  <FieldSpecification
    name="initialPressure"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/Reservoir/Reservoir_TETRA"
    fieldName="pressure"
    scale="1"
    functionName="initialPressureFunc"/>

  <FieldSpecification
    name="sourceTerm"
    objectPath="ElementRegions/Reservoir/Reservoir_TETRA"
    fieldName="pressure"
    scale="1e7"
    setNames="{ source }"
    functionName="timeInj"/>

  <FieldSpecification
    name="sinkTerm"
    objectPath="ElementRegions/Reservoir/Reservoir_TETRA"
    fieldName="pressure"
    scale="0.0"
    setNames="{ sink }"/>
</FieldSpecifications>
```

You may note :

- All static parameters and initial value fields must have `initialCondition` field set to 1.
- The `objectPath` refers to the `ElementRegion` in which the field has its value,
- The `setName` field points to the box previously defined to apply the fields,
- `name` and `fieldName` have a different meaning: `name` is used to give a name to the XML block. This name must be unique. `fieldName` is the name of the field registered in GEOSX. This value has to be set according to the expected input fields of each solver.

Note: GEOSX handles permeability as a diagonal matrix, so the three values of the permeability tensor are set individually using the `component` field,

2.3.9 Output

The **Outputs** XML tag is used to trigger the writing of visualization files. Here, we write files in a format natively readable by Paraview under the tar *VTK*

```

<Outputs>
  <!-- <Silo name="syntheticReservoirVizFile"/> -->
  <VTK
    name="syntheticReservoirVizFile"/>
</Outputs>

```

Note: The name keyword defines the name of the output directory.

2.3.10 Using functions to specify properties

Eventually, one can define varying properties using `TableFunction` (*Functions*) under the **Functions** tag:

```

<Functions>
  <TableFunction
    name="timeInj"
    inputVarNames="{ time }"
    coordinates="{ 200e6, 250e6, 500e6 }"
    values="{ 1, 0.01, 0.00001 }"/>

  <TableFunction
    name="initialPressureFunc"
    inputVarNames="{ elementCenter }"
    coordinateFiles="{ tables_FieldCaseTuto/xlin.geos, tables_FieldCaseTuto/ylin.
↪geos, tables_FieldCaseTuto/zlin.geos }"
    voxelFile="tables_FieldCaseTuto/pressure.geos"/>

  <TableFunction
    name="permxFunc"
    inputVarNames="{ elementCenter }"
    coordinateFiles="{ tables_FieldCaseTuto/xlin.geos, tables_FieldCaseTuto/ylin.
↪geos, tables_FieldCaseTuto/zlin.geos }"
    voxelFile="tables_FieldCaseTuto/permx.geos"
    interpolation="nearest"/>

  <TableFunction
    name="permyFunc"
    inputVarNames="{ elementCenter }"
    coordinateFiles="{ tables_FieldCaseTuto/xlin.geos, tables_FieldCaseTuto/ylin.
↪geos, tables_FieldCaseTuto/zlin.geos }"
    voxelFile="tables_FieldCaseTuto/permy.geos"
    interpolation="nearest"/>


  <TableFunction
    name="permzFunc"
    inputVarNames="{ elementCenter }"
    coordinateFiles="{ tables_FieldCaseTuto/xlin.geos, tables_FieldCaseTuto/ylin.
↪geos, tables_FieldCaseTuto/zlin.geos }"
    voxelFile="tables_FieldCaseTuto/permz.geos"
    interpolation="nearest"/>
</Functions>

```

Here, the injection pressure is set to vary with time. Attentive reader might have noticed that `sourceTerm` was bound to a `TableFunction` named *timeInj* under **FieldSpecifications** tag definition. The initial pressure is set based on the values contained in the table formed by the files which are specified. In particular, the files *xlin.geos*, *ylin.geos* and

zlin.geos define a regular meshing of the bounding box containing the reservoir. The *pressure.geos* file then defines the values of the pressure at those points.

We proceed in a similar manner as for *pressure.geos* to map a heterogeneous permeability field (here the 5th layer of the SPE 10 test case) onto our unstructured grid. This mapping will use a nearest point interpolation rule.



Note: The varying values imposed in *values* or passed through *voxelFile* are premultiplied by the *scale* attribute from **FieldSpecifications**.

2.3.11 Running GEOSX

The simulation can be launched with:

```
geosx -i FieldCaseTutorial11.xml
```

One can notice the correct load of the field function among the starting output messages

```
Adding Mesh: PAMELAMeshGenerator, SyntheticMesh
Adding Solver of type SinglePhaseFVM, named SinglePhaseFlow
Adding Geometric Object: Box, all
Adding Geometric Object: Box, source
Adding Geometric Object: Box, sink
Adding Output: VTK, syntheticReservoirVizFile
Adding Event: PeriodicEvent, solverApplications
Adding Event: PeriodicEvent, outputs
    TableFunction: timeInj
    TableFunction: initialPressureFunc
    TableFunction: permxFunc
```


(continues on next page)

(continued from previous page)


```
TableFunction: permyFunc
TableFunction: permzFunc
Adding Object CellElementRegion named Reservoir from ObjectManager::Catalog.
Adding Object CellElementRegion named Burden from ObjectManager::Catalog.
```

2.3.12 Visualization of results

We can open the file *syntheticReservoirVizFile.pvd* with Paraview to visualize the simulation results. The initial pressure field in the reservoir region is provided below as an example.



Since, in the event block, we have asked for the output to be generated at regular intervals throughout the simulation, we can also visualize the pressure distribution at different simulation times, showing the variation in the injection control.



2.3.13 To go further

Feedback on this tutorial

This concludes this tutorial. For any feedback, please submit a [GitHub issue on the project's GitHub page](#).

For more details

- More on meshes, please see [Meshes](#).
- More on events, please see [Event Management](#).

2.4 Tutorial 4: Boundary Conditions and Time-Dependent Functions

Context

In this tutorial, we use a small strain linear elastic based solid mechanics solver (see [Solid Mechanics Solver](#)) from GEOSX to solve for the bending problem of a three-dimensional cantilever beam. The beam is fixed at one end, and subjects to a traction force pointing to the y-positive direction on the other end. The beam is deformed in the x-y plane.

Objectives

At the end of this tutorial, you will know:

- how to use the solid mechanics solver to solve a quasistatic problem,
- how to set up displacement boundary condition at element nodes,
- how to set up traction boundary condition on element surfaces,

- how to use a table function to control time-dependent loading.


Input file

This tutorial uses no external input files and everything required is contained within a single GEOSX input file. The xml input file for this test case is located at:

```
src/coreComponents/physicsSolvers/solidMechanics/integratedTests/SSLE-QS-beamBending.  
↪xml
```

2.4.1 Discretized computational domain

The following mesh is used in this tutorial:



This mesh contains 80 x 8 x 4 eight-node brick elements in the x, y and z directions, respectively. Here, the InternalMesh is used to generate a structured three-dimensional mesh with C3D8 as the elementTypes. This mesh is defined as a cell block with the name cb1.

```
<Mesh>  
  <InternalMesh  
    name="mesh1"  
    elementTypes="{ C3D8 }"  
    xCoords="{ 0, 80 }"  
    yCoords="{ 0, 8 }"  
    zCoords="{ 0, 4 }"  
    nx="{ 80 }"  
    ny="{ 8 }"  
    nz="{ 4 }"  
    cellBlockNames="{ cb1 }"/>  
</Mesh>
```

2.4.2 Gravity

The gravity is turned off explicitly at the beginning of the input file:

```
<Solvers  
  gravityVector="{ 0.0, 0.0, 0.0 }">
```

2.4.3 Solid mechanics solver

The solid mechanics solver is based on the small strain Lagrangian finite element formulation. The problem is run as `QuasiStatic` without considering the beam inertial. The computational domain is discretized by `FE1`, which is defined in the `NumericalMethods` block. The material is designated as `shale`, whose properties are defined in the `Constitutive` block.

```
<SolidMechanicsLagrangianSSLE
  name="lagsolve"
  timeIntegrationOption="QuasiStatic"
  discretization="FE1"
  targetRegions="{ Region2 }"
  solidMaterialNames="{ shale }">
```

2.4.4 Finite element discretization

The computational domain is discretized by `C3D8` elements with the first order interpolation functions at each direction in the parent domain. The $2 \times 2 \times 2$ Gauss quadrature rule is adopted to be compatible with the first order interpolation functions.

```
<NumericalMethods>
  <FiniteElements>
    <FiniteElementSpace
      name="FE1"
      order="1"/>
    </FiniteElements>
  </NumericalMethods>
```

2.4.5 Constitutive model

Recall that in the `SolidMechanicsLagrangianSSLE` block, `shale` is designated as the material in the computational domain. Here, the material is defined as linear isotropic.

```
<ElasticIsotropic
  name="shale"
  defaultDensity="2700"
  defaultBulkModulus="5.5556e9"
  defaultShearModulus="4.16667e9"/>
```

2.4.6 Boundary conditions

As aforementioned, the beam is fixed on one end, and subjects to surface traction on the other end. These boundary conditions are set up through the `FieldSpecifications` block. Here, `nodeManager` and `faceManager` in the `objectPath` indicate that the boundary conditions are applied to the element nodes and faces, respectively. Component 0, 1, and 2 refer to the x, y, and z direction, respectively. And the non-zero values given by `Scale` indicate the magnitude of the loading. Some shorthands, such as `xneg` and `xpos`, are used as the locations where the boundary conditions are applied in the computational domain. For instance, `xneg` means the portion of the computational domain located at the left-most in the x-axis, while `xpos` refers to the portion located at the right-most area in the x-axis. Similar shorthands include `ypos`, `yneg`, `zpos`, and `zneg`. Particularly, the time-dependent loading applied at the beam tip is defined through a function with the name `timeFunction`.

```

<FieldSpecifications>
  <FieldSpecification
    name="xnegconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="0"
    scale="0.0"
    setNames="{ xneg }"/>

  <FieldSpecification
    name="yconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="1"
    scale="0.0"
    setNames="{ xneg }"/>

  <FieldSpecification
    name="zconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="2"
    scale="0.0"
    setNames="{ zneg, zpos }"/>

  <Traction
    name="xposconstraint"
    objectPath="faceManager"
    scale="1.0e6"
    direction="{ 0, 1, 0 }"
    functionName="timeFunction"
    setNames="{ xpos }"/>
</FieldSpecifications>

```

2.4.7 Table function

A table function is used to define the time-dependent loading at the beam tip. The `coordinates` and `values` form a time-magnitude pair for the loading time history. In this case, the loading magnitude increases linearly as the time evolves.

```

<Functions>
  <TableFunction
    name="timeFunction"
    inputVarNames="{ time }"
    coordinates="{ 0.0, 10.0 }"
    values="{ 0.0, 10.0 }"/>
</Functions>

```


2.4.8 Execution

Finally, the execution of the simulation is set up in the `Events` block, where `target` points to the solid mechanics solver defined in the `Solvers` block, and the time increment `forceDt` is set as 1.0s.

```
<PeriodicEvent
  name="solverApplications"
  forceDt="1.0"
  target="/Solvers/lagsolve"/>
```

2.4.9 Result

The deformed beam is shown as following (notice that the displacement is visually magnified):



2.4.10 To go further

Feedback on this tutorial

This concludes the solid mechanics for small-strain linear elasticity tutorial. For any feedback on this tutorial, please submit a [GitHub issue](#) on the project's [GitHub](#) page.

For more details

- More on meshes, please see *[Meshes](#)*.
- More on events, please see *[Event Management](#)*.

3.1 Multiphase Flow

Context

In this example, we set up a multiphase, multicomponent test case (see *Compositional Multiphase Flow Solver*). The permeability field corresponds to the two bottom layers (layers 84 and 85) of the SPE10 test case. The thermodynamic behavior of the fluid mixture is specified using a simple immiscible two-phase (Dead-Oil) model. Injection and production are simulated using boundary conditions.

Objective

The main objective of this example is to review the main elements of a simple two-phase simulation in GEOSX, including:

- the compositional multiphase flow solver,
- the multiphase constitutive models,
- the specifications of multiphase boundary conditions.

Input file

This example is based on the XML file located at

```
src/coreComponents/physicsSolvers/fluidFlow/benchmarks/SPE10/dead_oil_spe10_layers_84_  
↪ 85.xml
```

The XML file considered here follows the typical structure of the GEOSX input files:

1. *Solver*
2. *Mesh*
3. *Geometry*
4. *Events*
5. *NumericalMethods*

6. *ElementRegions*
7. *Constitutive*
8. *FieldSpecifications*
9. *Outputs*
10. Tasks

3.1.1 Multiphase flow solver

In GEOSX, the setup of a multiphase simulation starts in the **Solvers** XML block of the input file. This example relies on a solver of type **CompositionalMultiphaseFVM** that implements a fully implicit finite-volume scheme based on the standard two-point approximation of the flux (TPFA). More information on this solver can be found at [Compositional Multiphase Flow Solver](#).

Let us have a closer look at the **Solvers** XML block displayed below. The solver has a name (here, `compflow`) that can be chosen by the user and is not imposed by GEOSX. Note that this name is used in the **Events** XML block to trigger the application of the solver. Using the `targetRegions` attribute, the solver defines the target regions on which it is applied. In this example, there is only one region, named `reservoir`.

The constitutive models defined on these target regions must be listed in the **CompositionalMultiphaseFVM** block. This is done by passing the name of the fluid PVT model using the `fluidNames` attribute, the name of the solid compressibility model using the `solidNames` attribute, the name of the rock permeability model using the `permeabilityNames` attribute, and the name of the relative permeability model using the `relPermNames` attribute. If a capillary pressure model is employed in the simulation, its name must also be passed here, using the `capPressureNames` attribute. All the constitutive model names passed here must be defined in the **Constitutive** block of the XML file (see below).

The **CompositionalMultiphaseFVM** block contains two important sub-blocks, namely **NonlinearSolverParameters** and **LinearSolverParameters**. In **NonlinearSolverParameters**, one can finely tune the nonlinear tolerance, the application of the linear search algorithm, and the heuristics used to increase the time step size. In **LinearSolverParameters**, the user can specify the linear tolerance, the type of (direct or iterative) linear solver, and the type of preconditioner, if any. For large multiphase flow problems, we recommend using an iterative linear solver (`solverType="gmres"` or `solverType="fgmres"`) combined with the multigrid reduction (MGR) preconditioner (`preconditionerType="mgr"`). More information about the MGR preconditioner can be found in [Linear Solvers](#).

Note: For non-trivial simulations, we recommend setting the `initialDt` attribute to a small value (relative to the time scale of the problem) in seconds. If the simulation appears to be slow, use `logLevel="1"` in **CompositionalMultiphaseFVM** to detect potential Newton convergence problems. If the Newton solver struggles, please set `lineSearchAction="Attempt"` in **NonlinearSolverParameters**. If the Newton convergence is good, please add `logLevel="1"` in the **LinearSolverParameters** block to detect linear solver problems, especially if an iterative linear solver is used.

Note: To use the linear solver options of this example, you need to ensure that GEOSX is configured to use the Hypre linear solver package.

```
<Solvers>

  <CompositionalMultiphaseFVM
    name="compflow"
```

(continues on next page)

(continued from previous page)

```

    logLevel="1"
    discretization="fluidTPFA"
    fluidNames="{ fluid }"
    solidNames="{ rock }"
    relPermNames="{ relperm }"
    permeabilityNames="{ rockPerm }"
    targetRegions="{ reservoir }"
    temperature="300"
    useMass="1"
    initialDt="1e3"
    maxCompFractionChange="0.1"
    computeCFLNumbers="1">
    <NonlinearSolverParameters
      newtonTol="1.0e-4"
      newtonMaxIter="40"
      maxTimeStepCuts="10"
      lineSearchAction="None"/>
    <LinearSolverParameters
      solverType="fgmres"
      preconditionerType="mgr"
      krylovTol="1.0e-5"/>
    </CompositionalMultiphaseFVM>
  </Solvers>

```

3.1.2 Mesh

In this simulation, we define a simple mesh generated internally using the **InternalMesh** generator, as illustrated in the previous examples. The mesh dimensions and cell sizes are chosen to be those specified in the SPE10 test case, but are limited to the two bottom layers. The mesh description must be done in meters.

```

<Mesh>
  <InternalMesh
    name="mesh"
    elementTypes="{ C3D8 }"
    xCoords="{ 0, 365.76 }"
    yCoords="{ 0, 670.56 }"
    zCoords="{ 0, 1.22 }"
    nx="{ 60 }"
    ny="{ 220 }"
    nz="{ 2 }"
    cellBlockNames="{ block }"/>
  </InternalMesh>
</Mesh>

```

3.1.3 Geometry

As in the previous examples, the **Geometry** XML block is used to select the cells in which the boundary conditions are applied. To mimic the setup of the original SPE10 test case, we place a source term in the middle of the domain, and a sink term in each corner. The specification of the boundary conditions applied to the selected mesh cells is done in the **FieldSpecifications** block of the XML file using the names of the boxes defined here.

```
<Geometry>
```

(continues on next page)

(continued from previous page)

```

<Box
  name="source"
  xMin="{ 182.85, 335.25, -0.01 }"
  xMax="{ 189.00, 338.35, 2.00 }"/>
<Box
  name="sink1"
  xMin="{ -0.01, -0.01, -0.01 }"
  xMax="{ 6.126, 3.078, 2.00 }"/>
<Box
  name="sink2"
  xMin="{ -0.01, 667.482, -0.01 }"
  xMax="{ 6.126, 670.60, 2.00 }"/>
<Box
  name="sink3"
  xMin="{ 359.634, -0.01, -0.01 }"
  xMax="{ 365.8, 3.048, 2.00 }"/>
<Box
  name="sink4"
  xMin="{ 359.634, 667.482, -0.01 }"
  xMax="{ 365.8, 670.60, 2.00 }"/>

</Geometry>

```

3.1.4 Events

In the **Events** XML block of this example, we specify two types of **PeriodicEvents** serving different purposes, namely solver application and result output.

The periodic event named `solverApplications` triggers the application of the solver on its target region. This event must point to the solver by name. In this example, the name of the solver is `compflow` and was defined in the **Solvers** block. The time step is initialized using the `initialDt` attribute of the flow solver. Then, if the solver converges in more than a certain number of nonlinear iterations (by default, 40% of the maximum number of nonlinear iterations), the time step will be increased until it reaches the maximum time step size specified with `maxEventDt`. If the time step fails, the time step will be cut. The parameters defining the time stepping strategy can be finely tuned by the user in the flow solver block. Note that all times are in seconds.

The output event forces GEOSX to write out the results at the frequency specified by the attribute `timeFrequency`. Here, we choose to output the results using the VTK format (see [Tutorial 2: External Meshes](#) for a tutorial that uses the Silo output file format). Using `targetExactTimestep=1` in this XML block forces GEOSX to adapt the time stepping to ensure that an output is generated exactly at the time frequency requested by the user. In the `target` attribute, we must use the name defined in the **VTK** XML tag inside the **Output** XML section, as documented at the end of this example (here, `vtkOutput`).

More information about events can be found at [Event Management](#).

```

<Events
  maxTime="8e7">

  <PeriodicEvent
    name="outputs"
    timeFrequency="4e7"
    targetExactTimestep="1"
    target="/Outputs/vtkOutput"/>

  <PeriodicEvent

```

(continues on next page)

(continued from previous page)

```

name="solverApplications"
maxEventDt="1e6"
target="/Solvers/compflow"/>

```

```

</Events>

```

3.1.5 Numerical methods

In the **NumericalMethods** XML block, we select a two-point flux approximation (TPFA) finite-volume scheme to discretize the governing equations on the reservoir mesh. TPFA is currently the only numerical scheme that can be used with a flow solver of type **CompositionalMultiphaseFVM**.

```

<NumericalMethods>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="fluidTPFA"
      fieldName="pressure"
      targetRegions="{ reservoir }"
      coefficientName="permeability"
      coefficientModelNames="{ rockPerm }"/>
    </FiniteVolume>
  </NumericalMethods>

```

3.1.6 Reservoir region

In the **ElementRegions** XML block, we define a **CellElementRegion** named `reservoir` corresponding to the reservoir mesh. The attribute `cellBlocks` is set to `block` to point this element region to the hexahedral mesh defined internally.

The **CellElementRegion** must also point to the constitutive models that are used to update the dynamic rock and fluid properties in the cells of the reservoir mesh. The names `fluid`, `rockPorosity`, `rockPerm` and `relperm` used for this in the `materialList` correspond to the attribute name of the **Constitutive** block.

```

<ElementRegions>
  <CellElementRegion
    name="reservoir"
    cellBlocks="{ block }"
    materialList="{ fluid, rock, relperm, rockPerm, rockPorosity, nullSolid }"/>
  </ElementRegions>

```

3.1.7 Constitutive models

For a simulation performed with the **CompositionalMultiphaseFVM** physics solver, at least four types of constitutive models must be specified in the **Constitutive** XML block:

- a fluid model describing the thermodynamics behavior of the fluid mixture,
- a relative permeability model,
- a rock permeability model,
- a rock porosity model.

All these models use SI units exclusively. A capillary pressure model can also be specified in this block but is omitted here for simplicity.

Here, we introduce a fluid model describing a simplified mixture thermodynamic behavior. Specifically, we use an immiscible two-phase (Dead Oil) model by placing the XML tag **DeadOilFluid**. Other fluid models can be used with the **CompositionalMultiphaseFVM** solver, as explained in *Fluid Models*.

With the tag **BrooksCoreyRelativePermeability**, we define a relative permeability model. A list of available relative permeability models can be found at *Relative Permeability Models*.

The properties are chosen to match those of the original SPE10 test case.

Note: The names and order of the phases listed for the attribute `phaseNames` must be identical in the fluid model (here, **DeadOilFluid**) and the relative permeability model (here, **BrooksCoreyRelativePermeability**). Otherwise, GEOSX will throw an error and terminate.

We also introduce models to define rock compressibility and permeability. This step is similar to what is described in the previous examples (see for instance *Tutorial 1: First Steps*).

We remind the reader that the attribute `name` of the constitutive models defined here must be used in the **ElementRegions** and **Solvers** XML blocks to point the element regions and the physics solvers to their respective constitutive models.

```
<Constitutive>

  <DeadOilFluid
    name="fluid"
    phaseNames="{ oil, water }"
    surfaceDensities="{ 800.0, 1022.0 }"
    componentMolarWeight="{ 114e-3, 18e-3 }"
    hydrocarbonFormationVolFactorTableNames="{ B_o_table }"
    hydrocarbonViscosityTableNames="{ visc_o_table }"
    waterReferencePressure="30600000.1"
    waterFormationVolumeFactor="1.03"
    waterCompressibility="0.00000000041"
    waterViscosity="0.0003"/>

  <CompressibleSolidConstantPermeability
    name="rock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"/>

  <NullModel
    name="nullSolid"/>

  <PressurePorosity
    name="rockPorosity"
    defaultReferencePorosity="0.1"
    referencePressure="1.0e7"
    compressibility="1e-10"/>

  <BrooksCoreyRelativePermeability
    name="relperm"
    phaseNames="{ oil, water }"
    phaseMinVolumeFraction="{ 0.0, 0.0 }"
    phaseRelPermExponent="{ 2.0, 2.0 }"
```

(continues on next page)

(continued from previous page)

```

    phaseRelPermMaxValue="{ 1.0, 1.0 }"/>

    <ConstantPermeability
      name="rockPerm"
      permeabilityComponents="{ 1.0e-14, 1.0e-14, 1.0e-18 }"/>

  </Constitutive>

```

3.1.8 Initial and boundary conditions

In the **FieldSpecifications** section, we define the initial and boundary conditions as well as the geological properties (porosity, permeability). All this is done using SI units. Here, we focus on the specification of the initial and boundary conditions for a simulation performed with the **CompositionalMultiphaseFVM** solver. We refer to [Tutorial 1: First Steps](#) for a more general discussion on the **FieldSpecification** XML blocks.

For a simulation performed with the **CompositionalMultiphaseFVM** solver, we have to set the initial pressure as well as the initial global component fractions (in this case, the oil and water component fractions). The `component` attribute of the **FieldSpecification** XML block must use the order in which the `phaseNames` have been defined in the **DeadOilFluid** XML block. In other words, `component=0` is used to initialize the oil global component fraction and `component=1` is used to initialize the water global component fraction, because we previously set `phaseNames="{oil, water}"` in the **DeadOilFluid** XML block.

To specify the sink terms, we use the **FieldSpecification** mechanism in a similar fashion to impose the sink pressure and composition. This is done to mimic a pressure-controlled well (before breakthrough). To specify the source term, we use a **SourceFlux** block to impose a fixed mass injection rate of component 1 (water) to mimic a rate-controlled well.

```

<FieldSpecifications>
  <FieldSpecification
    name="permx"
    component="0"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/reservoir/block"
    fieldName="rockPerm_permeability"
    functionName="permxFunc"
    scale="9.869233e-16"/>
  <FieldSpecification
    name="permy"
    component="1"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/reservoir/block"
    fieldName="rockPerm_permeability"
    functionName="permyFunc"
    scale="9.869233e-16"/>
  <FieldSpecification
    name="permz"
    component="2"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/reservoir/block"
    fieldName="rockPerm_permeability"
    functionName="permzFunc"
    scale="9.869233e-16"/>

```

(continues on next page)

(continued from previous page)

```

<FieldSpecification
  name="referencePorosity"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="rockPorosity_referencePorosity"
  functionName="poroFunc"
  scale="1.0"/>

<FieldSpecification
  name="initialPressure"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="pressure"
  scale="4.1369e+7"/>

<FieldSpecification
  name="initialComposition_oil"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="globalCompFraction"
  component="0"
  scale="0.9999"/>

<FieldSpecification
  name="initialComposition_water"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="globalCompFraction"
  component="1"
  scale="0.0001"/>

<SourceFlux
  name="sourceTerm"
  objectPath="ElementRegions/reservoir/block"
  scale="-0.07279"
  component="1"
  setNames="{ source }"/>

<FieldSpecification
  name="sinkPressure"
  setNames="{ sink1, sink2, sink3, sink4 }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="pressure"
  scale="2.7579e+7"/>

<FieldSpecification
  name="sinkComposition_oil"
  setNames="{ sink1, sink2, sink3, sink4 }"
  objectPath="ElementRegions/reservoir/block"
  fieldName="globalCompFraction"
  component="0"
  scale="0.9999"/>

<FieldSpecification
  name="sinkComposition_water"
  setNames="{ sink1, sink2, sink3, sink4 }"

```

(continues on next page)

(continued from previous page)

```

objectPath="ElementRegions/reservoir/block"
fieldName="globalCompFraction"
component="1"
scale="0.0001"/>

```

```
</FieldSpecifications>
```

3.1.9 Output

In this section, we request an output of the results in VTK format. Note that the name defined here must match the names used in the **Events** XML block to define the output frequency.

```

<Outputs>
  <VTK
    name="vtkOutput"/>
</Outputs>

```

All elements are now in place to run GEOSX.

3.1.10 Running GEOSX

The first few lines appearing to the console are indicating that the XML elements are read and registered correctly:

```

Adding Solver of type CompositionalMultiphaseFVM, named compflow
Adding Mesh: InternalMesh, mesh
Adding Geometric Object: Box, source
Adding Geometric Object: Box, sink1
Adding Geometric Object: Box, sink2
Adding Geometric Object: Box, sink3
Adding Geometric Object: Box, sink4
Adding Event: PeriodicEvent, outputs
Adding Event: PeriodicEvent, solverApplications
TableFunction: permxFunc
TableFunction: permyFunc
TableFunction: permzFunc
TableFunction: poroFunc
TableFunction: B_o_table
TableFunction: visc_o_table
Adding Output: VTK, vtkOutput
Adding Object CellElementRegion named region from ObjectManager::Catalog.
region/block/fluid is allocated with 1 quadrature points.
region/block/rock is allocated with 1 quadrature points.
aaregion/block/relperm is allocated with 1 quadrature points.

```

At this point, we are done with the case set-up and the code steps into the execution of the simulation itself:

```

Time: 0s, dt:1000s, Cycle: 0

Attempt: 0, NewtonIter: 0
( Rfluid ) = ( 2.28e+00 ) ;      ( R ) = ( 2.28e+00 ) ;
Attempt: 0, NewtonIter: 1
( Rfluid ) = ( 8.83e-03 ) ;      ( R ) = ( 8.83e-03 ) ;
Last LinSolve(iter,res) = ( 2, 2.74e-03 ) ;

```

(continues on next page)

(continued from previous page)

```

Attempt: 0, NewtonIter: 2
( Rfluid ) = ( 8.86e-05 ) ;      ( R ) = ( 8.86e-05 ) ;
Last LinSolve(iter,res) = ( 2, 8.92e-03 ) ;

compflow: Max phase CFL number: 0.00399585
compflow: Max component CFL number: 0.152466
compflow: Newton solver converged in less than 16 iterations, time-step required will
↳be doubled.

```

3.1.11 Visualization

A file compatible with Paraview is produced in this example. It is found in the output folder, and usually has the extension *.pvd*. More details about this file format can be found [here](#). We can load this file into Paraview directly and visualize results:

3.1.12 To go further

Feedback on this example

This concludes the example on setting up an immiscible two-phase flow simulation in a channelized permeability field. For any feedback on this example, please submit a [GitHub issue on the project's GitHub page](#).

For more details

- A complete description of the reservoir flow solver is found here: *Compositional Multiphase Flow Solver*.
- The available constitutive models are listed at *Constitutive Models*.

3.2 Multiphase Flow with Wells

Context

In this example, we build on the concepts presented in *Multiphase Flow* to show how to set up a multiphase water injection problem with wells in the three-dimensional *Egg model*. The twelve wells (four producers and eight injectors) are placed according to the description of the original test case.

Objectives

In this example, we re-use many GEOSX features already presented in *Multiphase Flow*, but we now focus on:

- how to import an external mesh with embedded geological properties (permeability) in the GMSH format (*.msh*),
- how to set up the wells.

Input file

This example is based on the XML file located at

```
src/coreComponents/physicsSolvers/fluidFlow/benchmarks/Egg/dead_oil_egg.xml
```

The mesh file corresponding to the Egg model is stored in the GEOSXDATA repository. Therefore, you must first download the GEOSXDATA repository in the same folder as the GEOSX repository to run this test case.

Note: GEOSXDATA is a separate repository in which we store large mesh files in order to keep the main GEOSX repository lightweight.

The XML file considered here follows the typical structure of the GEOSX input files:

1. *Solver*
2. *Mesh*
3. *Events*
4. *NumericalMethods*
5. *ElementRegions*
6. *Constitutive*
7. *FieldSpecifications*
8. *Outputs*
9. *Tasks*

3.2.1 Coupling the flow solver with wells

In GEOSX, the simulation of reservoir flow with wells is set up by combining three solvers listed and parameterized in the **Solvers** XML block of the input file. We introduce separately a flow solver and a well solver acting on different regions of the domain—respectively, the reservoir region and the well regions. To drive the simulation and bind these single-physics solvers, we also specify a *coupling solver* between the reservoir flow solver and the well solver. This coupling of single-physics solvers is the generic approach used in GEOSX to define multiphysics problems. It is illustrated in *Poromechanics* for a poroelastic test case.

The three solvers employed in this example are:

- the single-physics reservoir flow solver, a solver of type **CompositionalMultiphaseFVM** named `compositionalMultiphaseFlow` (more information on this solver at [Compositional Multiphase Flow Solver](#)),
- the single-physics well solver, a solver of type **CompositionalMultiphaseWell** named `compositionalMultiphaseWell` (more information on this solver at [Compositional Multiphase Well Solver](#)),
- the coupling solver that binds the two single-physics solvers above, an object of type **CompositionalMultiphaseReservoir** named `coupledFlowAndWells`.

The **Solvers** XML block is shown below. The coupling solver points to the two single-physics solvers using the attributes `flowSolverName` and `wellSolverName`. These names can be chosen by the user and are not imposed by GEOSX. The flow solver is applied to the reservoir and the well solver is applied to the wells, as specified by their respective `targetRegions` attributes.

The simulation is fully coupled and driven by the coupled solver. Therefore, the time stepping information (here, `initialDt`, but there may be other parameters used to fine-tune the time stepping strategy), the nonlinear solver parameters, and the linear solver parameters must be specified at the level of the coupling solver. There is no need to specify these parameters at the level of the single-physics solvers. Any solver information specified in the single-physics XML blocks will not be taken into account.

Note: It is worth repeating the `logLevel="1"` parameter at the level of the well solver to make sure that a notification is issued when the well control is switched (from rate control to BHP control, for instance).

Here, we instruct GEOSX to perform at most `newtonMaxIter = "10"` Newton iterations. GEOSX will adjust the time step size as follows:

- if the Newton solver converges in `dtIncIterLimit x newtonMaxIter = 5` iterations or fewer, GEOSX will double the time step size for the next time step,
- if the Newton solver converges in `dtCutIterLimit x newtonMaxIter = 8` iterations or more, GEOSX will reduce the time step size for the next time step by a factor `timestepCutFactor = 0.1`,
- if the Newton solver fails to converge in `newtonMaxIter = 10`, GEOSX will cut the time step size by a factor `timestepCutFactor = 0.1` and restart from the previous converged time step.

The maximum number of time step cuts is specified by the attribute `maxTimeStepCuts`. Note that a backtracking line search can be activated by setting the attribute `lineSearchAction` to `Attempt` or `Require`. If `lineSearchAction = "Attempt"`, we accept the nonlinear iteration even if the line search does not reduce the residual norm. If `lineSearchAction = "Require"`, we cut the time step if the line search does not reduce the residual norm.

Note: To use the linear solver options of this example, you need to ensure that GEOSX is configured to use the Hypr linear solver package.

```
<Solvers>
  <CompositionalMultiphaseReservoir
    name="coupledFlowAndWells"
    flowSolverName="compositionalMultiphaseFlow"
    wellSolverName="compositionalMultiphaseWell"
    logLevel="1"
    initialDt="1e4"
    targetRegions="{ reservoir, wellRegion1, wellRegion2, wellRegion3, wellRegion4,
↪wellRegion5, wellRegion6, wellRegion7, wellRegion8, wellRegion9, wellRegion10,
↪wellRegion11, wellRegion12 }">
    <NonlinearSolverParameters
      newtonTol="1.0e-4"
      newtonMaxIter="25"
      dtCutIterLimit="0.9"
      dtIncIterLimit="0.6"
      timestepCutFactor="0.1"
      maxTimeStepCuts="10"
      lineSearchAction="None"/>
    <LinearSolverParameters
      solverType="fgmres"
      preconditionerType="mgr"
      krylovTol="1e-4"
      krylovAdaptiveTol="1"
      krylovWeakestTol="1e-2"/>
  </CompositionalMultiphaseReservoir>

  <CompositionalMultiphaseFVM
    name="compositionalMultiphaseFlow"
    targetRegions="{ reservoir }"
    discretization="fluidTPFA"
    fluidNames="{ fluid }"
    solidNames="{ rock }"
    permeabilityNames="{ rockPerm }"
    relPermNames="{ relperm }"
    temperature="297.15"
    maxCompFractionChange="0.3"
```

(continues on next page)

(continued from previous page)

```

    logLevel="1"
    useMass="1"/>

    <CompositionalMultiphaseWell
      name="compositionalMultiphaseWell"
      targetRegions="{ wellRegion1, wellRegion2, wellRegion3, wellRegion4,
↪wellRegion5, wellRegion6, wellRegion7, wellRegion8, wellRegion9, wellRegion10,
↪wellRegion11, wellRegion12 }"
      fluidNames="{ fluid }"
      relPermNames="{ relperm }"
      wellTemperature="297.15"
      maxCompFractionChange="0.5"
      logLevel="1"
      useMass="1">
        <WellControls
          name="wellControls1"
          type="producer"
          control="BHP"
          referenceElevation="28"
          targetBHP="3.9e7"
          targetPhaseRate="1e6"
          targetPhaseName="oil"/>
        <WellControls
          name="wellControls2"
          type="producer"
          control="BHP"
          referenceElevation="28"
          targetBHP="3.9e7"
          targetPhaseRate="1e6"
          targetPhaseName="oil"/>

```

3.2.2 Mesh definition and well geometry

In the presence of wells, the **Mesh** block of the XML input file includes two parts:


- a sub-block **PAMELAMeshGenerator** defining the reservoir mesh (see [Tutorial 2: External Meshes](#) for more on this),
- a collection of sub-blocks **InternalWell** defining the geometry of the wells.

The reservoir mesh is imported from a `.msh` file that contains the mesh geometry and also includes the permeability values in the x, y, and z directions. These quantities must be specified using the metric unit system, i.e., in meters for the well geometry and square meters for the permeability field. We note that the mesh file only contains the active cells, so there is no keyword needed in the XML file to define them.

Each well is defined internally (i.e., not imported from a file) in a separate **InternalWell** XML sub-block. An **InternalWell** sub-block must point to the reservoir mesh that the well perforates using the attribute `meshName`, to the region corresponding to this well using the attribute `wellRegionName`, and to the control of this well using the attribute `wellControl`. Each block **InternalWell** must point to the reservoir mesh (using the attribute `meshName`), the corresponding well region (using the attribute `wellRegionName`), and the corresponding well control (using the attribute `wellControlName`).

Each well is defined using a vertical polyline going through the seven layers of the mesh, with a perforation in each layer. The well placement implemented here follows the pattern of the original test case. The well geometry must be specified in meters.

The location of the perforations is found internally using the linear distance along the wellbore from the top of the well, specified by the attribute `distanceFromHead`. It is the responsibility of the user to make sure that there is a perforation in the bottom cell of the well mesh otherwise an error will be thrown and the simulation will terminate. For each perforation, the well transmissibility factors employed to compute the perforation rates are calculated internally using the Peaceman formulation.



```
<Mesh>
  <PAMELAMeshGenerator
    name="mesh"
    file="../../../../../../GEOSXDATA/DataSets/Egg/egg.msh"
    fieldsToImport="{ PERM }"
    fieldNamesInGEOSX="{ rockPerm_permeability }"/>

  <InternalWell
    name="wellProducer1"
    wellRegionName="wellRegion1"
    wellControlsName="wellControls1"
    meshName="mesh"
    polylineNodeCoords="{ { 124, 340, 28 },
                          { 124, 340, 0 } }"
    polylineSegmentConn="{ { 0, 1 } }"
    radius="0.1"
    numElementsPerSegment="7">
    <Perforation
      name="producer1_perf1"
      distanceFromHead="2"/>
    <Perforation
      name="producer1_perf2"
      distanceFromHead="6"/>
    <Perforation
      name="producer1_perf3"
      distanceFromHead="10"/>
    <Perforation
      name="producer1_perf4"
      distanceFromHead="14"/>
    <Perforation
      name="producer1_perf5"
      distanceFromHead="18"/>
```

(continues on next page)

(continued from previous page)

```

    <Perforation
      name="producer1_perf6"
      distanceFromHead="22"/>
    <Perforation
      name="producer1_perf7"
      distanceFromHead="26"/>
  </InternalWell>

  <InternalWell
    name="wellProducer2"
    wellRegionName="wellRegion2"
    wellControlsName="wellControls2"
    meshName="mesh"
    polylineNodeCoords="{ { 276, 316, 28 },
                          { 276, 316, 0 } }"
    polylineSegmentConn="{ { 0, 1 } }"
    radius="0.1"
    numElementsPerSegment="7">
    <Perforation
      name="producer2_perf1"
      distanceFromHead="2"/>
    <Perforation
      name="producer2_perf2"
      distanceFromHead="6"/>
    <Perforation
      name="producer2_perf3"
      distanceFromHead="10"/>
    <Perforation
      name="producer2_perf4"
      distanceFromHead="14"/>
    <Perforation
      name="producer2_perf5"
      distanceFromHead="18"/>
    <Perforation
      name="producer2_perf6"
      distanceFromHead="22"/>
    <Perforation
      name="producer2_perf7"
      distanceFromHead="26"/>
  </InternalWell>

```

3.2.3 Events

In the **Events** XML block, we specify four types of **PeriodicEvents**.

The periodic event named `solverApplications` notifies GEOSX that the coupled solver `coupledFlowAndWells` has to be applied to its target regions (here, reservoir and wells) at every time step. The time stepping strategy has been fully defined in the **CompositionalMultiphaseReservoir** coupling block using the `initialDt` attribute and the **NonlinearSolverParameters** nested block.

We also define an output event instructing GEOSX to write out `.vtk` files at the time frequency specified by the attribute `timeFrequency`. Here, we choose to output the results using the VTK format (see [Tutorial 2: External Meshes](#) for a example that uses the Silo output file format). The `target` attribute must point to the **VTK** sub-block of the **Outputs** block (defined at the end of the XML file) by name (here, `vtkOutput`).

We define the events involved in the collection and output of the well production rates following the procedure defined

in *Tasks Manager*. The time history collection events trigger the collection of the well rates at the desired frequency, while the time history output events trigger the output of the HDF5 files containing the time series. These events point by name to the corresponding blocks of the **Tasks** and **Outputs** XML blocks, respectively. Here, these names are `wellRateCollection1` and `timeHistoryOutput1`.

```
<Events
  maxTime="1.5e7">
  <PeriodicEvent
    name="vtk"
    timeFrequency="2e6"
    target="/Outputs/vtkOutput"/>

  <PeriodicEvent
    name="timeHistoryOutput1"
    timeFrequency="1.5e7"
    target="/Outputs/timeHistoryOutput1"/>

  <PeriodicEvent
    name="timeHistoryOutput2"
    timeFrequency="1.5e7"
    target="/Outputs/timeHistoryOutput2"/>

  <PeriodicEvent
    name="timeHistoryOutput3"
    timeFrequency="1.5e7"
    target="/Outputs/timeHistoryOutput3"/>

  <PeriodicEvent
    name="timeHistoryOutput4"
    timeFrequency="1.5e7"
    target="/Outputs/timeHistoryOutput4"/>

  <PeriodicEvent
    name="solverApplications"
    maxEventDt="5e5"
    target="/Solvers/coupledFlowAndWells"/>

  <PeriodicEvent
    name="timeHistoryCollection1"
    timeFrequency="1e6"
    target="/Tasks/wellRateCollection1"/>

  <PeriodicEvent
    name="timeHistoryCollection2"
    timeFrequency="1e6"
    target="/Tasks/wellRateCollection2"/>

  <PeriodicEvent
    name="timeHistoryCollection3"
    timeFrequency="1e6"
    target="/Tasks/wellRateCollection3"/>

  <PeriodicEvent
    name="timeHistoryCollection4"
    timeFrequency="1e6"
    target="/Tasks/wellRateCollection4"/>
</Events>
```

3.2.4 Numerical methods

In the `NumericalMethods` XML block, we instruct GEOSX to use a TPFA finite-volume numerical scheme. This part is similar to the corresponding section of *Multiphase Flow*, and has been adapted to match the specifications of the Egg model.

```
<NumericalMethods>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="fluidTPFA"
      fieldName="pressure"
      coefficientName="permeability"/>
    </FiniteVolume>
  </NumericalMethods>
```

3.2.5 Reservoir and well regions

In this section of the input file, we follow the procedure already described in *Multiphase Flow* for the definition of the reservoir region with multiphase constitutive models.

We associate a **CellElementRegion** named `reservoir` to the reservoir mesh. Since we have imported a mesh with one region consisting of hexahedral cells, we must set the attribute `cellBlocks` to `DEFAULT_HEX`.

Note: If you use a name that is not `DEFAULT_HEX` for this attribute, GEOSX will throw an error at the beginning of the simulation.

We also associate a **WellElementRegion** to each well. As the **CellElementRegion**, it contains a `materialList` that must point (by name) to the constitutive models defined in the **Constitutive** XML block.

```
<ElementRegions>
  <CellElementRegion
    name="reservoir"
    cellBlocks="{ DEFAULT_HEX }"
    materialList="{ fluid, rock, relperm, rockPerm, rockPorosity, nullSolid }"/>

  <WellElementRegion
    name="wellRegion1"
    materialList="{ fluid, relperm }"/>

  <WellElementRegion
    name="wellRegion2"
    materialList="{ fluid, relperm }"/>
```

3.2.6 Constitutive models

The **CompositionalMultiphaseFVM** physics solver relies on at least four types of constitutive models listed in the **Constitutive** XML block:

- a fluid model describing the thermodynamics behavior of the fluid mixture,
- a relative permeability model,
- a rock permeability model,
- a rock porosity model.

All the parameters must be provided using the SI unit system.

This part is identical to that of *Multiphase Flow*.

```
<Constitutive>
  <DeadOilFluid
    name="fluid"
    phaseNames="{ oil, water }"
    surfaceDensities="{ 848.9, 1025.2 }"
    componentMolarWeight="{ 114e-3, 18e-3 }"
    tableFiles="{ pvdo.txt, pvtw.txt }"/>

  <BrooksCoreyRelativePermeability
    name="relperm"
    phaseNames="{ oil, water }"
    phaseMinVolumeFraction="{ 0.1, 0.2 }"
    phaseRelPermExponent="{ 4.0, 3.0 }"
    phaseRelPermMaxValue="{ 0.8, 0.75 }"/>

  <CompressibleSolidConstantPermeability
    name="rock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"/>

  <NullModel
    name="nullSolid"/>

  <PressurePorosity
    name="rockPorosity"
    defaultReferencePorosity="0.2"
    referencePressure="0.0"
    compressibility="1.0e-13"/>

  <ConstantPermeability
    name="rockPerm"
    permeabilityComponents="{ 1.0e-12, 1.0e-12, 1.0e-12 }"/>
</Constitutive>
```

3.2.7 Initial conditions

We are ready to specify the reservoir initial conditions of the problem in the **FieldSpecifications** XML block. The well variables do not have to be initialized here since they will be defined internally.

The formulation of the **CompositionalMultiphaseFVM** physics solver (documented at *Compositional Multiphase Flow Solver*) requires the definition of the initial pressure field and initial global component fractions. We define here a uniform pressure field that does not satisfy the hydrostatic equilibrium, but a hydrostatic initialization of the pressure field is possible using *Functions*. For the initialization of the global component fractions, we remind the user that their `component` attribute (here, 0 or 1) is used to point to a specific entry of the `phaseNames` attribute in the **DeadOilFluid** block.

Note that we also define the uniform porosity field here since it is not included in the mesh file imported by the **PAMELAMeshGenerator**.

```
<FieldSpecifications>
  <FieldSpecification
    name="initialPressure"
```

(continues on next page)

(continued from previous page)

```

        initialCondition="1"
        setNames="{ all }"
        objectPath="ElementRegions/reservoir/DEFAULT_HEX"
        fieldName="pressure"
        scale="4e7"/>

<FieldSpecification
  name="initialComposition_oil"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/DEFAULT_HEX"
  fieldName="globalCompFraction"
  component="0"
  scale="0.9"/>

<FieldSpecification
  name="initialComposition_water"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir/DEFAULT_HEX"
  fieldName="globalCompFraction"
  component="1"
  scale="0.1"/>
</FieldSpecifications>

```

3.2.8 Outputs

In this section, we request an output of the results in VTK format and an output of the rates for each producing well. Note that the name defined here must match the name used in the **Events** XML block to define the output frequency.

```

<Outputs>
  <VTK
    name="vtkOutput"/>

  <TimeHistory
    name="timeHistoryOutput1"
    sources="{ /Tasks/wellRateCollection1 }"
    filename="wellRateHistory1"/>

  <TimeHistory
    name="timeHistoryOutput2"
    sources="{ /Tasks/wellRateCollection2 }"
    filename="wellRateHistory2"/>

  <TimeHistory
    name="timeHistoryOutput3"
    sources="{ /Tasks/wellRateCollection3 }"
    filename="wellRateHistory3"/>

  <TimeHistory
    name="timeHistoryOutput4"
    sources="{ /Tasks/wellRateCollection4 }"
    filename="wellRateHistory4"/>
</Outputs>

```

3.2.9 Tasks

In the **Events** block, we have defined four events requesting that a task periodically collects the rate for each producing well. This task is defined here, in the **PackCollection** XML sub-block of the **Tasks** block. The task contains the path to the object on which the field to collect is registered (here, a `WellElementSubRegion`) and the name of the field (here, `wellElementMixtureConnectionRate`). The details of the history collection mechanism can be found in *Tasks Manager*.

```
<Tasks>
  <PackCollection
    name="wellRateCollection1"
    objectPath="ElementRegions/wellRegion1/wellRegion1uniqueSubRegion"
    fieldName="wellElementMixtureConnectionRate"/>

  <PackCollection
    name="wellRateCollection2"
    objectPath="ElementRegions/wellRegion2/wellRegion2uniqueSubRegion"
    fieldName="wellElementMixtureConnectionRate"/>

  <PackCollection
    name="wellRateCollection3"
    objectPath="ElementRegions/wellRegion3/wellRegion3uniqueSubRegion"
    fieldName="wellElementMixtureConnectionRate"/>

  <PackCollection
    name="wellRateCollection4"
    objectPath="ElementRegions/wellRegion4/wellRegion4uniqueSubRegion"
    fieldName="wellElementMixtureConnectionRate"/>
</Tasks>
```

All elements are now in place to run GEOSX.

3.2.10 Running GEOSX

The first few lines appearing to the console are indicating that the XML elements are read and registered correctly:

```
Adding Solver of type CompositionalMultiphaseReservoir, named coupledFlowAndWells
Adding Solver of type CompositionalMultiphaseFVM, named compositionalMultiphaseFlow
Adding Solver of type CompositionalMultiphaseWell, named compositionalMultiphaseWell
Adding Mesh: PAMELAMeshGenerator, mesh
Adding Mesh: InternalWell, wellProducer1
Adding Mesh: InternalWell, wellProducer2
Adding Mesh: InternalWell, wellProducer3
Adding Mesh: InternalWell, wellProducer4
Adding Mesh: InternalWell, wellInjector1
Adding Mesh: InternalWell, wellInjector2
Adding Mesh: InternalWell, wellInjector3
Adding Mesh: InternalWell, wellInjector4
Adding Mesh: InternalWell, wellInjector5
Adding Mesh: InternalWell, wellInjector6
Adding Mesh: InternalWell, wellInjector7
Adding Mesh: InternalWell, wellInjector8
Adding Event: PeriodicEvent, solverApplications
Adding Event: PeriodicEvent, vtk
Adding Output: VTK, vtkOutput
Adding Object CellElementRegion named reservoir from ObjectManager::Catalog.
```

(continues on next page)

(continued from previous page)

```

Adding Object WellElementRegion named wellRegion1 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion2 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion3 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion4 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion5 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion6 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion7 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion8 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion9 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion10 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion11 from ObjectManager::Catalog.
Adding Object WellElementRegion named wellRegion12 from ObjectManager::Catalog.

```

This is followed by the creation of the 18553 hexahedral cells of the imported mesh:

```

0 >>> *****
0 >>> PAMELA Library Import tool
0 >>> *****
0 >>> GMSH FORMAT IDENTIFIED
0 >>> *** Importing Gmsh mesh format...
0 >>> Reading nodes...
0 >>> Done0
0 >>> Reading elements...
0 >>> Reading element data...
0 >>> Number of nodes = 22227
0 >>> Number of triangles = 0
0 >>> Number of quadrilaterals = 0
0 >>> Number of tetrahedra = 0
0 >>> Number of hexahedra = 18553
0 >>> Number of pyramids = 0
0 >>> Number of prisms = 0
0 >>> *** Done
0 >>> *** Creating Polygons from Polyhedra...
0 >>> 59205 polygons have been created
0 >>> *** Done
0 >>> *** Perform partitioning...
0 >>> TRIVIAL partitioning...
0 >>> Ghost elements...
0 >>> Clean mesh...
0 >>> *** Done...
0 >>> Clean Adjacency...
0 >>> *** Done...

```

At this point, we are done with the case set-up and the code steps into the execution of the simulation itself:

```

Time: 0s, dt:10000s, Cycle: 0

Attempt: 0, NewtonIter: 0
( Rfluid ) = ( 9.39e+01 ) ;      ( R ) = ( 1.06e+02 ) ;
Attempt: 0, NewtonIter: 1
( Rfluid ) = ( 2.14e+00 ) ;      ( R ) = ( 2.20e+00 ) ;
Last LinSolve(iter,res) = ( 2, 3.62e-03 ) ;
Attempt: 0, NewtonIter: 2
( Rfluid ) = ( 3.23e-01 ) ;      ( R ) = ( 3.37e-01 ) ;
Last LinSolve(iter,res) = ( 4, 1.82e-03 ) ;
Attempt: 0, NewtonIter: 3
( Rfluid ) = ( 1.07e-02 ) ;      ( R ) = ( 1.16e-02 ) ;

```

(continues on next page)

(continued from previous page)


```
Last LinSolve(iter,res) = ( 2, 6.13e-03 ) ;
Attempt: 0, NewtonIter: 4
( Rfluid ) = (7.46e-05) ;      ( R ) = ( 7.50e-05 ) ;
Last LinSolve(iter,res) = ( 3, 5.09e-03 ) ;

coupledFlowAndWells: Newton solver converged in less than 15 iterations, time-step_
↪required will be doubled.
```

3.2.11 Visualization

A file compatible with Paraview is produced in this example. It is found in the output folder, and usually has the extension *.pvd*. More details about this file format can be found [here](#). We can load this file into Paraview directly and visualize results:

We have instructed GEOSX to output the time series of rates for each producer. The data contained in the corresponding hdf5 files can be extracted and plotted as shown below.



3.2.12 To go further

Feedback on this example

This concludes the example on setting up a Dead-Oil simulation in the Egg model. For any feedback on this example, please submit a [GitHub issue on the project's GitHub page](#).

For more details

- A complete description of the reservoir flow solver is found here: *Compositional Multiphase Flow Solver*.
- The well solver is description at *Compositional Multiphase Well Solver*.
- The available constitutive models are listed at *Constitutive Models*.

3.3 CO₂ Injection

Context

In this example, we show how to set up a multiphase simulation of CO₂ injection.

Objectives

At the end of this example you will know:

- how to set up a CO₂ injection scenario with a well,
- how to run a case using MPI-parallelism.

Input file

The XML file for this test case is located at :

```
src/coreComponents/physicsSolvers/multiphysics/integratedTests/SimpleCo2InjTutorial.  
↪xml
```

This mesh is a simple internally generated regular grid (50 x 1 x 150). A single CO₂ injection well is at the center of the reservoir.

The XML file considered here follows the typical structure of the GEOSX input files:

1. *Solver*
2. *Mesh*
3. *Events*
4. *NumericalMethods*
5. *ElementRegions*
6. *Constitutive*
7. *FieldSpecifications*
8. *Outputs*
9. *Tasks*

3.3.1 Multiphase flow and well solvers

Let us inspect the **Solver** XML tags. They consist of three blocks **CompositionalMultiphaseFVM**, **Compositional-MultiphaseWell** and **CompositionalMultiphaseReservoir**, which are respectively handling the solution from multiphase flow in the reservoir, multiphase flow in the wells, and coupling between those two parts.

```

<Solvers>
  <CompositionalMultiphaseReservoir
    name="coupledFlowAndWells"
    flowSolverName="compositionalMultiphaseFlow"
    wellSolverName="compositionalMultiphaseWell"
    logLevel="1"
    initialDt="1e2"
    targetRegions="{ reservoir, wellRegion }">
    <NonlinearSolverParameters
      newtonTol="1.0e-4"
      lineSearchAction="Attempt"
      maxTimeStepCuts="10"
      newtonMaxIter="40"/>
    <LinearSolverParameters
      solverType="fgmres"
      preconditionerType="mgr"
      krylovTol="1e-4"
      krylovAdaptiveTol="1"
      krylovWeakestTol="1e-2"/>
    </CompositionalMultiphaseReservoir>

  <CompositionalMultiphaseFVM
    name="compositionalMultiphaseFlow"
    targetRegions="{ reservoir }"
    discretization="fluidTPFA"
    fluidNames="{ fluid }"
    solidNames="{ rock }"
    permeabilityNames="{ rockPerm }"
    relPermNames="{ relperm }"
    temperature="368.15"
    maxCompFractionChange="0.2"
    logLevel="1"
    useMass="1"/>

  <CompositionalMultiphaseWell
    name="compositionalMultiphaseWell"
    targetRegions="{ wellRegion }"
    fluidNames="{ fluid }"
    relPermNames="{ relperm }"
    wellTemperature="368.15"
    logLevel="1"
    useMass="1">
    <WellControls
      name="wellControls"
      type="injector"
      control="totalVolRate"
      referenceElevation="7500"
      targetBHP="1e8"
      targetTotalRate="5e-2"
      injectionStream="{ 0.995, 0.005 }"/>
    </CompositionalMultiphaseWell>
</Solvers>

```

In the **CompositionalMultiphaseFVM** (*Compositional Multiphase Flow Solver*), a classical multiphase compositional solver with a TPFA discretization is described. The flow solver definition includes a list of names to point to fluid constitutive data through `fluidNames`, solid constitutive data through `solidNames`, permeability data through `permeabilityNames` and relative permeability constitutive data through `relPermNames` attributes.

The **CompositionalMultiphaseWell** (*Compositional Multiphase Well Solver*) consists of wellbore specifications (see *Multiphase Flow with Wells* for detailed example). As its reservoir counterpart, it includes references to fluid and relative permeability models, but also defines a **WellControls** sub-tag. This sub-tag specifies the CO₂ injector control mode: the well is initially rate-controlled, with a rate specified in `targetRate` and a maximum pressure specified in `targetBHP`. The injector-specific attribute, `injectionStream`, describes the composition of the injected mixture (here, pure CO₂).

The **CompositionalMultiphaseReservoir** coupling section describes the binding between those two previous elements (see *Poromechanics* for detailed example on coupling physics in GEOSX). In addition to being bound to the previously described blocks through `flowSolverName` and `wellSolverName` sub-tags, it contains the `initialDt` starting time-step size value and defines the **NonlinearSolverParameters** and **LinearSolverParameters** that are used to control Newton-loop and linear solver behaviors (see *Linear Solvers* for a detailed description of linear solver attributes).

Note: To use the linear solver options of this example, you need to ensure that GEOSX is configured to use the Hypr linear solver package.

3.3.2 Mesh and well geometry

In this example, the **Mesh** tag is used to generate the reservoir mesh internally (*Tutorial 1: First Steps*). The internal generation of well is defined with the **InternalWell** sub-tag. Apart from the name identifier attribute and their `wellRegionName` (*ElementRegions*) and `wellControlsName` (*Solver*) binding attributes, `polylineNodeCoords` and `polylineSegmentConn` attributes are used to define the path of the wellbore and connections between its nodes. The `numElementsPerSegment` discretizes the wellbore segments while the `radius` attribute specifies the wellbore radius (*Multiphase Flow with Wells* for details on wells). Once the wellbore is defined and discretized, the position of **Perforations** is defined using the linear distance from the head of the wellbore (`distanceFromHead`).

```
<Mesh>
  <InternalMesh
    name="cartesianMesh"
    elementTypes="{ C3D8 }"
    xCoords="{ 0, 1000 }"
    yCoords="{ 450, 550 }"
    zCoords="{ 6500, 7700 }"
    nx="{ 50 }"
    ny="{ 1 }"
    nz="{ 150 }"
    cellBlockNames="{ cellBlock }"/>

  <InternalWell
    name="wellInjector1"
    wellRegionName="wellRegion"
    wellControlsName="wellControls"
    meshName="cartesianMesh"
    polylineNodeCoords="{ { 500.0, 500.0, 6600.00 },
                          { 500.0, 500.0, 6650.00 } }"
    polylineSegmentConn="{ { 0, 1 } }"
    radius="0.1"
    numElementsPerSegment="2">
    <Perforation
      name="injector1_perf1"
      distanceFromHead="45"/>
```

(continues on next page)

(continued from previous page)

```

</InternalWell>
</Mesh>

```

Note: It is the responsibility of the user to make sure that there is a perforation in the bottom cell of the well mesh, otherwise an error will be thrown and the simulation will terminate.

3.3.3 Events

The solver is applied as a periodic event whose target is referred to as `coupledFlowAndWells` nametag. Using the `maxEventDt` attribute, we specify a max time step size of 5×10^6 seconds.

The output event triggers a VTK output every 10^7 seconds, constraining the solver schedule to match exactly these dates. The output path to data is specified as a `target` of this **PeriodicEvent**.

Another periodic event is defined under the name `restarts`. It consists of saved checkpoints every 5×10^7 seconds, whose physical output folder name is defined under the **Output** tag.

Finally, the time history collection and output events are used to trigger the mechanisms involved in the generation of a time series of well pressure (see the procedure outlined in *Tasks Manager*, and the example in *Multiphase Flow with Wells*).

```

<Events
  maxTime="2e8">

  <PeriodicEvent
    name="outputs"
    timeFrequency="1e7"
    targetExactTimestep="1"
    target="/Outputs/simpleReservoirViz"/>

  <PeriodicEvent
    name="restarts"
    timeFrequency="5e7"
    targetExactTimestep="1"
    target="/Outputs/restartOutput"/>

  <PeriodicEvent
    name="timeHistoryCollection"
    timeFrequency="1e7"
    targetExactTimestep="1"
    target="/Tasks/wellPressureCollection" />

  <PeriodicEvent
    name="timeHistoryOutput"
    timeFrequency="2e8"
    targetExactTimestep="1"
    target="/Outputs/timeHistoryOutput" />

  <PeriodicEvent
    name="solverApplications"
    maxEventDt="5e5"
    target="/Solvers/coupledFlowAndWells"/>

</Events>

```


3.3.4 Numerical methods

The **TwoPointFluxApproximation** is chosen for the fluid equation discretization. The tag specifies:

- A primary field to solve for as `fieldName`. For a flow problem, this field is pressure.
- A set of target regions in `targetRegions`.
- A `coefficientName` pointing to the field used for TPFA transmissibilities construction.
- A `coefficientModelNames` used to specify the permeability constitutive model(s).

```
<NumericalMethods>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="fluidTPFA"
      targetRegions="{ reservoir }"
      fieldName="pressure"
      coefficientName="permeability"
      coefficientModelNames="{ rockPerm }"/>
    </FiniteVolume>
  </NumericalMethods>
```

3.3.5 Element regions

We define a **CellElementRegion** pointing to the cell block defining the reservoir mesh, and a **WellElementRegion** for the well. The two regions contain a list of constitutive model names.

```
<ElementRegions>
  <CellElementRegion
    name="reservoir"
    cellBlocks="{ cellBlock }"
    materialList="{ fluid, rock, relperm, rockPerm, rockPorosity, nullSolid }"/>

  <WellElementRegion
    name="wellRegion"
    materialList="{ fluid, relperm, rockPerm }"/>
</ElementRegions>
```

3.3.6 Constitutive laws

Under the **Constitutive** tag, four items can be found:

- **CO2BrineFluid** : this tag defines phase names, component molar weights, and fluid behaviors such as CO₂ solubility in brine and viscosity/density dependencies on pressure and temperature.
- **PressurePorosity** : this tag contains all the data needed to model rock compressibility.
- **BrooksCoreyRelativePermeability** : this tag defines the relative permeability model for each phase, its end-point values, residual volume fractions (saturations), and the Corey exponents.
- **ConstantPermeability** : this tag defines the permeability model that is set to a simple constant diagonal tensor, whose values are defined in `permeabilityComponent`. Note that these values will be overwritten by the permeability field imported in **FieldSpecifications**.

```

<Constitutive>
  <CO2BrineFluid
    name="fluid"
    phaseNames="{ gas, water }"
    componentNames="{ co2, water }"
    componentMolarWeight="{ 44e-3, 18e-3 }"
    phasePVTParaFiles="{ pvt_tables/pvtgas.txt, pvt_tables/pvtliquid.txt }"
    flashModelParaFile="pvt_tables/co2flash.txt"/>

  <CompressibleSolidConstantPermeability
    name="rock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"/>

  <NullModel
    name="nullSolid"/>

  <PressurePorosity
    name="rockPorosity"
    defaultReferencePorosity="0.1"
    referencePressure="1.0e7"
    compressibility="4.5e-10"/>

  <BrooksCoreyRelativePermeability
    name="relperm"
    phaseNames="{ gas, water }"
    phaseMinVolumeFraction="{ 0.05, 0.30 }"
    phaseRelPermExponent="{ 2.0, 2.0 }"
    phaseRelPermMaxValue="{ 1.0, 1.0 }"/>

  <ConstantPermeability
    name="rockPerm"
    permeabilityComponents="{ 1.0e-17, 1.0e-17, 3.0e-17 }"/>

</Constitutive>

```

The PVT data specified by **CO2BrineFluid** is set to model the behavior of the CO₂-brine system as a function of pressure, temperature, and salinity. We currently rely on a two-phase, two-component (CO₂ and H₂O) model in which salinity is a constant parameter in space and in time. The model is described in detail in *CO₂-brine model*. The model definition requires three text files:

In *co2flash.txt*, we define the CO₂ solubility model used to compute the amount of CO₂ dissolved in the brine phase as a function of pressure (in Pascal), temperature (in Kelvin), and salinity (in units of molality):

```
FlashModel CO2Solubility 1e6 1.5e7 5e4 367.15 369.15 1 0
```

The first keyword is an identifier for the model type (here, a flash model). It is followed by the model name. Then, the lower, upper, and step increment values for pressure and temperature ranges are specified. The trailing 0 defines a zero-salinity in the model. Note that the water component is not allowed to evaporate into the CO₂-rich phase.

The *pvtgas.txt* and *pvtliquid.txt* files define the models used to compute the density and viscosity of the two phases, as follows:

```
DensityFun SpanWagnerCO2Density 1e6 1.5e7 5e4 94 96 1
ViscosityFun FenghourCO2Viscosity 1e6 1.5e7 5e4 94 96
```

```
DensityFun BrineCO2Density 1e6 1.5e7 5e4 94 96 1 0
ViscosityFun BrineViscosity 0
```

In these files, the first keyword of each line is an identifier for the model type (either a density or a viscosity model). It is followed by the model name. Then, the lower, upper, and step increment values for pressure and temperature ranges are specified. The trailing 0 for BrineCO2Density and BrineViscosity entry is the salinity of the brine, set to zero.

Note: It is the responsibility of the user to make sure that the pressure and temperature values encountered in the simulation (in the reservoir and in the well) are within the bounds specified in the PVT files. GEOSX will not throw an error if a value outside these bounds is encountered, but the (nonlinear) behavior of the simulation and the quality of the results will likely be negatively impacted.

3.3.7 Property specification

The **FieldSpecifications** tag is used to declare fields such as directional permeability, reference porosity, initial pressure, and compositions. Here, these fields are homogeneous, except for the permeability field that is taken as an heterogeneous log-normally distributed field and specified in **Functions** as in *Tutorial 3: Regions and Property Specifications*.

```
<FieldSpecifications>
  <FieldSpecification
    name="permx"
    initialCondition="1"
    component="0"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rockPerm_permeability"
    scale="1e-15"
    functionName="permxFunc"/>

  <FieldSpecification
    name="permy"
    initialCondition="1"
    component="1"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rockPerm_permeability"
    scale="1e-15"
    functionName="permyFunc"/>

  <FieldSpecification
    name="permz"
    initialCondition="1"
    component="2"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rockPerm_permeability"
    scale="3e-15"
    functionName="permzFunc"/>

  <FieldSpecification
    name="initialPressure"
```

(continues on next page)

(continued from previous page)

```

        initialCondition="1"
        setNames="{ all }"
        objectPath="ElementRegions/reservoir"
        fieldName="pressure"
        scale="1.25e7"/>

<FieldSpecification
  name="initialComposition_co2"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir"
  fieldName="globalCompFraction"
  component="0"
  scale="0.0"/>

<FieldSpecification
  name="initialComposition_water"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions/reservoir"
  fieldName="globalCompFraction"
  component="1"
  scale="1.0"/>
</FieldSpecifications>

```

Note: In this case, we are using the same permeability field (*perm.geos*) for all the directions. Note also that the `fieldName` values are set to *rockPerm_permeability* to access the permeability field handled as a **Constitutive** law. These permeability values will overwrite the values already set in the **Constitutive** block.

3.3.8 Output

The **Outputs** XML tag is used to write visualization, restart, and time history files.

Here, we write visualization files in a format natively readable by Paraview under the tag **VTK**. A **Restart** tag is also be specified. In conjunction with a **PeriodicEvent**, a restart file allows to resume computations from a set of checkpoints in time. Finally, we require an output of the well pressure history using the **TimeHistory** tag.

```

<Outputs>
  <VTK
    name="simpleReservoirViz"/>

  <Restart
    name="restartOutput"/>

  <TimeHistory
    name="timeHistoryOutput"
    sources="{/Tasks/wellPressureCollection}"
    filename="wellPressureHistory" />

</Outputs>

```

3.3.9 Tasks

In the **Events** block, we have defined an event requesting that a task periodically collects the pressure at the well. This task is defined here, in the **PackCollection** XML sub-block of the **Tasks** block. The task contains the path to the object on which the field to collect is registered (here, a `WellElementSubRegion`) and the name of the field (here, `pressure`). The details of the history collection mechanism can be found in [Tasks Manager](#).

```
<Tasks>
  <PackCollection
    name="wellPressureCollection"
    objectPath="ElementRegions/wellRegion/wellRegionuniqueSubRegion"
    fieldName="pressure" />
</Tasks>
```

3.3.10 Running GEOSX

The simulation can be launched with 4 cores using MPI-parallelism:

```
mpirun -np 4 geosx -i SimpleCo2InjTutorial.xml -x 1 -y 1 -z 4
```

A restart from a checkpoint file *SimpleCo2InjTutorial_restart_000000024.root* is always available thanks to the following command line :

```
mpirun -np 4 geosx -i SimpleCo2InjTutorial.xml -r SimpleCo2InjTutorial_restart_
↪000000024 -x 1 -y 1 -z 4
```


The output then shows the loading of HDF5 restart files by each core.

```
Loading restart file SimpleCo2InjTutorial_restart_000000024
Rank 0: rankFilePattern = SimpleCo2InjTutorial_restart_000000024/rank_%07d.hdf5
Rank 0: Reading in restart file at SimpleCo2InjTutorial_restart_000000024/rank_
↪0000000.hdf5
Rank 1: Reading in restart file at SimpleCo2InjTutorial_restart_000000024/rank_
↪0000001.hdf5
Rank 3: Reading in restart file at SimpleCo2InjTutorial_restart_000000024/rank_
↪0000003.hdf5
Rank 2: Reading in restart file at SimpleCo2InjTutorial_restart_000000024/rank_
↪0000002.hdf5
```


and the simulation restarts from this point in time.

3.3.11 Visualization

Using Paraview, we can observe the CO₂ plume moving upward under buoyancy effects and forming a gas cap at the top of the domain,



The heterogeneous values of the log permeability field can also be visualized in Paraview as shown below:



3.3.12 To go further

Feedback on this example

This concludes the CO₂ injection field case example. For any feedback on this example, please submit a [GitHub issue](#) on the project's [GitHub page](#).

For more details

- A complete description of the reservoir flow solver is found here: *Compositional Multiphase Flow Solver*.
- The well solver is described at *Compositional Multiphase Well Solver*.
- The available fluid constitutive models are listed at *Fluid Models*.

3.4 Poromechanics

Context

In this example, we use a coupled solver to solve a poroelastic Terzaghi-type problem, a classic benchmark in poroelasticity. We do so by coupling a single phase flow solver with a small-strain Lagrangian mechanics solver.

Objectives

At the end of this example you will know:

- how to use multiple solvers for poromechanical problems,
- how to define finite elements and finite volume numerical methods.

Input file

This example uses no external input files and everything required is contained within a single GEOSX input file. The xml input file for this test case is located at:

```
src/coreComponents/physicsSolvers/multiphysics/integratedTests/PoroElastic_Terzaghi_
↪ FIM.xml
```

3.4.1 Description of the case

We simulate the consolidation of a poroelastic fluid-saturated column of height L having unit cross-section. The column is instantaneously loaded at time $t = 0$ s with a constant compressive traction w applied on the face highlighted in red in the figure below. Only the loaded face is permeable to fluid flow, with the remaining parts of the boundary subject to roller constraints and impervious.

Fig. 3.1: Sketch of the setup for Terzaghi's problem.

GEOSX will calculate displacement and pressure fields along the column as a function of time. We will use the analytical solution for pressure to check the accuracy of the solution obtained with GEOSX, namely

$$p(x, t) = \frac{4}{\pi} p_0 \sum_{m=0}^{\infty} \frac{1}{2m+1} \exp \left[-\frac{(2m+1)^2 \pi^2 c_c t}{4L^2} \right] \sin \left[\frac{(2m+1)\pi x}{2L} \right],$$

where $p_0 = \frac{b}{K_v S_e + b^2} |w|$ is the initial pressure, constant throughout the column, and $c_c = \frac{\kappa}{\mu} \frac{K_v}{K_v S_e + b^2}$ is the consolidation coefficient (or diffusion coefficient), with

- b Biot's coefficient
- $K_v = \frac{E(1-\nu)}{(1+\nu)(1-2\nu)}$ the uniaxial bulk modulus, E Young's modulus, and ν Poisson's ratio
- $S_e = \frac{(b-\phi)(1-b)}{K} + \phi c_f$ the constrained specific storage coefficient, ϕ porosity, $K = \frac{E}{3(1-2\nu)}$ the bulk modulus, and c_f the fluid compressibility
- κ the isotropic permeability
- μ the fluid viscosity

The characteristic consolidation time of the system is defined as $t_c = \frac{L^2}{c_c}$. Knowledge of t_c is useful for choosing appropriately the timestep sizes that are used in the discrete model.

3.4.2 Coupled solvers

GEOSX is a multi-physics tool. Different combinations of physics solvers available in the code can be applied in different regions of the mesh at different moments of the simulation. The `XML Solvers` tag is used to list and parameterize these solvers.

We define and characterize each single-physics solver separately. Then, we define a *coupling solver* between these single-physics solvers as another, separate, solver. This approach allows for generality and flexibility in our multi-physics resolutions. The order in which these solver specifications is done is not important. It is important, though, to instantiate each single-physics solver with meaningful names. The names given to these single-physics solver instances will be used to recognize them and create the coupling.

To define a poromechanical coupling, we will effectively define three solvers:

- the single-physics flow solver, a solver of type `SinglePhaseFVM` called here `SinglePhaseFlowSolver` (more information on these solvers at [Singlephase Flow Solver](#)),
- the small-stress Lagrangian mechanics solver, a solver of type `SolidMechanicsLagrangianSSLE` called here `LinearElasticitySolver` (more information here: [Solid Mechanics Solver](#)),
- the coupling solver that will bind the two single-physics solvers above, an object of type `SinglePhasePoromechanics` called here `PoroelasticitySolver` (more information at [Poromechanics Solver](#)).

Note that the name attribute of these solvers is chosen by the user and is not imposed by GEOSX.

The two single-physics solvers are parameterized as explained in their respective documentation.

Let us focus our attention on the coupling solver. This solver (`PoroelasticitySolver`) uses a set of attributes that specifically describe the coupling for a poromechanical framework. For instance, we must point this solver to the correct fluid solver (here: `SinglePhaseFlowSolver`), the correct solid solver (here: `LinearElasticitySolver`). Now that these two solvers are tied together inside the coupling solver, we have a coupled multiphysics problem defined. More parameters are required to characterize a coupling. Here, we specify the discretization method (`FE1`, defined further in the input file), and the target regions (here, we only have one, `Domain`).

```
<SinglePhasePoromechanics
  name="PoroelasticitySolver"
  solidSolverName="LinearElasticitySolver"
  fluidSolverName="SinglePhaseFlowSolver"
  porousMaterialNames="{ porousRock }"
  logLevel="1"
  discretization="FE1"
  targetRegions="{ Domain }">
  <LinearSolverParameters
    directParallel="0"/>
</SinglePhasePoromechanics>

<SolidMechanicsLagrangianSSLE
  name="LinearElasticitySolver"
  timeIntegrationOption="QuasiStatic"
  logLevel="1"
  discretization="FE1"
  targetRegions="{ Domain }"
  solidMaterialNames="{ skeleton }"/>

<SinglePhaseFVM
  name="SinglePhaseFlowSolver"
  logLevel="1"
  discretization="singlePhaseTPFA"
  targetRegions="{ Domain }"
  fluidNames="{ fluid }"
  solidNames="{ porousRock }"
  permeabilityNames="{ skeletonPerm }"/>
</Solvers>
```

3.4.3 Multiphysics numerical methods

Numerical methods in multiphysics settings are similar to single physics numerical methods. All can be defined under the same `NumericalMethods` XML tag. In this problem, we use finite volume for flow and finite elements for solid mechanics. Both methods require additional parameterization attributes to be defined here.

As we have seen before, the coupling solver and the solid mechanics solver require the specification of a discretization method called `FE1`. This discretization method is defined here as a finite element method using linear basis functions and Gaussian quadrature rules. For more information on defining finite elements numerical schemes, please see the dedicated `FiniteElement` section.

The finite volume method requires the specification of a discretization scheme. Here, we use a two-point flux approximation as described in the dedicated documentation (found here: `FiniteVolume`).

```
<NumericalMethods>
  <FiniteElements>
```

(continues on next page)

(continued from previous page)

```

<FiniteElementSpace
  name="FE1"
  order="1"/>
</FiniteElements>

<FiniteVolume>
  <TwoPointFluxApproximation
    name="singlePhaseTPFA"
    fieldName="pressure"
    coefficientName="permeability"
    coefficientModelNames="{ skeletonPerm }"/>
  </FiniteVolume>
</NumericalMethods>

```

3.4.4 Mesh, material properties, and boundary conditions

Last, let us take a closer look at the geometry of this simple problem. We use the internal mesh generator to create a beam-like mesh, with one single element along the Y and Z axes, and 21 elements along the X axis. All the elements are hexahedral elements (C3D8) of the same dimension (1x1x1 meters).

```

<Mesh>
  <InternalMesh
    name="mesh1"
    elementTypes="{ C3D8 }"
    xCoords="{ 0, 10 }"
    yCoords="{ 0, 1 }"
    zCoords="{ 0, 1 }"
    nx="{ 25 }"
    ny="{ 1 }"
    nz="{ 1 }"
    cellBlockNames="{ cb1 }"/>
  </Mesh>

```

The parameters used in the simulation are summarized in the following table.

Symbol	Parameter	Units	Value
E	Young's modulus	[Pa]	10^4
ν	Poisson's ratio	[-]	0.2
b	Biot's coefficient	[-]	1.0
ϕ	Porosity	[-]	0.3
ρ_f	Fluid density	[kg/m ³]	1.0
c_f	Fluid compressibility	[Pa ⁻¹]	0.0
κ	Permeability	[m ²]	10^{-4}
μ	Fluid viscosity	[Pa s]	1.0
$ w $	Applied compression	[Pa]	1.0
L	Column length	[m]	10.0

Material properties and boundary conditions are specified in the `Constitutive` and `FieldSpecifications` sections. For such set of parameters we have $p_0 = 1.0$ Pa, $c_c = 1.111 \text{ m}^2 \text{ s}^{-1}$, and $t_c = 90$ s. Therefore, as shown in the `Events` section, we run this simulation for 90 seconds.

3.4.5 Running GEOSX

To run the case, use the following command:

```
path/to/geosx -i src/coreComponents/physicsSolvers/multiphysics/
integratedTests/PoroElastic_Terzaghi_FIM.xml
```


Here, we see for instance the `RSolid` and `RFluid` at a representative timestep (residual values for solid and fluid mechanics solvers, respectively)

```
Attempt: 0, NewtonIter: 0
( RSolid ) = ( 5.00e-01 ) ;      ( Rsolid, Rfluid ) = ( 5.00e-01, 0.00e+00 )
( R ) = ( 5.00e-01 ) ;
Attempt: 0, NewtonIter: 1
( RSolid ) = ( 4.26e-16 ) ;      ( Rsolid, Rfluid ) = ( 4.26e-16, 4.22e-17 )
( R ) = ( 4.28e-16 ) ;
```

As expected, since we are dealing with a linear problem, the fully implicit solver converges in a single iteration.

3.4.6 Inspecting results

This plot compares the analytical pressure solution (continuous lines) at selected times with the numerical solution (markers).



3.4.7 To go further

Feedback on this example

This concludes the poroelastic example. For any feedback on this example, please submit a [GitHub issue](#) on the project's [GitHub](#) page.

For more details

- More on poroelastic multiphysics solvers, please see *Poromechanics Solver*.
- More on numerical methods, please see *Numerical Methods*.
- More on functions, please see *Functions*.

3.5 Hydraulic Fracturing

Context

In this example, we use a fully coupled hydrofracture solver from GEOSX to solve for the propagation of a single fracture within a reservoir with heterogeneous in-situ properties. Advanced xml features will be used throughout the example.

Objectives

At the end of this example you will know:

- how to use multiple solvers for hydraulic fracturing problems,
- how to specify pre-existing fractures and where new fractures can develop,
- how to construct a mesh with bias,
- how to specify heterogeneous in-situ properties and initial conditions,
- how to use parameters, symbolic math, and units in xml files.

Input files


This example uses a set of input files and table files located at:

```
examples/hydraulicFracturing/heterogeneousInSituProperties
```

Note: because these files use the advanced xml features, they must be preprocessed using the `geosx_xml_tools` package. To install `geosx_xml_tools`, see *Advanced XML Features (geosx_xml_tools)*

3.5.1 Description of the case

Here, our goal is to demonstrate how hydraulic fractures are modeled in a typical environment. The in-situ properties and initial conditions are based upon a randomly generated, fractal, 1D layer-cake model.



The inputs for this case are contained inside a case-specific (`heterogeneousInSitu_singleFracture.xml`) and base (`heterogeneousInSitu_base.xml`) XML files. The `tables` directory contains the pre-constructed geologic model. This example will first focus on the case-specific input file, which contains the key parameter definitions, then consider the base xml file.

3.5.2 Included: including external xml files

At the head of the case-specific xml file is a block that will instruct GEOSX to include an external file. In our case, this points to the base hydraulic fracturing input file.

```
<Included>
  <File
    name="./heterogeneousInSitu_base.xml"/>
  </Included>
```

3.5.3 Parameters: defining variables to be used throughout the file

The `Parameters` block defines a series of variables that can be used throughout the input file. These variables allow a given input file to be easily understood and/or modified for a specific environment, even by non-expert users. Parameters are specified in pairs of names and values. The names should only contain alphanumeric characters and underlines. The values can contain any type (strings, doubles, etc.).

Parameters can be used throughout the input file (or an included input file) by placing them in-between dollar signs. Barring any circular-definition errors, parameters can be used within other parameters. For example, see the parameter `mu_upscaled`. The value of this parameter is a symbolic expression, which is denoted by the surrounding back-ticks, and is dependent upon two other parameters. During pre-processing, `geosx_xml_tools` will substitute the parameter definitions, and evaluate the symbolic expression using a python-derived syntax.

A number of the input parameters include optional unit definitions, which are denoted by the square brackets following a value. For example, the parameter `t_max` is used to set the maximum time for the simulation to 20 minutes.

```
<Parameters>
  <!-- Use the swarm upscaling law -->
  <Parameter
    name="Nperf"
    value="5"/>

  <Parameter
    name="Nswarm"
    value="8"/>

  <Parameter
    name="mu_init"
    value="0.005"/>

  <Parameter
    name="K_init"
    value="1e6"/>

  <Parameter
    name="mu_upscaled"
    value="`$mu_init*$($Nswarm**2)`"/>

  <Parameter
    name="K_upscaled"
    value="`$K_init*$($Nswarm**0.5)`"/>

  <Parameter
    name="ContactStiffness"
    value="1e10"/>

  <!-- Event timing -->
  <Parameter
    name="t_init_a"
    value="1 [min]"/>

  <Parameter
    name="dt_max_a"
    value="10 [s]"/>

  <Parameter
    name="t_init_b"
    value="2 [min]"/>

  <Parameter
    name="dt_max_b"
    value="2 [s]"/>

  <Parameter
    name="t_max"
```

(continues on next page)

(continued from previous page)

```

    value="20 [min]"/>

    <Parameter
      name="dt_max_c"
      value="4 [s]"/>

    <!-- Etc. -->
    <Parameter
      name="table_root"
      value="./tables"/>

    <Parameter
      name="t_allocation"
      value="28 [min]"/>
  </Parameters>

```

3.5.4 Mesh with biased boundaries

The mesh block for this example uses a biased mesh along the simulation boundaries to reduce the size of the problem, while maintaining the desired spatial extents. For a given element region with bias, the left-hand element will be $x\%$ smaller and the right-hand element will be $x\%$ larger than the average element size. Along the x -axis of the mesh, we select a value of zero for the first region to indicate that we want a uniform-sized mesh, and we select a bias of -0.6 for the second region to indicate that we want the element size to smoothly increase in size as it moves in the $+x$ direction. The other dimensions of the mesh follow a similar pattern.

```

<Mesh>
  <InternalMesh
    name="mesh1"
    xCoords="{ 0, 200, 250 }"
    yCoords="{ -100, 0, 100 }"
    zCoords="{ -150, -100, 0, 100, 150 }"
    nx="{ 50, 5 }"
    ny="{ 10, 10 }"
    nz="{ 5, 25, 25, 5 }"
    xBias="{ 0, -0.6 }"
    yBias="{ 0.6, -0.6 }"
    zBias="{ 0.6, 0, 0, -0.6 }"
    cellBlockNames="{ cb1 }"
    elementTypes="{ C3D8 }"/>
  </InternalMesh>
</Mesh>

```

3.5.5 Defining a fracture nodeset

For this example, we want to propagate a single hydraulic fracture along the plane defined by $y = 0$. To achieve this, we need to define three nodesets:

- `source_a`: The location where we want to inject fluid. Typically, we want this to be a single face in the x - z plane.
- `perf_a`: This is the initial fracture for the simulation. This nodeset needs to be at least two-faces wide in the x - z plane (to represent the fracture at least one internal node needs to be open).
- `fracturable_a`: This is the set of faces where we will allow the fracture to grow. For a problem where we expect the fracture to curve out of the plane defined by $y = 0$, this could be replaced.

```

<Geometry>
  <Box
    name="source_a"
    xMin="{ -0.1, -0.1, -0.1 }"
    xMax="{ 4.1, 0.1, 4.1 }"/>

  <Box
    name="perf_a"
    xMin="{ -4.1, -0.1, -4.1 }"
    xMax="{ 4.1, 0.1, 4.1 }"/>

  <ThickPlane
    name="fracturable_a"
    normal="{ 0, 1, 0 }"
    origin="{ 0, 0, 0 }"
    thickness="0.1"/>
</Geometry>

```

3.5.6 Boundary conditions

The boundary conditions for this problem are defined in the case-specific and the base xml files. The case specific block includes four instructions:

- `frac`: this marks the initial perforation.
- `separableFace`: this marks the set of faces that are allowed to break during the simulation.
- `waterDensity`: this initializes the fluid in the perforation.
- `sourceTerm`: this instructs the code to inject fluid into the `source_a` nodeset. Note the usage of the symbolic expression and parameters in the scale. This boundary condition is also driven by a function, which we will define later.

```

<FieldSpecifications>
  <!-- Fracture-related nodesets -->
  <FieldSpecification
    name="frac"
    fieldName="ruptureState"
    initialCondition="1"
    objectPath="faceManager"
    scale="1"
    setNames="{ perf_a }"/>

  <FieldSpecification
    name="separableFace"
    fieldName="isFaceSeparable"
    initialCondition="1"
    objectPath="faceManager"
    scale="1"
    setNames="{ fracturable_a }"/>

  <!-- Fluid Initial Conditions -->
  <FieldSpecification
    name="waterDensity"
    fieldName="water_density"
    initialCondition="1"
    objectPath="ElementRegions"

```

(continues on next page)

(continued from previous page)

```

    scale="1000"
    setNames="{ perf_a }"/>

    <!-- Fluid BC's -->
    <!-- Note: the units of the source flux BC are in kg/s not m3/s -->
    <SourceFlux
      name="sourceTerm"
      objectPath="ElementRegions/Fracture"
      scale="-1000.0/$Nperf$"
      functionName="flow_rate"
      setNames="{ source_a }"/>
  </FieldSpecifications>

```

The base block includes instructions to set the initial in-situ properties and stresses. It is also used to specify the external mechanical boundaries on the system. In this example, we are using roller-boundary conditions (zero normal-displacement). Depending upon how close they are to the fracture, they can significantly affect its growth. Therefore, it is important to test whether the size of the model is large enough to avoid this.

```

<FieldSpecifications>
  <FieldSpecification
    name="bulk_modulus"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rock_BulkModulus"
    functionName="bulk_modulus"
    scale="1.0"/>

  <FieldSpecification
    name="shear_modulus"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rock_ShearModulus"
    functionName="shear_modulus"
    scale="1.0"/>

  <FieldSpecification
    name="sigma_xx"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rock_stress"
    component="0"
    functionName="sigma_xx"
    scale="1.0"/>

  <FieldSpecification
    name="sigma_yy"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions"
    fieldName="rock_stress"
    component="1"
    functionName="sigma_yy"
    scale="1.0"/>

```

(continues on next page)

(continued from previous page)

```

<FieldSpecification
  name="sigma_zz"
  initialCondition="1"
  setNames="{ all }"
  objectPath="ElementRegions"
  fieldName="rock_stress"
  component="2"
  functionName="sigma_zz"
  scale="1.0"/>

<!-- Mechanical BC's -->
<FieldSpecification
  name="x_constraint"
  component="0"
  fieldName="TotalDisplacement"
  objectPath="nodeManager"
  scale="0.0"
  setNames="{ xneg, xpos }"/>

<FieldSpecification
  name="y_constraint"
  component="1"
  fieldName="TotalDisplacement"
  objectPath="nodeManager"
  scale="0.0"
  setNames="{ yneg, ypos }"/>

<FieldSpecification
  name="z_constraint"
  component="2"
  fieldName="TotalDisplacement"
  objectPath="nodeManager"
  scale="0.0"
  setNames="{ zneg, zpos }"/>
</FieldSpecifications>

```

3.5.7 Coupled hydraulic fracturing solver

The Solvers block is located in the base xml file. Note that the `gravityVector` attribute indicates that we are applying gravity in the z-direction in this problem.

Similar to other coupled physics solvers, the Hydrofracture solver is specified in three parts:

- Hydrofracture: this is the primary solver, which will be called by the event manager. Two of its key attributes are the names of the dependent solid and fluid solvers.
- SolidMechanicsLagrangianSSLE: this is the solid mechanics solver.
- SinglePhaseFVM: this is the fluid solver.

The final solver present in this example is the SurfaceGenerator, which manages how faces in the model break.

```

<Solvers
  gravityVector="{ 0.0, 0.0, -9.81 }">
  <Hydrofracture
    name="hydrofracture"

```

(continues on next page)

(continued from previous page)

```

solidSolverName="lagsolve"
fluidSolverName="SinglePhaseFlow"
surfaceGeneratorName="SurfaceGen"
porousMaterialNames="{ fractureFilling }"
couplingTypeOption="FIM"
logLevel="1"
discretization="FE1"
targetRegions="{ Domain, Fracture }"
contactRelationName="fractureContact">
<NonlinearSolverParameters
  newtonTol="1.0e-5"
  newtonMaxIter="50"
  lineSearchMaxCuts="10"/>
<LinearSolverParameters
  solverType="gmres"
  preconditionerType="amg"/>
</Hydrofracture>

<SolidMechanicsLagrangianSSLE
  name="lagsolve"
  timeIntegrationOption="QuasiStatic"
  discretization="FE1"
  targetRegions="{ Domain }"
  solidMaterialNames="{ rock }"
  contactRelationName="fractureContact">
<NonlinearSolverParameters
  newtonTol="1.0e-6"/>
<LinearSolverParameters
  solverType="gmres"
  krylovTol="1.0e-10"/>
</SolidMechanicsLagrangianSSLE>

<SinglePhaseFVM
  name="SinglePhaseFlow"
  discretization="singlePhaseTPFA"
  targetRegions="{ Fracture }"
  fluidNames="{ water }"
  solidNames="{ fractureFilling }"
  permeabilityNames="{ fracPerm }">
<NonlinearSolverParameters
  newtonTol="1.0e-5"
  newtonMaxIter="10"/>
<LinearSolverParameters
  solverType="gmres"
  krylovTol="1.0e-12"/>
</SinglePhaseFVM>

<SurfaceGenerator
  name="SurfaceGen"
  targetRegions="{ Domain }"
  solidMaterialNames="{ rock }"
  rockToughness="$K_upscaled$"
  nodeBasedSIF="1"
  mpiCommOrder="1"/>
</Solvers>

```

3.5.8 Events

Rather than explicitly specify the desired timestep behavior, this example uses a flexible approach for timestep-ping. The hydrofracture solver is applied in three segments, where `maxEventDt` indicates the maximum allowable timestep:

- `solverApplications_a`: this corresponds to the problem initialization, where we request `dt_max_a=10`.
- `solverApplications_b`: this corresponds to the period of initial fluid injection, where we request `dt_max_b=2`.
- `solverApplications_c`: this corresponds to rest of the problem where we request `dt_max_c=4`.

Depending upon how well the solution converges, the timestep may be smaller than the maximum requested value. Other key events in this problem include:

- `preFracture`: this calls the surface generator at the beginning of the problem and helps to initialize the fracture.
- `outputs`: this produces output silo files.
- `restarts`: this is a `HaltEvent`, which tracks the external clock. When the runtime exceeds the specified value (here `$t_allocation$=28` minutes), the code will call the target (which writes a restart file) and instruct the code to exit.

```
<Events
  maxTime="$t_max$">
  <!-- Generate the initial fractures -->
  <SoloEvent
    name="preFracture"
    target="/Solvers/SurfaceGen"/>

  <!-- Primary outputs -->
  <PeriodicEvent
    name="outputs"
    target="/Outputs/siloOutput"
    timeFrequency="1 [min]"/>

  <!-- Apply the hydrofracture solver, limiting the timesteps during certain time_
  <!-- intervals -->
  <PeriodicEvent
    name="solverApplications_a"
    maxEventDt="$dt_max_a$"
    endTime="$t_init_a$"
    target="/Solvers/hydrofracture"/>

  <PeriodicEvent
    name="solverApplications_b"
    maxEventDt="$dt_max_b$"
    beginTime="$t_init_a$"
    endTime="$t_init_b$"
    target="/Solvers/hydrofracture"/>

  <PeriodicEvent
    name="solverApplications_c"
    maxEventDt="$dt_max_c$"
    beginTime="$t_init_b$"
    target="/Solvers/hydrofracture"/>

  <!-- Watch the wall-clock, write a restart, and exit gracefully if necessary -->
  <HaltEvent
    name="restarts"
```

(continues on next page)

(continued from previous page)

```

    maxRuntime="$t_allocation$"
    target="/Outputs/restartOutput"/>
</Events>

```

3.5.9 Functions to set in-situ properties

The function definitions are in the base xml file, and rely upon the files in the tables directory. The functions in this example include the flow rate over time, the in-situ principal stress, and the bulk/shear moduli of the rock. Note the use of the `table_root` parameter, which contains the root path to the table files.

The `flow_rate` TableFunction is an example of a 1D function. It has a single input, which is time. The table is defined using a single `coordinateFile`:

```

0.00000e+00
6.00000e+01
1.20000e+02
3.72000e+03
3.78000e+03
1.00000e+09

```

And a single `voxelFile`:

```

0.00000e+00
0.00000e+00
5.00000e-02
5.00000e-02
0.00000e+00
0.00000e+00

```

Given the specified linear interpolation method, these values define a simple trapezoidal function. Note: since this is a 1D table, these values could alternately be given within the xml file using the `coordinates` and `values` attributes.

The `sigma_xx` TableFunction is an example of a 3D function. It uses `elementCenter` as its input, which is a vector. It is specified using a set of three coordinate files (one for each axis), and a single voxel file. The geologic model in this example is a layer-cake, which was randomly generated, so the size of the x and y axes are 1. The interpolation method used here is `upper`, so the values in the table indicate those at the top of each layer.

```

<Functions>
  <!-- Pumping Schedule -->
  <TableFunction
    name="flow_rate"
    inputVarNames="{ time }"
    coordinateFiles="{ $table_root$/flowRate_time.csv }"
    voxelFile="$table_root$/flowRate.csv"/>

  <!-- Geologic Model -->
  <TableFunction
    name="sigma_xx"
    inputVarNames="{ elementCenter }"
    coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
    voxelFile="$table_root$/sigma_xx.csv"
    interpolation="upper"/>

  <TableFunction
    name="sigma_yy"

```

(continues on next page)

(continued from previous page)

```

inputVarNames="{ elementCenter }"
coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
voxelFile="$table_root$/sigma_yy.csv"
interpolation="upper"/>

<TableFunction
  name="sigma_zz"
  inputVarNames="{ elementCenter }"
  coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
  voxelFile="$table_root$/sigma_zz.csv"
  interpolation="upper"/>

<TableFunction
  name="init_pressure"
  inputVarNames="{ elementCenter }"
  coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
  voxelFile="$table_root$/porePressure.csv"
  interpolation="upper"/>

<TableFunction
  name="bulk_modulus"
  inputVarNames="{ elementCenter }"
  coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
  voxelFile="$table_root$/bulkModulus.csv"
  interpolation="upper"/>

<TableFunction
  name="shear_modulus"
  inputVarNames="{ elementCenter }"
  coordinateFiles="{ $table_root$/x.csv, $table_root$/y.csv, $table_root$/z.csv }"
  voxelFile="$table_root$/shearModulus.csv"
  interpolation="upper"/>
</Functions>

```

3.5.10 Preprocessing the input file

Because we are using advanced xml features in this example, the input file must be pre-processed using `geosx_xml_tools`. To build the final input file `hydrofracture_processed.xml`, run the following:

```

geosx_bin_dir/preprocess_xml examples/hydraulicFracturing/
heterogeneousInSituProperties/heterogeneousInSitu_singleFracture.xml -o
hydrofracture_processed.xml

```

3.5.11 Running GEOSX

This is a moderate-sized example, so it is recommended to run this problem in parallel. For example, this will run the code on the debug partition using a total of 36 cores:

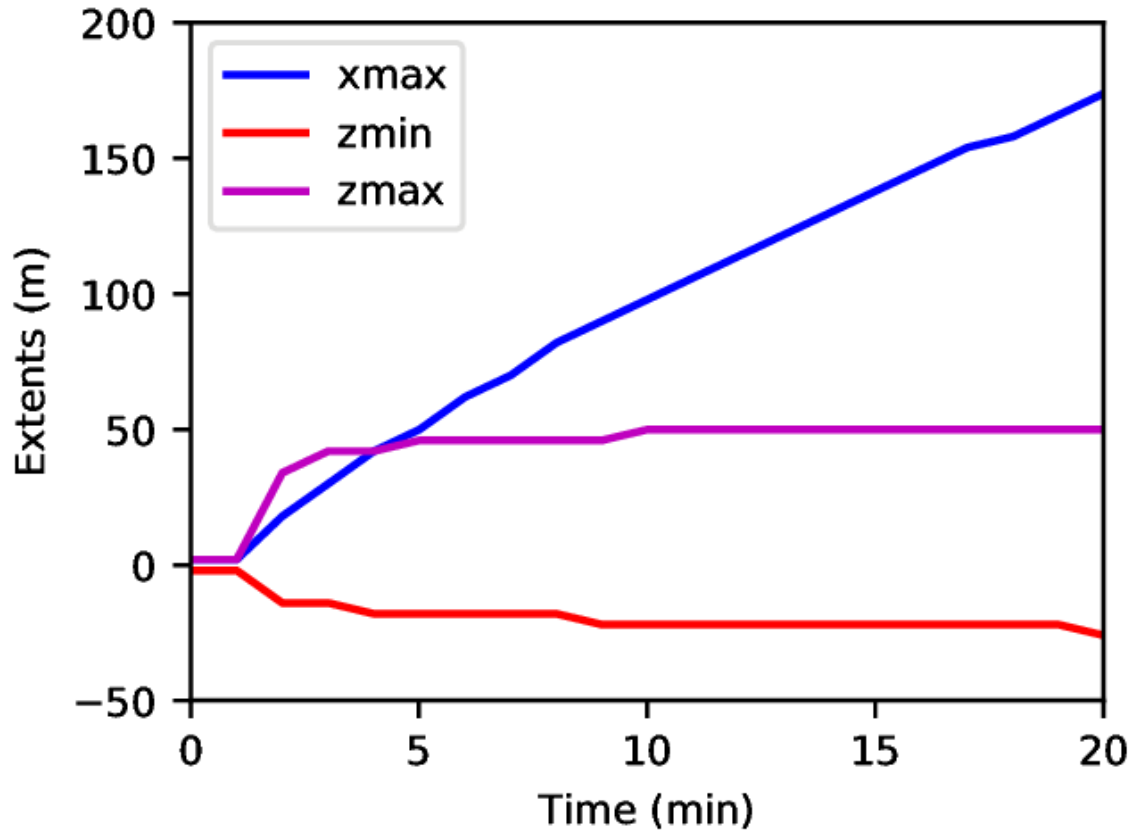
```

srun -n 36 -ppdebug geosx_bin_dir/geosx -i hydrofracture_processed.xml -x 6 -y
2 -z 3

```


3.5.12 Inspecting results


In the above example, we requested silo-format output files every minute. We can therefore import these into VisIt or python and visualize the outcome. The following figure shows the extents of the generated fracture over time:



Because we did not explicitly specify any fracture barriers in this example, the fracture dimensions are controlled by the in-situ stresses. During the first couple of minutes of growth, the fracture quickly reaches its maximum/minimum height, which corresponds to a region of low in-situ minimum stress.

The following figures show the aperture and pressure of the hydraulic fracture (near the source) over time:





3.5.13 To go further

Feedback on this example

This concludes the hydraulic fracturing example. For any feedback on this example, please submit a [GitHub issue](#) on the project's [GitHub page](#).

For more details

- More on advanced xml features, please see *Advanced XML Features (geosx_xml_tools)*.
- More on functions, please see *Functions*.
- More on biased meshes, please see *Mesh Bias*.

4.1 Validation and Verification Studies

4.1.1 Sneddon's Problem

Objectives

At the end of this example you will know:

- how to define embedded fractures in the porous domain,
- how to use the SolidMechanicsEmbeddedFractures solver to solve mechanics problems with embedded fractures.

Input file

This example uses no external input files and everything required is contained within a single GEOSX input file. The xml input file for this test case is located at:

```
src/coreComponents/physicsSolvers/solidMechanics/benchmarks/Sneddon-Validation.xml
```

Description of the case

We compute the displacement field induced by the presence of a pressurized fracture, of length L_f , in a porous medium.

GEOSX will calculate the displacement field in the porous matrix and the displacement jump at the fracture surface. We will use the analytical solution for the fracture aperture, w_n (normal component of the jump) to, i.e.

$$w_n(s) = \frac{4(1 - \nu^2)p_f}{E} \sqrt{\frac{L_f^2}{4} - s^2}$$

where - E is the Young's modulus - ν is the Poisson's ratio - p_f is the fracture pressure - s is the local fracture coordinate in $[-\frac{L_f}{2}, \frac{L_f}{2}]$

All inputs for this case are contained inside a single XML file. In this example, we focus our attention on the `Solvers` tags, the `ElementRegions` tags and the `Geometry` tags.

Embedded fractures mechanics solver

To define a mechanics solver capable of including embedded fractures, we will define two solvers:

- a `SolidMechanicsEmbeddedFractures` solver, called `mechSolve`
- a small-strain Lagrangian mechanics solver, of type `SolidMechanicsLagrangianSSLE` called here `matrixSolver` (see: *Solid Mechanics Solver*)

Note that the name attribute of these solvers is chosen by the user and is not imposed by GEOSX. It is important to make sure that the `solidSolverName` specified in the embedded fractures solver corresponds to the small-strain Lagrangian solver used in the matrix.

The two single-physics solvers are parameterized as explained in their respective documentation, each with their own tolerances, verbosity levels, target regions, and other solver-specific attributes.

Additionally, we need to specify another solver of type, `EmbeddedSurfaceGenerator`, which is used to discretize the fracture planes.

```
<SolidMechanicsEmbeddedFractures
  name="mechSolve"
  targetRegions="{ Domain, Fracture }"
  fractureRegionName="Fracture"
  initialDt="10"
  solidSolverName="matrixSolver"
  contactRelationName="fractureContact"
  logLevel="1">
  <NonlinearSolverParameters
    newtonTol="1.0e-6"
    newtonMaxIter="2"
    maxTimeStepCuts="1"/>
  <LinearSolverParameters
    directParallel="0"/>
</SolidMechanicsEmbeddedFractures>

<SolidMechanicsLagrangianSSLE
  name="matrixSolver"
  timeIntegrationOption="QuasiStatic"
  logLevel="1"
  discretization="FE1"
  targetRegions="{ Domain }"
  solidMaterialNames="{ rock }"/>

<EmbeddedSurfaceGenerator
  name="SurfaceGenerator"
  solidMaterialNames="{ rock }"
  targetRegions="{ Domain, Fracture }"
  fractureRegion="Fracture"
  logLevel="1"/>
</Solvers>
```

Events

For this problem we will add two events defining solver applications:

- an event specifying the execution of the `EmbeddedSurfaceGenerator` to generate the fracture elements.
- a periodic even specifying the execution of the embedded fractures solver.

```
<Events
  maxTime="10">
  <SoloEvent
    name="preFracture"
    target="/Solvers/SurfaceGenerator"/>

  <PeriodicEvent
    name="solverApplications"
    forceDt="10"
    target="/Solvers/mechSolve"/>
```

Mesh, material properties, and boundary conditions

Last, let us take a closer look at the geometry of this simple problem. We use the internal mesh generator to create a large domain ($1000\text{ m} \times 1001\text{ m} \times 1\text{ m}$), with one single element along the Z axes, 420 elements along the X axis and 121 elements along the Y axis. All the elements are hexahedral elements (C3D8) and that refinement is performed around the fracture.

```
<Mesh>
  <InternalMesh
    name="mesh1"
    elementTypes="{ C3D8 }"
    xCoords="{ 0, 400, 600, 1000 }"
    yCoords="{ 0, 400, 601, 1001 }"
    zCoords="{ 0, 100 }"
    nx="{ 10, 400, 10 }"
    ny="{ 10, 101, 10 }"
    nz="{ 1 }"
    cellBlockNames="{ cb1 }"/>
  </Mesh>
```

The parameters used in the simulation are summarized in the following table.

Symbol	Parameter	Units	Value
E	Young's modulus	[Pa]	10^4
ν	Poisson's ratio	[-]	0.2
L_f	Fracture length	[m]	20
p_f	Fracture pressure	[Pa]	10^5

Material properties and boundary conditions are specified in the `Constitutive` and `FieldSpecifications` sections.

Adding an embedded fracture

```
<Geometry>
  <BoundedPlane
    name="FracturePlane"
    normal="{ 0, 1, 0 }"
    origin="{ 500, 500.5, 50 }"
    lengthVector="{ 1, 0, 0 }"
```

(continues on next page)

(continued from previous page)

```
widthVector="{ 0, 0, 1 }"  
dimensions="{ 20, 100 }"/>  
</Geometry>
```

Running GEOSX

To run the case, use the following command:

```
path/to/geosx -i src/coreComponents/physicsSolvers/solidMechanics/benchmarks/  
Sneddon-Validation.xml
```

Inspecting results

This plot compares the analytical pressure solution (continuous lines) at selected times with the numerical solution (markers).

To go further

Feedback on this example

This concludes the Sneddon example. For any feedback on this example, please submit a [GitHub issue on the project's GitHub page](#).

4.1.2 Elasto-Plastic Near-Well Deformation

Context

The main goal of this example is to learn how to use the internal wellbore mesh generator and an elasto-plastic model to handle wellbore problems in GEOSX. The Extended Drucker-Prager model (see *Model: Extended Drucker-Prager*) is applied to solve for elastoplastic deformation within the vicinity of a vertical wellbore. For the presented example, an analytical solution is employed to verify the accuracy of the numerical results. The resulting model can be used as a base for more complex analysis (e.g., wellbore drilling, fluid injection and storage scenarios).

Objectives

At the end of this example you will know:

- how to construct meshes for wellbore problems with the internal mesh generator,
- how to specify initial and boundary conditions, such as in-situ stresses and variation of traction at the wellbore wall,
- how to use a plastic model for mechanical problems in the near wellbore region.

Input file

This example uses no external input files and everything required is contained within a single xml file that is located at:

```
examples/plasticity/DruckerPrager_Verification.xml
```

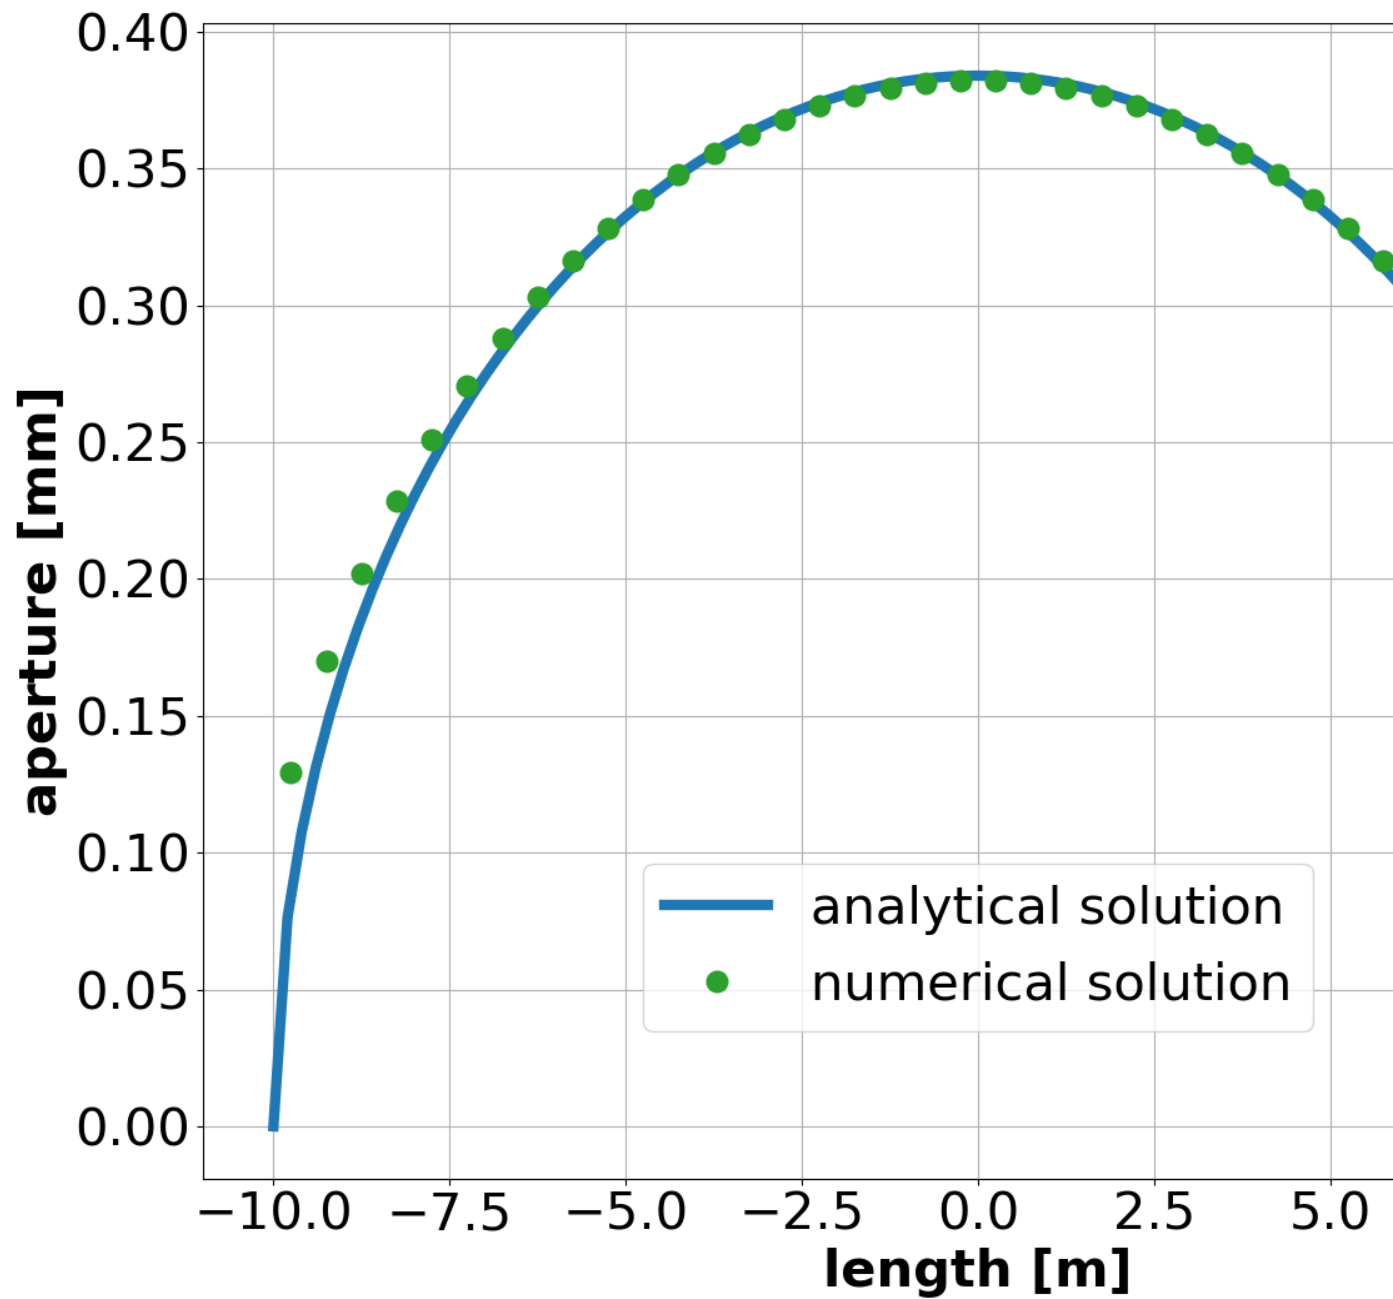


Fig. 4.1: Comparing GEOSX results with analytical solution

Description of the case

We simulate a drained wellbore problem subjected to isotropic horizontal stress (σ_h) and vertical stress (σ_v). By lowering the wellbore supporting pressure (P_w), the wellbore contracts, and the reservoir rock experiences elastoplastic deformation. A plastic zone develops in the near wellbore region, as shown below.




Fig. 4.2: Sketch of the wellbore problem (Chen and Abousleiman, 2017)

To simulate this phenomenon, the strain hardening Extended Drucker-Prager model with an associated plastic flow rule in GEOSX is used in this example. Displacement and stress fields around the wellbore are numerically calculated. These numerical predictions are then compared with the corresponding analytical solutions (Chen and Abousleiman, 2017) from the literature.

All inputs for this case are contained inside a single XML file. In this example, we focus our attention on the Mesh tags, the Constitutive tags, and the FieldSpecifications tags.

Mesh

Following figure shows the generated mesh that is used for solving this 3D wellbore problem

Let us take a closer look at the geometry of this wellbore problem. We use the internal mesh generator `InternalWellbore` to create a rock domain ($10\text{ m} \times 5\text{ m} \times 2\text{ m}$), with a wellbore of initial radius equal to 0.1 m . Coordinates of trajectory defines the wellbore trajectory, which represents a vertical well in this example. By turning on `autoSpaceRadialElems="{ 1 }"`, the internal mesh generator automatically sets number and spacing of elements in the radial direction, which overrides the values of `nr`. With `useCartesianOuterBoundary="0"`, a Cartesian aligned outer boundary on the outer block is enforced. In




Fig. 4.3: Generated mesh for the wellbore problem

this way, a structured three-dimensional mesh is created with $100 \times 80 \times 2$ elements in the radial, tangential and z directions, respectively. All the elements are eight-node hexahedral elements (C3D8) and refinement is performed to conform with the wellbore geometry. This mesh is defined as a cell block with the name `cb1`.

```
<Mesh>
  <InternalWellbore
    name="mesh1"
    elementTypes="{ C3D8 }"
    radius="{ 0.1, 5.0 }"
    theta="{ 0, 180 }"
    zCoords="{ -1, 1 }"
    nr="{ 40 }"
    nt="{ 80 }"
    nz="{ 2 }"
    trajectory="{ { 0.0, 0.0, -1.0 },
                  { 0.0, 0.0, 1.0 } }"
    autoSpaceRadialElems="{ 1 }"
    useCartesianOuterBoundary="0"
    cellBlockNames="{ cb1 }"/>
  </InternalWellbore>
</Mesh>
```

Solid mechanics solver

For the drained wellbore problem, the pore pressure variation is omitted and can be subtracted from the analysis. Therefore, we just need to define a solid mechanics solver, which is called `mechanicsSolver`. This solid mechanics solver (see *Solid Mechanics Solver*) is based on the Lagrangian finite element formulation. The problem is run as `QuasiStatic` without considering inertial effects. The computational domain is discretized by `FE1`, which is defined in the `NumericalMethods` section. The material is named as `rock`, whose mechanical properties are specified in the `Constitutive` section.

```
<Solvers
  gravityVector="{ 0.0, 0.0, 0.0 }">
  <SolidMechanics_LagrangianFEM
    name="mechanicsSolver"
    timeIntegrationOption="QuasiStatic"
    logLevel="1"
    discretization="FE1"
    targetRegions="{ Omega }"
    solidMaterialNames="{ rock }">
    <NonlinearSolverParameters
      newtonTol="1.0e-5"
      newtonMaxIter="15"/>
    </SolidMechanics_LagrangianFEM>
  </Solvers>
```

Constitutive laws

For this drained wellbore problem, we simulate the elastoplastic deformation caused by wellbore contraction. A homogeneous domain with one solid material is assumed, whose mechanical properties are specified in the Constitutive section:

```
<Constitutive>
  <ExtendedDruckerPrager
    name="rock"
    defaultDensity="2700"
    defaultBulkModulus="5.0e8"
    defaultShearModulus="3.0e8"
    defaultCohesion="0.0"
    defaultInitialFrictionAngle="15.27"
    defaultResidualFrictionAngle="23.05"
    defaultDilationRatio="1.0"
    defaultHardening="0.01"/>
  </Constitutive>
```

Recall that in the `SolidMechanics_LagrangianFEM` section, `rock` is designated as the material in the computational domain. Here, Extended Drucker Prager model `ExtendedDruckerPrager` is used to simulate the elastoplastic behavior of rock. As for the material parameters, `defaultInitialFrictionAngle`, `defaultResidualFrictionAngle` and `defaultCohesion` denote the initial friction angle, the residual friction angle, and cohesion, respectively, as defined by the Mohr-Coulomb failure envelope. As the residual friction angle `defaultResidualFrictionAngle` is larger than the initial one `defaultInitialFrictionAngle`, a strain hardening model is adopted, whose hardening rate is given as `defaultHardening="0.01"`. If the residual friction angle is set to be less than the initial one, strain weakening will take place. Setting `defaultDilationRatio="1.0"` corresponds to an associated flow rule. The constitutive parameters such as the density, the bulk modulus, and the shear modulus are specified in the International System of Units.

Initial and boundary conditions

The next step is to specify fields, including:

- The initial value (the in-situ stresses and traction at the wellbore wall have to be initialized)
- The boundary conditions (the reduction of wellbore pressure and constraints of the outer boundaries have to be set)

In this example, we need to specify isotropic horizontal stress ($\sigma_h = -11.25$ MPa) and vertical stress ($\sigma_v = -15.0$ MPa). To reach equilibrium, a compressive traction $P_w = -11.25$ MPa is instantaneously applied at the wellbore wall `rneg` at time $t = 0$ s, which will then be gradually reduced to a lower value (-2.0 MPa) to let wellbore contract. The remaining parts of the outer boundaries are subjected to roller constraints. These boundary conditions are set up through the `FieldSpecifications` section.

```
<FieldSpecifications>
  <FieldSpecification
    name="stressXX"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="0"
    scale="-11250000.0"/>

  <FieldSpecification
    name="stressYY"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="1"
    scale="-11250000.0"/>

  <FieldSpecification
    name="stressZZ"
    initialCondition="1"
    setNames="{ all }"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="2"
    scale="-15000000.0"/>

  <Traction
    name="ExternalLoad"
    setNames="{ rneg }"
    objectPath="faceManager"
    scale="-11.25e6"
    tractionType="normal"
    functionName="timeFunction"/>

  <FieldSpecification
    name="xconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="0"
    scale="0.0"
    setNames="{ xneg, xpos }"/>

  <FieldSpecification
    name="yconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="1"
    scale="0.0"
    setNames="{ tneg, tpos, ypos }"/>
```

(continues on next page)

(continued from previous page)

```

<FieldSpecification
  name="zconstraint"
  objectPath="nodeManager"
  fieldName="TotalDisplacement"
  component="2"
  scale="0.0"
  setNames="{ zneg, zpos }"/>
</FieldSpecifications>

```

With `tractionType="normal"`, traction is applied to the wellbore wall `rneg` as a pressure specified from the product of scale `scale="-11.25e6"` and the outward face normal. A table function `timeFunction` is used to define the time-dependent traction `ExternalLoad`. The coordinates and values form a time-magnitude pair for the loading time history. In this case, the loading magnitude decreases linearly as the time evolves.

```

<Functions>
  <TableFunction
    name="timeFunction"
    inputVarNames="{ time }"
    coordinates="{ 0.0, 1.0 }"
    values="{ 1.0, 0.1778 }"/>
</Functions>

```

You may note :

- All initial value fields must have `initialCondition` field set to 1;
- The `setName` field points to the previously defined box to apply the fields;
- `nodeManager` and `faceManager` in the `objectPath` indicate that the boundary conditions are applied to the element nodes and faces, respectively;
- `fieldName` is the name of the field registered in GEOSX;
- Component 0, 1, and 2 refer to the x, y, and z direction, respectively;
- And the non-zero values given by `Scale` indicate the magnitude of the loading;
- Some shorthand, such as `xneg` and `xpos`, are used as the locations where the boundary conditions are applied in the computational domain. For instance, `xneg` means the portion of the computational domain located at the left-most in the x-axis, while `xpos` refers to the portion located at the right-most area in the x-axis. Similar shorthand include `ypos`, `yneg`, `zpos`, and `zneg`;
- The mud pressure loading has a negative value due to the negative sign convention for compressive stress in GEOSX.

The parameters used in the simulation are summarized in the following table.

Symbol	Parameter	Unit	Value
K	Bulk modulus	[MPa]	500
G	Shear Modulus	[MPa]	300
C	Cohesion	[MPa]	0.0
ϕ_i	Initial Friction Angle	[degree]	15.27
ϕ_r	Residual Friction Angle	[degree]	23.05
c_h	Hardening Rate	[-]	0.01
σ_h	Horizontal Stress	[MPa]	-11.25
σ_v	Vertical Stress	[MPa]	-15.0
a_0	Initial Well Radius	[m]	0.1
P_w	Mud Pressure	[MPa]	-2.0

Inspecting results

In the above example, we requested silo-format output files. We can therefore import these into VisIt and use python scripts to visualize the outcome. Below figure shows the comparisons between the numerical predictions (marks) and the corresponding analytical solutions (solid curves) with respect to the distributions of normal stress components, stress path, the supporting wellbore pressure and wellbore size. It is clear that the GEOSX predictions are in excellent agreement with the analytical results.

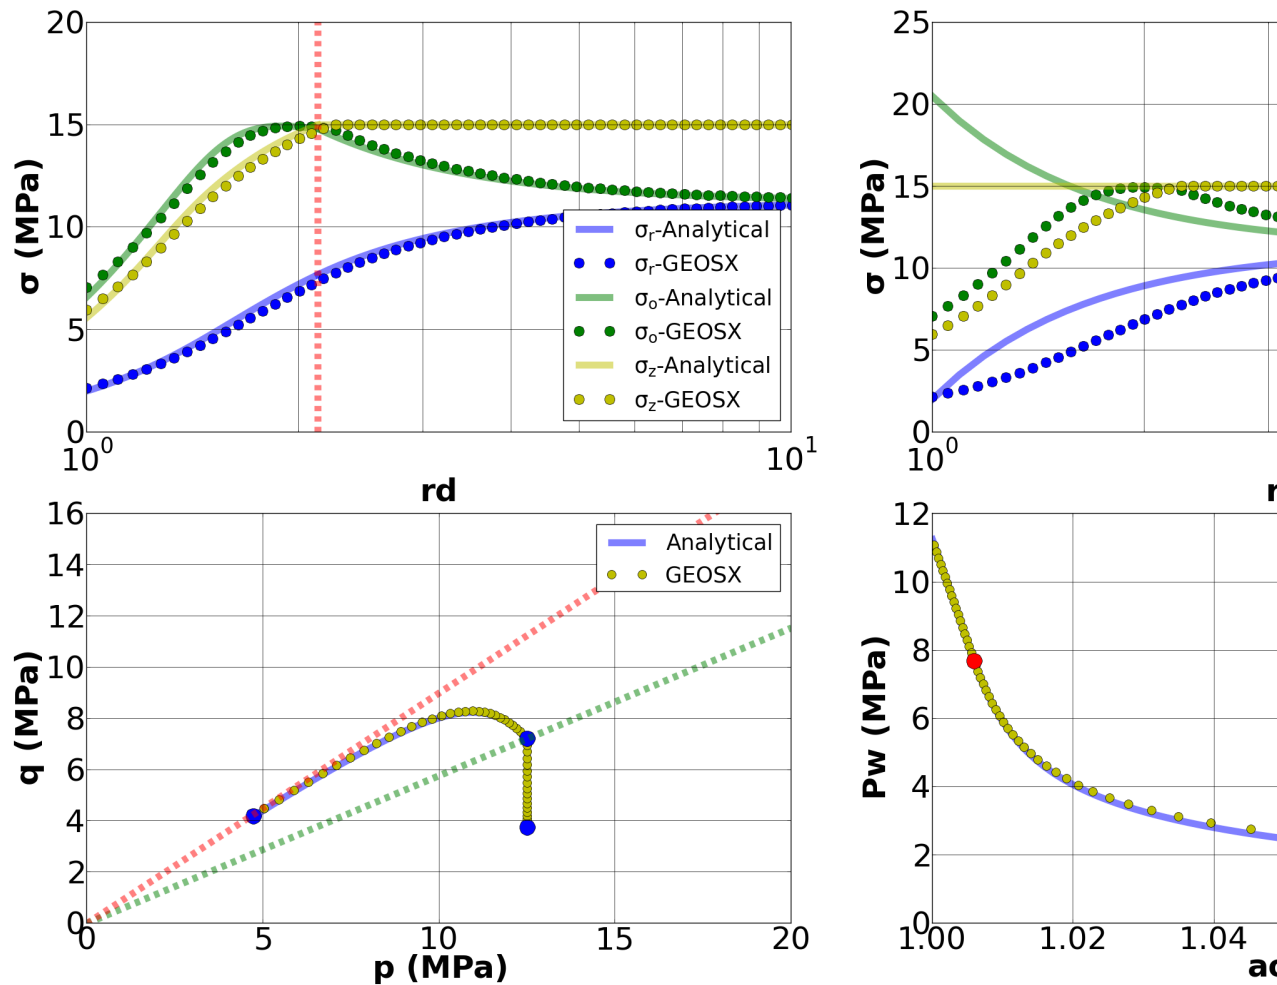


Fig. 4.4: Comparing GEOSX results with analytical solutions

For the same wellbore problem, using different constitutive models (plastic vs. elastic), obviously, distinct differences in rock deformation and distribution of resultant stresses is also observed and highlighted.

To go further

[Feedback on this example](#)

This concludes the example on Plasticity Model for Wellbore Problems. For any feedback on this example, please submit a [GitHub issue on the project's GitHub page](#).

For more details

- More on plasticity models, please see *Model: Extended Drucker-Prager*.
- More on functions, please see *Functions*.

4.1.3 Cased Elastic Wellbore Problem

Problem description

This example uses the solid mechanics solver to handle a cased wellbore problem subjected to a pressure test. The completed wellbore is composed of a steel casing, a cement sheath and rock formation. Isotropic linear elastic behavior is assumed for all the three materials. No separation is allowed for the casing-cement and cement-rock contact interfaces.

Analytical results of the radial and hoop stresses, σ_{rr} and $\sigma_{\theta\theta}$, in casing, cement sheath and rock are expressed as (Hervé and Zaoui, 1995) :

$$\sigma_{rr} = (2.0 * \lambda + 2.0 * G) * A - \frac{2.0 * G * B}{r^2} \quad \sigma_{\theta\theta} = (2.0 * \lambda + 2.0 * G) * A + \frac{2.0 * G * B}{r^2}$$

where λ and G are the Lamé moduli, r is the radial coordinate, A and B are piecewise constants that are obtained by solving the boundary and interface conditions, as detailed in the postprocessing script.

Input file

This benchmark example uses no external input files and everything required is contained within a single GEOSX input file. The xml input file for this test case is located at:

```
src/coreComponents/physicsSolvers/solidMechanics/benchmarks/CasedElasticWellbore.xml
```

All inputs for this case are contained inside a single XML file. In this example, we would focus our attention on the `Solvers`, `Mesh` and `Constitutive` tags.

Solid mechanics solver

As fluid flow is not considered, only the solid mechanics `SolidMechanicsLagrangianSSLE` solver is required for solving this linear elastic problem. In this solver, the three regions and three materials associated to casing, cement sheath and rock are respectively defined by `targetRegions` and `solidMaterialNames`.

```
<SolidMechanicsLagrangianSSLE
  name="lagsolve"
  timeIntegrationOption="QuasiStatic"
  discretization="FE1"
  logLevel="0"
  targetRegions="{ casing, cement, rock }"
  solidMaterialNames="{ casing, cement, rock }">
```

Cased wellbore mesh

The internal wellbore mesh generator `InternalWellbore` is employed to create the mesh of this wellbore problem. The radii of the casing cylinder, the cement sheath cylinder and the far-field boundary of the surrounding rock formation are defined by a vector `radius`. In the tangent direction, `theta` angle is specified from 0 to 360 degree

for a full geometry of the domain. Note that a half or a quarter of the domain can be defined by a `theta` angle from 0 to 180 or 90 degree, respectively. The trajectory of the well is defined by `trajectory`, which is vertical in this case. The `autoSpaceRadialElems` parameters allow optimally increasing the element size from local zone around the wellbore to the far-field zone. In this example, the auto spacing option is only applied for the rock formation. The `useCartesianOuterBoundary` transforms the far-field boundary to a squared shape to enforce a Cartesian aligned outer boundary, which eases the loading of the boundary conditions. The `cellBlockNames` and `elementTypes` define the regions and related element types associated to casing, cement sheath and rock.

```
<Mesh>
  <InternalWellbore
    name="mesh1"
    elementTypes="{ C3D8, C3D8, C3D8 }"
    radius="{ 0.1, 0.106, 0.133, 2.0 }"
    theta="{ 0, 360 }"
    zCoords="{ 0, 1 }"
    nr="{ 10, 20, 10 }"
    nt="{ 320 }"
    nz="{ 1 }"
    trajectory="{ { 0.0, 0.0, 0.0 },
                  { 0.0, 0.0, 1.0 } }"
    autoSpaceRadialElems="{ 0, 0, 1 }"
    useCartesianOuterBoundary="2"
    cellBlockNames="{ casing, cement, rock }"
  />
</Mesh>
```

Steel, cement, and rock constitutive laws

Isotropic linear elastic constitutive behavior is considered for all the three materials. Note that the default density is useless for this case.

```
<ElasticIsotropic
  name="casing"
  defaultDensity="2700"
  defaultBulkModulus="175e9"
  defaultShearModulus="80.8e9"/>

<ElasticIsotropic
  name="cement"
  defaultDensity="2700"
  defaultBulkModulus="10.3e9"
  defaultShearModulus="6.45e9"/>

<ElasticIsotropic
  name="rock"
  defaultDensity="2700"
  defaultBulkModulus="5.5556e9"
  defaultShearModulus="4.16667e9"/>
```

Boundary conditions

Far-field boundary are subjected to roller constraints. The normal traction on the inner face of the casing is defined by `Traction` field specification. The nodeset generated by the internal wellbore generator for this face is named as `rneg`. The traction type is normal to mimic a casing test pressure that is applied normal to the casing inner face . The negative sign of the scale is attributed to the negative sign convention for compressive stress in GEOSX.




Fig. 4.5: Mesh for simulating a cased wellbore

```

<FieldSpecifications>

  <FieldSpecification
    name="xConstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="0"
    scale="0.0"
    setNames="{ xneg, xpos }"/>

  <FieldSpecification
    name="yConstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="1"
    scale="0.0"
    setNames="{ yneg, ypos }"/>

  <FieldSpecification
    name="zconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="2"
    scale="0.0"
    setNames="{ zneg, zpos }"/>

  <Traction
    name="innerPressure"
    objectPath="faceManager"
    tractionType="normal"
    scale="-10.0e6"
    setNames="{ rneg }"/>
</FieldSpecifications>

```

Results and benchmark

A good agreement between the GEOSX results and analytical results is shown in the figure below

To go further

Feedback on this example

This concludes the cased wellbore example. For any feedback on this example, please submit a [GitHub issue](#) on the project's [GitHub](#) page.

4.1.4 Poromechanical Wellbore Problem

Context

The main objective of this example is to demonstrate how to use the internal wellbore mesh generator and poromechanical solvers in GEOSX to tackle wellbore problems in porous media. In this example, a poroplastic model is applied to find the solution of rock deformation within the vicinity of a vertical wellbore, considering elastoplastic deformation, fluid diffusion and poromechanical coupling effect. To do so, a single phase flow solver is fully coupled with a Lagrangian mechanics solver and the Extended Drucker-Prager model (see `TwoInvariantPlasticity`) is chosen as





Fig. 4.6: Comparing GEOSX results with analytical solution



the material model for the solid domain. We first solve this problem with a poroelastic model and verify the modeling results with the corresponding analytical solutions. Then, the verified case is modified to test a poroplastic version, whose results are compared with the ones obtained from the poroelastic case to highlight the impact of plasticity in this specific problem.

Objectives

At the end of this example you will know:

- how to construct meshes for wellbore problems with the internal wellbore mesh generator,
- how to specify initial and boundary conditions, such as reservoir properties, in-situ stresses, mixed loading (mechanical and fluid) at wellbore wall and far-field constraints,
- how to use multiple solvers in GEOSX for predicting poroplastic deformations in the near wellbore region.

Input file

This example uses no external input files and everything required is contained within a single GEOSX input file.

The xml input file for the test case with poroelasticity is located at:

```
examples/plasticity/WellboreProblem_PoroElastic.xml
```

The xml input file for the test case with poroplasticity is located at:

```
examples/plasticity/WellboreProblem_PoroDruckerPrager.xml
```

Description of the case

We simulate the wellbore problem subjected to anisotropic horizontal stress (σ_h and σ_H) and vertical stress (σ_v), as shown below. This is a vertical wellbore, which is drilled in a porous medium. By changing the wellbore supporting pressure, the mechanical deformation of the reservoir rock will be induced and evolve with time, due to fluid diffusion and coupling effect. Considering inelastic constitutive behavior, the reservoir rock in the near wellbore region will experience elastoplastic deformation and a plastic zone will be developed and expand with time. To setup the base case, a poroelastic version is employed to find the poroelastic solutions of this wellbore problem, which are verified with the analytical solution (Detournay and Cheng, 1993) from the literature. Following that, a poroplastic version is built and used to obtain the temporal and spatial solutions of pore pressure, displacement and stress fields around the wellbore, considering induced plastic deformation.

All inputs for this case are contained inside a single XML file. In this example, we focus our attention on the `Mesh` tags, the `Solver` tags, the `Constitutive` tags, and the `FieldSpecifications` tags.

Mesh

The following figure shows the generated mesh that is used for solving this wellbore problem

Let us take a closer look at the geometry of this wellbore problem. We use the internal mesh generator `InternalWellbore` to create a rock domain ($10\text{ m} \times 5\text{ m} \times 2\text{ m}$), with a wellbore of initial radius equal to 0.1 m. Coordinates of `trajectory` defines the wellbore trajectory, which represents a perfect vertical well in this example. By turning on `autoSpaceRadialElems="{ 1 }"`, the internal mesh generator automatically sets number and spacing of elements in the radial direction, which overrides the values of `nr`. With `useCartesianOuterBoundary="0"`, a Cartesian aligned outer boundary on the outer block is enforced. In this way, a structured three-dimensional mesh is created with $100 \times 80 \times 2$ elements in the radial, tangential and z directions, respectively. All the elements are eight-node hexahedral elements (C3D8) and refinement is performed to conform with the wellbore geometry. This mesh is defined as a cell block with the name `cb1`.




Fig. 4.7: Sketch of the wellbore problem




Fig. 4.8: Generated mesh for the wellbore problem

```

<Mesh>
  <InternalWellbore
    name="mesh1"
    elementTypes="{ C3D8 }"
    radius="{ 0.1, 5.0 }"
    theta="{ 0, 180 }"
    zCoords="{ -1, 1 }"
    nr="{ 40 }"
    nt="{ 80 }"
    nz="{ 2 }"
    trajectory="{ { 0.0, 0.0, -1.0 },
                  { 0.0, 0.0, 1.0 } }"
    autoSpaceRadialElems="{ 1 }"
    useCartesianOuterBoundary="0"
    cellBlockNames="{ cb1 }"/>
</Mesh>

```

Solid mechanics solver

GEOSX is a multi-physics platform. Different combinations of physics solvers available in the code can be applied in different regions of the domain and be functional at different stages of the simulation. The `Solvers` tag in the XML file is used to list and parameterize these solvers.

To specify a coupling between two different solvers, we define and characterize each single-physics solver separately. Then, we customize a *coupling solver* between these single-physics solvers as an additional solver. This approach allows for generality and flexibility in constructing multi-physics solvers. The order of specifying these solvers is not restricted in GEOSX. Note that end-users should give each single-physics solver a meaningful and distinct name, as GEOSX will recognize these single-physics solvers based on their customized names and create user-expected coupling.

As demonstrated in this example, to setup a poromechanical coupling, we need to define three different solvers in the XML file:

- the mechanics solver, a solver of type `SolidMechanics_LagrangianFEM` called here `mechanicsSolver` (more information here: [Solid Mechanics Solver](#)),

```

<SolidMechanics_LagrangianFEM
  name="mechanicsSolver"
  timeIntegrationOption="QuasiStatic"
  logLevel="1"
  discretization="FE1"
  targetRegions="{ Omega }"
  solidMaterialNames="{ rock }">
  <NonlinearSolverParameters
    newtonTol = "1.0e-5"
    newtonMaxIter = "15"
  />
</SolidMechanics_LagrangianFEM>

```

- the single-phase flow solver, a solver of type `SinglePhaseFVM` called here `SinglePhaseFlowSolver` (more information on these solvers at [Singlephase Flow Solver](#)),

```

<SinglePhaseFVM
  name="SinglePhaseFlowSolver"
  logLevel="1"
  discretization="singlePhaseTPFA"

```

(continues on next page)

(continued from previous page)

```

targetRegions="{Omega}"
fluidNames="{ water }"
solidNames="{ porousRock }"
permeabilityNames="{rockPerm}">
<NonlinearSolverParameters
  newtonTol = "1.0e-6"
  newtonMaxIter = "8"
/>
</SinglePhaseFVM>
</Solvers>

```

- the coupling solver (`SinglePhasePoromechanics`) that will bind the two single-physics solvers above, which is named as `PoromechanicsSolver` (more information at `SinglePhasePoromechanics`).

The two single-physics solvers are parameterized as explained in their corresponding documents.

In this example, let us focus on the coupling solver. This solver (`PoromechanicsSolver`) uses a set of attributes that specifically describe the coupling process within a poromechanical framework. For instance, we must point this solver to the designated fluid solver (here: `SinglePhaseFlowSolver`) and solid solver (here: `mechanicsSolver`). These solvers are forced to interact through the `porousMaterialNames="{porousRock}"` with all the constitutive models. We specify the discretization method (`FE1`, defined in the `NumericalMethods` section), and the target regions (here, we only have one, `Omega`). More parameters are required to characterize a coupling procedure (more information at [Poromechanics Solver](#)). In this way, the two single-physics solvers will be simultaneously called and executed for solving the wellbore problem here.

Discretization methods for multiphysics solvers

Numerical methods in multiphysics settings are similar to single physics numerical methods. In this problem, we use finite volume for flow and finite elements for solid mechanics. All necessary parameters for these methods are defined in the `NumericalMethods` section.

As mentioned before, the coupling solver and the solid mechanics solver require the specification of a discretization method called `FE1`. In GEOSX, this discretization method represents a finite element method using linear basis functions and Gaussian quadrature rules. For more information on defining finite elements numerical schemes, please see the dedicated `FiniteElement` section.

The finite volume method requires the specification of a discretization scheme. Here, we use a two-point flux approximation scheme (`singlePhaseTPFA`), as described in the dedicated documentation (found here: `FiniteVolume`).

```

<NumericalMethods>
  <FiniteElements>
    <FiniteElementSpace
      name="FE1"
      order="1"/>
    </FiniteElements>
  <FiniteVolume>
    <TwoPointFluxApproximation
      name="singlePhaseTPFA"
      fieldName="pressure"
      coefficientName="permeability"
      coefficientModelNames="{rockPerm}"
    />
  </FiniteVolume>
</NumericalMethods>

```

Constitutive laws

For this test problem, the solid and fluid materials are named as `rock` and `water` respectively, whose mechanical properties are specified in the `Constitutive` section. In this example, different material models, linear elastic isotropic model (see `LinearElasticIsotropic`) and Extended Drucker-Prager model (see `TwoInvariantPlasticity`), are used to solve the mechanical deformation, which is the only difference between the poroelastic and poroplastic cases in this example.

For the poroelastic case, `PorousElasticIsotropic` model is used to describe the linear elastic isotropic response of rock to loading. And the single-phase fluid model `CompressibleSinglePhaseFluid` is selected to simulate the flow of water upon injection:

```
<Constitutive>
  <PorousElasticIsotropic
    name="porousRock"
    solidModelName="rock"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"
  />
  <ElasticIsotropic
    name="rock"
    defaultDensity="2700"
    defaultBulkModulus="1.1111e10"
    defaultShearModulus="8.3333e9"
  />
  <CompressibleSinglePhaseFluid
    name="water"
    defaultDensity="1000"
    defaultViscosity="0.001"
    referencePressure="0e6"
    referenceDensity="1000"
    compressibility="2.09028227021e-10"
    referenceViscosity="0.001"
    viscosibility="0.0"
  />
  <BiotPorosity
    name="rockPorosity"
    grainBulkModulus="1.0e27"
    defaultReferencePorosity="0.3"
  />
  <ConstantPermeability
    name="rockPerm"
    permeabilityComponents="{1.0e-20, 1.0e-20, 1.0e-20}"
  />
</Constitutive>
```

For the poroplastic case, `PorousExtendedDruckerPrager` model is used to simulate the elastoplastic behavior of rock. And the single-phase fluid model `CompressibleSinglePhaseFluid` is employed to handle the storage and flow of water:

```
<Constitutive>
  <PorousExtendedDruckerPrager
    name="porousRock"
    solidModelName="rock"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPerm"
  />
```

(continues on next page)

(continued from previous page)

```

<ExtendedDruckerPrager
  name="rock"
  defaultDensity="2700"
  defaultBulkModulus="1.1111e10"
  defaultShearModulus="8.3333e9"
  defaultCohesion="1.0e6"
  defaultInitialFrictionAngle="15.27"
  defaultResidualFrictionAngle="23.05"
  defaultDilationRatio="1.0"
  defaultHardening="0.01"
/>
<CompressibleSinglePhaseFluid
  name="water"
  defaultDensity="1000"
  defaultViscosity="0.001"
  referencePressure="0e6"
  referenceDensity="1000"
  compressibility="2.09028227021e-10"
  referenceViscosity="0.001"
  viscosibility="0.0"
/>
<BiotPorosity
  name="rockPorosity"
  grainBulkModulus="1.0e27"
  defaultReferencePorosity="0.3"
/>
<ConstantPermeability
  name="rockPerm"
  permeabilityComponents="{1.0e-20, 1.0e-20, 1.0e-20}"
/>
</Constitutive>

```

As for the material parameters, `defaultInitialFrictionAngle`, `defaultResidualFrictionAngle` and `defaultCohesion` denote the initial friction angle, the residual friction angle, and cohesion, respectively, as defined by the Mohr-Coulomb failure envelope. As the residual friction angle `defaultResidualFrictionAngle` is larger than the initial one `defaultInitialFrictionAngle`, a strain hardening model is automatically chosen, whose hardening rate is given as `defaultHardening="0.01"`. If the residual friction angle is set to be less than the initial one, strain weakening will take place. `defaultDilationRatio="1.0"` corresponds to an associated flow rule. If using an incompressible fluid, the user can lower the fluid compressibility `compressibility` to 0. The constitutive parameters such as the density, the bulk modulus, and the shear modulus are specified in the International System of Units. A stress-dependent porosity model `rockPorosity` and constant permeability `rockPerm` model are defined in this section.

Initial and boundary conditions

The next step is to specify fields, including:

- The initial value (the in-situ stresses and pore pressure have to be initialized)
- The boundary conditions (traction and fluid loading at the wellbore wall and constraints of the outer boundaries have to be set)

In this example, we need to specify anisotropic horizontal stress ($\sigma_h = -9.0$ MPa and $\sigma_H = -11.0$ MPa) and vertical stress ($\sigma_v = -12.0$ MPa). A compressive traction (InnerMechanicalLoad) $P_w = -10$ MPa and fluid loading (InnerFluidLoad) $P_f = 10$ MPa are applied at the wellbore wall `rneg`. The remaining parts of the outer boundaries are subjected to roller constraints. These boundary conditions are set up through the `FieldSpecifications`

section.

```

<FieldSpecifications>
  <FieldSpecification
    name="stressXX"
    initialCondition="1"
    setNames="{all}"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="0"
    scale="-9.0e6"
  />

  <FieldSpecification
    name="stressYY"
    initialCondition="1"
    setNames="{all}"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="1"
    scale="-11.0e6"
  />

  <FieldSpecification
    name="stressZZ"
    initialCondition="1"
    setNames="{all}"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="rock_stress"
    component="2"
    scale="-12.0e6"
  />

  <FieldSpecification
    name="initialPressure"
    initialCondition="1"
    setNames="{all}"
    objectPath="ElementRegions/Omega/cb1"
    fieldName="pressure"
    scale="0e6"
  />

  <FieldSpecification
    name="xconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="0"
    scale="0.0"
    setNames="{xneg, xpos}"
  />

  <FieldSpecification
    name="yconstraint"
    objectPath="nodeManager"
    fieldName="TotalDisplacement"
    component="1"
    scale="0.0"
    setNames="{tneg, tpos, ypos}"
  />

```

(continues on next page)

(continued from previous page)

```

/>

<FieldSpecification
  name="zconstraint"
  objectPath="nodeManager"
  fieldName="TotalDisplacement"
  component="2"
  scale="0.0"
  setNames="{zneg, zpos}"
/>

<Traction
  name="InnerMechanicalLoad"
  setNames="{ rneg }"
  objectPath="faceManager"
  scale="-10.0e6"
  tractionType="normal"
  functionName="timeFunction"
/>

<FieldSpecification
  name="InnerFluidLoad"
  setNames="{ rneg }"
  objectPath="faceManager"
  fieldName="pressure"
  scale="10e6"
  functionName="timeFunction"
/>
</FieldSpecifications>

```

With `tractionType="normal"`, traction is applied to the wellbore wall `rneg` as a pressure specified from the product of scale `scale="-10.0e6"` and the outward face normal. A table function `timeFunction` is used to define the time-dependent loading. The coordinates and values form a time-magnitude pair for the loading time history. In this case, the loading magnitude is given as:

```

<Functions>
  <TableFunction
    name="timeFunction"
    inputVarNames="{time}"
    coordinates="{0.0, 0.1, 1e6}"
    values="{0.0, 1.0, 1.0}"
  />
</Functions>

```

You may note :

- All initial value fields must have `initialCondition` field set to 1;
- The `setName` field points to the previously defined box to apply the fields;
- `nodeManager` and `faceManager` in the `objectPath` indicate that the boundary conditions are applied to the element nodes and faces, respectively;
- `fieldName` is the name of the field registered in GEOSX;
- Component 0, 1, and 2 refer to the x, y, and z direction, respectively;
- And the non-zero values given by `scale` indicate the magnitude of the loading;

- Some shorthands, such as `xneg` and `xpos`, are used as the locations where the boundary conditions are applied in the computational domain. For instance, `xneg` means the portion of the computational domain located at the left-most in the x-axis, while `xpos` refers to the portion located at the right-most area in the x-axis. Similar shorthands include `ypos`, `yneg`, `zpos`, and `zneg`;
- The mud pressure loading has a negative value due to the negative sign convention for compressive stress in GEOSX.

The parameters used in the simulation are summarized in the following table, which are specified in the `Constitutive` and `FieldSpecifications` sections.

Symbol	Parameter	Unit	Value
K	Bulk Modulus	[GPa]	11.11
G	Shear Modulus	[GPa]	8.33
C	Cohesion	[MPa]	1.0
ϕ_i	Initial Friction Angle	[degree]	15.27
ϕ_r	Residual Friction Angle	[degree]	23.05
c_h	Hardening Rate	[-]	0.01
σ_h	Min Horizontal Stress	[MPa]	-9.0
σ_H	Max Horizontal Stress	[MPa]	-11.0
σ_v	Vertical Stress	[MPa]	-12.0
a_0	Initial Well Radius	[m]	0.1
P_w	Traction at Well	[MPa]	-10.0
P_f	Fluid Pressure at Well	[MPa]	10.0
ρ_f	Fluid Density	[kg/m ³]	1000.0
μ	Fluid Viscosity	[Pa s]	0.001
c_f	Fluid Compressibility	[Pa ⁻¹]	2.09×10^{-10}
κ	Matrix Permeability	[m ²]	1.0×10^{-20}
ϕ	Porosity	[-]	0.3

Inspecting results

As defined in the `Events` section, we run this simulation for 497640 seconds. In the above examples, we requested `silo-format` output files. We can therefore import these into `VisIt` and use python scripts to visualize the outcome. Please note that a non-dimensional time is used in the analytical solution, and the end time here leads to a non-dimensional end time of $t^* = 4.62$.

Using the poroelastic solver, below figure shows the prediction of pore pressure distribution upon fluid injection.

For the above poroelastic example, an analytical solution (Detournay and Cheng, 1993) is hereby employed to verify the accuracy of the numerical results. Following figure shows the comparisons between the numerical predictions (marks) and the corresponding analytical solutions (solid curves) with respect to the distributions of pore pressure, radial displacement, effective radial and tangential stresses along the minimum horizontal stress direction (x-axis). One can observe that GEOSX results correlate very well with the analytical solutions for the poroelastic case.

For the same 3D wellbore problem, the poroplastic case is thereafter tested and compared with the poroelastic one. The figure below shows the distribution of σ_{yy} in the near wellbore region for both cases. As expected, a relaxation of the tangential stress along the direction of minimum horizontal stress is detected, which can be attributed to the plastic response of the rock.

By using python scripts, we can extract the simulation results along any direction and provide detailed comparisons between different cases. Here, the pore pressure, radial displacement, radial and tangential effective stresses along the direction of minimum horizontal stress are obtained at different time steps and plotted against the corresponding ones of the poroelastic case. Because of fluid diffusion and coupling effect, following figure shows that these solutions evolve with time for both cases. As mentioned above, a plastic zone is developed in the vicinity of the wellbore, due

DB: plot_00103200
Cycle: 1032 Time:497640




Fig. 4.9: Simulation result of pore pressure distribution




Fig. 4.10: Comparing GEOSX results with analytical solutions

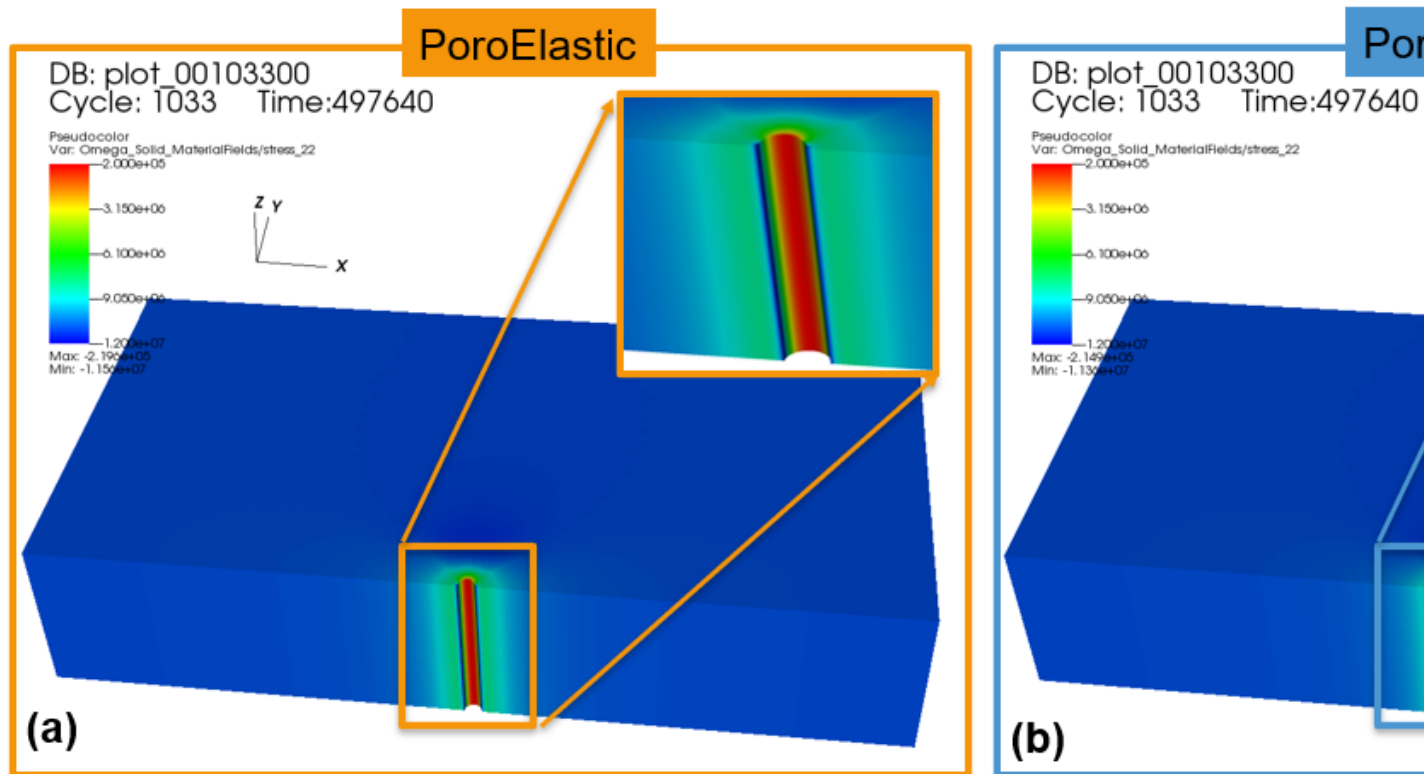


Fig. 4.11: Simulation result of S_{yy} : PoroElastic vs. PoroPlastic

to stress concentration. As for the far field region, these two cases become almost identical, with the rock deformation governed by poroelasticity.

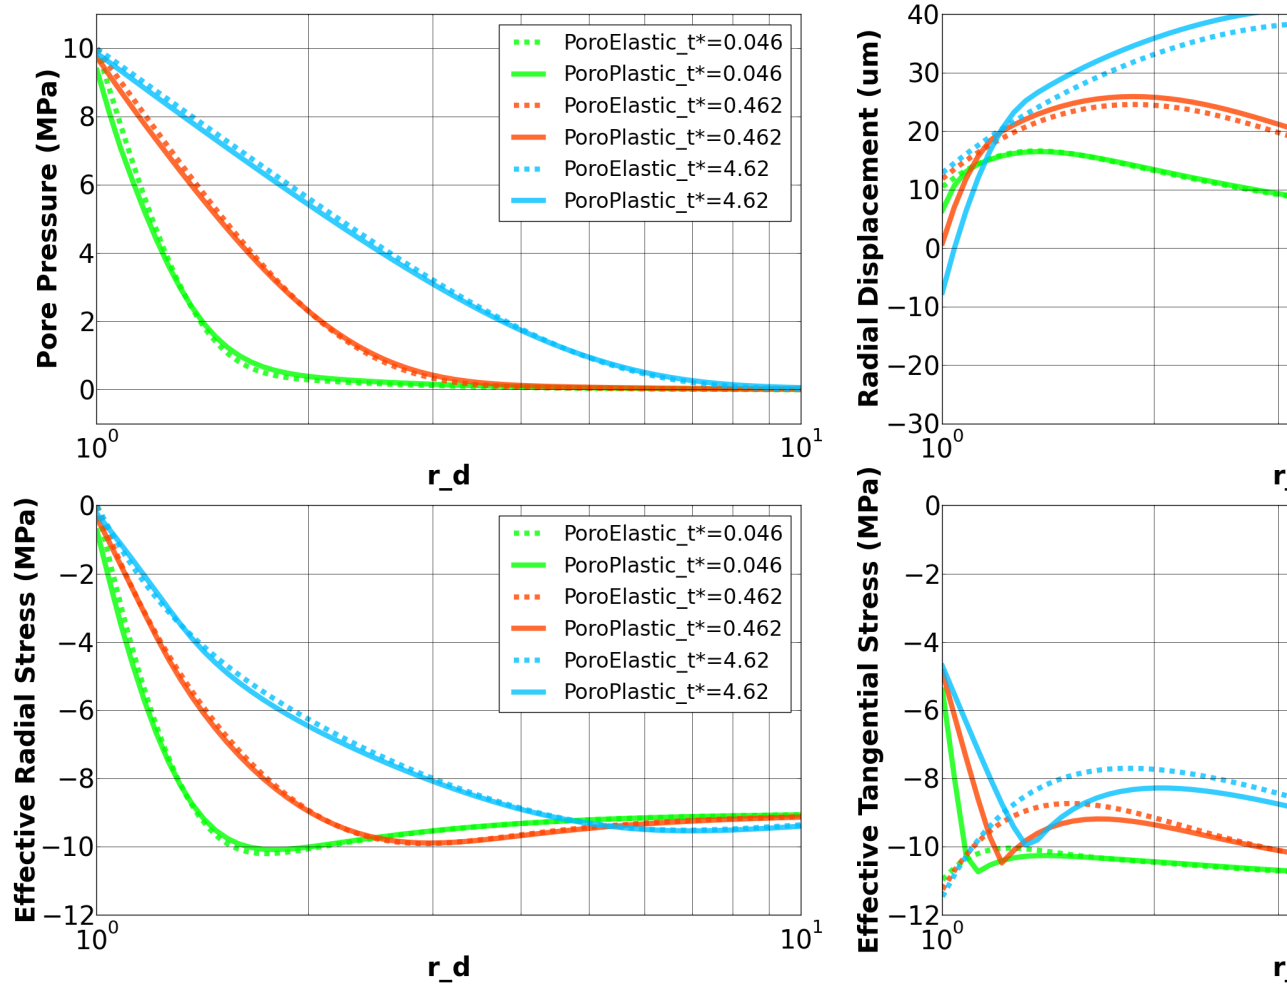


Fig. 4.12: Comparing the PoroPlastic case with the PoroElastic case at different times

To go further

Feedback on this example

This concludes the example on PoroPlasticity Model for Wellbore Problems. For any feedback on this example, please submit a [GitHub issue](#) on the project's GitHub page.

For more details

- More on plasticity models, please see [Model: Extended Drucker-Prager](#).
- More on multiphysics solvers, please see [SinglePhasePoromechanics](#).

4.2 Performance Benchmarks

4.3 Application Studies

Welcome to the GEOSX user guide.

5.1 Input Files

5.1.1 XML

GEOSX is configured via one (or more) [Extensible Markup Language \(XML\)](#) files. These files contain a set of elements and attributes that closely follow the internal datastructure of GEOSX. When running GEOSX, these files are specified using the *-i* argument:

```
geosx -i input.xml
```

XML Components

The following illustrates some of the key features of a GEOSX-format xml file:

```
<?xml version="1.0" ?>

<Problem>
  <BlockA
    someAttribute="1.234">

    <!-- Some comment -->
    <BlockB
      name="firstNamedBlock"
      anotherAttribute="0"/>
    <BlockB
      name="secondNamedBlock"
      anotherAttribute="1"/>
```

(continues on next page)

(continued from previous page)

```
</BlockA>
</Problem>
```

The two basic components of an xml file are blocks, which are specified using angle brackets (“<BlockA> </BlockA>”), and attributes that are attached to blocks (attributeName=“attributeValue”). Block and attributes can use any ASCII character aside from <, &, ‘, and “ (if necessary, use <, &, ', or "). Comments are indicated as follows: <!-- Some comment -->.

At the beginning of a GEOSX input file, you will find an optional xml declaration (<?xml version=“1.0” ?>) that is used to indicate the format to certain text editors. You will also find the root *Problem* block, where the GEOSX configuration is placed. Note that, aside from these elements and commented text, the xml format requires that no other objects exist at the first level.

In the example above, there is a single element within the *Problem* block: *BlockA*. *BlockA* has an attribute *someAttribute*, which has a value of 1.234, and has three children: a commented string “Some comment” and two instances of *BlockB*. The *name* attribute is required for blocks that allow multiple instances, and should include a unique string to avoid potential errors. Where applicable these blocks will be executed in the order in which they are specified in input file.

5.1.2 Input Validation

The optional *xmlns:xsi* and *xsi:noNamespaceSchemaLocation* attributes in the *Problem* block can be used to indicate the type of document and the location of the xml schema to the text editor:

```
<Problem
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="/path/to/schema.xsd" />
```

The schema contains a list of xml blocks and attributes that are supported by GEOSX, indicates whether a given object is optional or required, and defines the format of the object (string, floating point number, etc.). A copy of the schema is included in the GEOSX source code (/path/to/GEOSX/src/coreComponents/schema/schema.xsd). It can also be generated using GEOSX: `geosx -s schema.xsd`

Many text editors can use the schema to help in the construction of an xml file and to indicate whether it is valid. Using a validation tool is highly recommended for all users. The following instructions indicate how to turn on validation for a variety of tools:

xmllint

xmllint is a command-line tool that is typically pre-installed on UNIX-like systems. To check whether an input file is valid, run the following command:

```
xmllint --schema /path/to/schema.xsd input_file.xml
```

Sublime Text

We recommend using the [Exalt](#) or [SublimeLinter_xmllint](#) plug-ins to validate xml files within sublime. If you have not done so already, install the sublime [Package Control](#). To install the package, press ctrl + shift + p, type and select [Package Control: Install Package](#), and search for [exalt](#) or [SublimeLinter / SublimeLinter-xmllint](#). Note that, depending on the circumstances, these tools may indicate only a subset of the validation errors at a given time. Once resolved, the tools should re-check the document to look for any additional errors.

As an additional step for SublimLinter-xmlint, you will need to add a linter configuration. To do so, go to Preferences/Package Settings/SublimeLinter/Settings. In the right-hand side of the new window, add the xmlint configuration:

```
{
  "linters": {
    "xmlint": {
      {
        "args": "--schema /path/to/schema.xsd",
        "styles": [
          {
            "mark_style": "fill",
            "scope": "region.bluish",
            "types": ["error"],
            "icon": "stop",
          }
        ]
      },
    ],
  },
}
```

Eclipse

The Eclipse Web Develop Tools includes features for validating xml files. To install them, go to Help -> Eclipse Marketplace, search for the Eclipse Web Developer Tools, install the package, and restart Eclipse. Finally, configure the xml validation preferences under Window -> Preferences -> XML -> XML Files -> Validation. Eclipse will automatically fetch the schema, and validate an active xml file. The editor will highlight any lines with errors, and underline the specific errors.

GEOSX XML Tools

The `geosx_xml_tools` package, which is used to enable advanced features such as parameters, symbolic math, etc., contains tools for validating xml files. To do so, call the command-line script with the `-s` argument, i.e.: `preprocess_xml input_file.xml -s /path/to/schema.xsd`. After compiling the final xml file, `pygeosx` will fetch the designated schema, validate, and print any errors to the screen.

Note: Attributes that are using advanced xml features will likely contain characters that are not allowed by their corresponding type pattern. As such, file editors that are configured to use other validation methods will likely identify errors in the raw input file.

5.1.3 XML Schema

An XML schema definition (XSD) file lays out the expected structure of an input XML file. During the build process, GEOSX automatically constructs a comprehensive schema from the code's data structure, and updates the version in the source (`GEOSX/src/coreComponents/schema/schema.xsd`).

Schema Components

The first entry in the schema are a set of headers the file type and version. Following this, the set of available simple types for attributes are laid out. Each of these includes a variable type name, which mirrors those used in the main code, and a regular expression, which is designed to match valid inputs. These patterns are defined and documented in `DataTypes::typeRegex`. The final part of the schema is the file layout, beginning with the root `Problem`. Each

complex type defines an element, its children, and its attributes. Each attribute defines the input name, type, default value, and/or usage. Comments preceding each attribute are used to relay additional information to the users.

Automatic Schema Generation

A schema may be generated by calling the main code with the `-s` argument , e.g.: `geosx -s schema.xsd` (Note: this is done automatically during the build process). To do this, GEOSX does the following:

- 1) Initialize the GEOSX data structure.
- 2) Initialize objects that are registered to catalogs via `ManagedGroup::ExpandObjectCatalogs()`.
- 3) Recursively write element and attribute definitions to the schema using information stored in GEOSX groups and wrappers.
- 4) Define any expected deviations from the schema via `ManagedGroup::SetSchemaDeviations()`.

5.2 Meshes

The purpose of this document is to explain how users and developers interact with mesh data. This section describes how meshes are handled and stored in GEOSX.

There are two possible methods for generating a mesh: either by using GEOSX's internal mesh generator (for Cartesian meshes only), or by importing meshes from various common mesh file formats. This latter options allows one to work with more complex geometries, such as unstructured meshes comprised of a variety of element types (polyhedral elements).

5.2.1 Internal Mesh Generation

Basic Example

The Internal Mesh Generator allows one to quickly build simple cartesian grids and divide them into several regions. The following attributes are supported in the input block for `InternalMesh`:


Name	Type	Default	Description
cellBlock-Names	string_array	required	Names of each mesh block
element-Types	string_array	required	Element types of each mesh block
name	string	required	A name is required for any non-unique nodes
nx	integer_array	required	Number of elements in the x-direction within each mesh block
ny	integer_array	required	Number of elements in the y-direction within each mesh block
nz	integer_array	required	Number of elements in the z-direction within each mesh block
positionTolerance	real64	1e-10	A position tolerance to verify if a node belong to a nodeset
trianglePattern	integer	0	Pattern by which to decompose the hex mesh into prisms (more explanation required)
xBias	real64_array{1}		Bias of element sizes in the x-direction within each mesh block ($dx_left=(1+b)*L/N$, $dx_right=(1-b)*L/N$)
xCords	real64_array	required	x-coordinates of each mesh block vertex
yBias	real64_array{1}		Bias of element sizes in the y-direction within each mesh block ($dy_left=(1+b)*L/N$, $dy_right=(1-b)*L/N$)
yCords	real64_array	required	y-coordinates of each mesh block vertex
zBias	real64_array{1}		Bias of element sizes in the z-direction within each mesh block ($dz_left=(1+b)*L/N$, $dz_right=(1-b)*L/N$)
zCords	real64_array	required	z-coordinates of each mesh block vertex

The following is an example XML <mesh> block, which will generate a vertical beam with two CellBlocks (one in red and one in blue in the following picture).

```
<Mesh>
  <InternalMesh name="mesh"
    elementTypes="C3D8"
    xCords="0, 1"
    yCords="0, 1"
    zCords="0, 2, 6"
    nx="1"
    ny="1"
    nz="2, 4"
    cellBlockNames="cb1 cb2"/>
</Mesh>
```

- name the name of the mesh body
- elementTypes the type of the elements that will be generated.
- xCoord List of x coordinates of the boundaries of the CellBlocks
- yCoord List of y coordinates of the boundaries of the CellBlocks
- zCoord List of z coordinates of the boundaries of the CellBlocks

- `nx` List containing the number of cells in `x` direction within the `CellBlocks`
- `ny` List containing the number of cells in `y` direction within the `CellBlocks`
- `nz` List containing the number of cells in `z` direction within the `CellBlocks`
- `cellBlockNames` List containing the names of the `CellBlocks`



Mesh Bias

The internal mesh generator is capable of producing meshes with element sizes that vary smoothly over space. This is achieved by specifying `xBias`, `yBias`, and/or `zBias` fields. (Note: if present, the length of these must match `nx`, `ny`, and `nz`, respectively, and each individual value must be in the range $(-1, 1)$.)

For a given element block, the average element size will be

$$dx_{average}[i] = \frac{xCoords[i+1] - xCoords[i]}{nx[i]},$$

the element on the left-most side of the block will have size

$$dx_{left}[i] = (1 + xBias[i]) \cdot dx_{average}[i],$$

and the element on the right-most side will have size

$$dx_{right}[i] = (1 - xBias[i]) \cdot dx_{average}[i].$$

The following are the two most common scenarios that occur while designing a mesh with bias:

1. The size of the block and the element size on an adjacent region are known. Assuming that we are to the left of the target block, the appropriate bias would be:


$$xBias[i] = 1 - \frac{nx[i] \cdot dx_{left}[i + 1]}{xCoords[i + 1] - xCoords[i]}$$

2. The bias of the block and the element size on an adjacent region are known. Again, assuming that we are to the left of the target block, the appropriate size for the block would be:

$$xCoords[i + 1] - xCoords[i] = \frac{nx[i] \cdot dx_{left}[i + 1]}{1 - xBias[i]}$$

The following is an example of a mesh block along each dimension, and an image showing the corresponding mesh. Note that there is a core region of elements with zero bias, and that the transitions between element blocks are smooth.

```
<Mesh>
  <InternalMesh
    name="mesh1 "
    elementTypes="{ C3D8 }"
    xCoords="{ -10, -1, 0, 1, 10 }"
    yCoords="{ -10, -1, 0, 1, 10 }"
    zCoords="{ -10, -1, 0, 1, 10 }"
    nx="{ 4, 1, 1, 4 }"
    ny="{ 5, 1, 1, 5 }"
    nz="{ 6, 1, 1, 6 }"
    xBias="{ 0.555, 0, 0, -0.555 }"
    yBias="{ 0.444, 0, 0, -0.444 }"
    zBias="{ 0.333, 0, 0, -0.333 }"
    cellBlockNames="{ cb1 }"/>
  </InternalMesh>
</Mesh>
```



Advanced Cell Block Specification

It's possible to generate more complex CellBlock using the InternalMeshGenerator. For instance, the staircase example is a model which is often used in GEOSX as an integrated test. It defines CellBlocks in the three directions to generate a staircase-like model with the following code.

```
<Mesh>
  <InternalMesh name="mesh1"
    elementTypes="{C3D8}"
    xCoords="{0, 5, 10}"
    yCoords="{0, 5, 10}"
    zCoords="{0, 2.5, 5, 7.5, 10}"
    nx="{5, 5}"
    ny="{5, 5}"
    nz="{3, 3, 3, 3}"
    cellBlockNames="{b00,b01,b02,b03,b04,b05,b06,b07,b08,b09,b10,b11,b12,
↪ b13,b14,b15}" />
```

(continues on next page)

(continued from previous page)

```

</Mesh>

<ElementRegions>
  <CellElementRegion name="Channel"
    cellBlocks="{b08,b00,b01,b05,b06,b14,b15,b11}"
    materialList="{fluid1, rock, relperm}"/>
  <CellElementRegion name="Barrier"
    cellBlocks="{b04,b12,b13,b09,b10,b02,b03,b07}"
    materialList="{}/>
</ElementRegions>

```

Thus, the generated mesh will be :

5.2.2 Using an External Mesh

Supported Formats

GEOSX provides features to run simulations on unstructured meshes. It uses [PAMELA](#) to read the external meshes and its API to write it into the GEOSX mesh data structure.

The supported mesh format are:

- The [GMSH](#) file format (.msh v2).
- The ECLIPSE file formats (.egrid, .grdecl)

The supported mesh elements for volume elements consist of the following:

- 4 nodes tetrahedra,
- 5 nodes pyramids,
- 6 nodes wedges,
- 8 nodes hexahedra,

The mesh can be divided in several regions. These regions are intended to support different physics or to define different constitutive properties.

- For the GMSH file format, the regions are defined using the [physical entity names](#) provided by GMSH.
- For the ECLIPSE file formats, the regions have to be first defined using the ECLIPSE software.

Importing the Mesh

Importing regions

Several blocks are involved to import an external mesh into GEOSX, defined in the XML input file. These are the `<Mesh>` block and the `<ElementRegions>` block.

The mesh block has the following syntax:

```

<Mesh>
  <PAMELAMeshGenerator name="MyMeshName"
    file="/path/to/the/mesh/file.msh"/>
</Mesh>

```

We advise users to use absolute path to the mesh file.

GEOSX uses `ElementRegions` to support different physics or to define different constitutive properties. An `ElementRegion` is defined as a set of `CellBlocks`. A `CellBlock` is an ensemble of elements with the same element geometry.

In the example presented above, the mesh is composed of two regions (*Top* and *Bot*). Each region contains 3 `CellBlocks`.

The `ElementRegions` are defined as below :

```
<ElementRegions>
  <ElementRegion name="Top" cellBlocks="Top_HEX Top_WEDGE Top_TETRA" materialList=
  ↪ "water rock"/>
  <ElementRegion name="Bot" cellBlocks="Bot_HEX Bot_WEDGE Bot_TETRA" materialList=
  ↪ "water rock"/>
</ElementRegions>
```

You have to use the following syntax to declare your `CellBlocks` :

```
nameOfTheRegionWithinTheMesh_typeOfTheElement
```

The keywords for the element types are :

- TETRA
- WEDGE
- PYR
- HEX

If the regions are not named in the file (it happens with all the eclipse grids and several GMSH mesh files), the name of the region is `DEFAULT`, e.g:

```
<ElementRegions>
  <ElementRegion name="Default" cellBlocks="DEFAULT_HEX" materialList="water rock"/>
</ElementRegions>
```

Using the `gmsh` file format, regions can be easily named as a preprocessed step using the `gmsh` software or directly editing the file following the syntax defined in the [documentation](#).

An example of a `gmsh` file with all the physical regions defined is used in [Tutorial 3: Regions and Property Specifications](#).

Importing surfaces

Surfaces are imported through point sets in GEOSX. This feature is supported using only the `gmsh` file format. In the same way than the regions, the surfaces of interests can be defined using the [physical entity names](#). The surfaces are automatically imported in GEOSX if they exist in the `gmsh` file. Within GEOSX, the point set will have the same name than the one given in the file. This name can be used again to impose boundary condition. For instance, if a surface is named “Bottom” and the user wants to impose a Dirichlet boundary condition of 0 on it, it can be easily done using this syntax.

```
<FieldSpecification
  name="zconstraint"
  objectPath="nodeManager"
```

(continues on next page)

(continued from previous page)

```
fieldName="Velocity"
component="2"
scale="0.0"
setNames="{ Bottom }"/>
```

The name of the surface of interest appears under the keyword `setNames`. Again, an example of a gmsh file with the surfaces fully defined is available within *Tutorial 3: Regions and Property Specifications*.

5.3 Physics Solvers

The `<Solvers>` section of the input file specifies one or several physics solvers to be included in the simulation.

5.3.1 Solution Strategy

All physics solvers share a common solution strategy for nonlinear time-dependent problems. Here, we briefly describe the nonlinear solver and the timestepping strategy employed.

Nonlinear Solver

At each time-step, the nonlinear system of discrete residual equations, i.e.

$$r(x) = 0$$

is solved by employing the Newton-Raphson method. Here, x is the vector of primary unknowns. Thus, each physics solver is responsible for assembling the Jacobian matrix J containing the analytical derivatives of the residual vector r with respect to the primary variables. Then, at each Newton iteration ν , the following linear system is solved

$$J^\nu \delta x^{\nu+1} = -r^\nu,$$

where, $\delta x^{\nu+1}$ is the Newton update. This linear system can be solved with a variety of different linear solvers described in *Linear Solvers*. The Newton update, $\delta x^{\nu+1}$ is then applied to the primary variables:

$$x^{\nu+1} = x^\nu + \delta x^{\nu+1}.$$

This procedure is repeated until convergence is achieved or until the maximum number of iterations is reached.

Line Search

A line search method can be applied along with the Newton's method to facilitate Nonlinear convergence. After the Newton update, if the residual norm has increased instead of decreased, a line search algorithm is employed to correct the Newton update.

The user can choose between two different behaviors in case the line search fails to provide a reduced residual norm:

1. accept the solution and move to the next Newton iteration;
2. reject the solution and request a timestep cut;

Timestepping Strategy

The actual timestep size employed is determined by a combination of several factors. In particular, specific output events may have timestep requirements that force a specific timestep to be used. However, physics solvers do have the possibility of requesting a specific timestep size to the event manager based on their specific requirements. In particular, in case of fast convergence indicated by a small number of Newton iterations, i.e.

$$\text{numIterations} < \text{dtIncIterLimit} \cdot \text{newtonMaxIter},$$

the physics solver will require to double the timestep size. On the other hand, if a large number of nonlinear iterations are necessary to find the solution at timestep n

$$\text{numIterations} > \text{dtCutIterLimit} \cdot \text{newtonMaxIter},$$

the physics solver will request the next timestep, $n + 1$, to be half the size of timestep n . Here,

Additionally, in case the nonlinear solver fails to converge with the timestep provided by the event manager, the timestep size is cut, i.e.

$$\text{dt} = \text{timestepCutFactor} \cdot \text{dt},$$

and the nonlinear loop is repeated with the new timestep size.

Parameters

All parameters defining the behavior of the nonlinear solver and determining the timestep size requested by the physics solver are defined in the `NonlinearSolverParameters` and are presented in the following table.

Name	Type	Default	Description
allowNonConverged	integer	0	Allow non-converged solution to be accepted. (i.e. exit from the Newton loop without achieving the desired tolerance)
dtCutIterLimit	real64	0.7	Fraction of the Max Newton iterations above which the solver asks for the time-step to be cut for the next dt.
dtIncIterLimit	real64	0.4	Fraction of the Max Newton iterations below which the solver asks for the time-step to be doubled for the next dt.
lineSearchAction	geosx_NonlinearSolverParameters.lineSearchAction	Attempt	How the line search is to be used. Options are: <ul style="list-style-type: none"> * None - Do not use line search. * Attempt - Use line search. Allow exit from line search without achieving smaller residual than starting residual. * Require - Use line search. If smaller residual than starting residual is not achieved, cut time step.
lineSearchCutFactor	real64	0.5	Line search cut factor. For instance, a value of 0.5 will result in the effective application of the last solution by a factor of (0.5, 0.25, 0.125, ...)
lineSearchMaxCuts	integer	4	Maximum number of line search cuts.
logLevel	integer	0	Log level
maxSubSteps	integer	10	Maximum number of time sub-steps allowed for the solver
maxTimeStepCuts	integer	2	Max number of time step cuts
newtonMaxIter	integer	5	Maximum number of iterations that are allowed in a Newton loop.
newtonMinIter	integer	1	Minimum number of iterations that are required before exiting the Newton loop.
newtonTol	real64	1e-06	The required tolerance in order to exit the Newton iteration loop.
5.3. Physics Solvers			141
timestepCutFactor	real64	0.5	Factor by which the time step will be cut if a timestep cut is required.

5.3.2 Solid Mechanics Solver

List of Symbols

i, j, k	\equiv indices over spatial dimensions
a, b, c	\equiv indices over nodes
l	\equiv indices over volumetric elements
q, r, s	\equiv indices over faces
n	\equiv indices over time
$kiter$	\equiv iteration count for non-linear solution scheme
Ω	\equiv Volume of continuum body
Ω_{crack}	\equiv Volume of open crack
Γ	\equiv External surface of Ω
Γ_t	\equiv External surface where tractions are applied
Γ_u	\equiv External surface where kinematics are specified
Γ_{crack}	\equiv entire surface of crack
$\Gamma_{cohesive}$	\equiv surface of crack subject to cohesive tractions
η_0	\equiv set of all nodes
η_f	\equiv set of all nodes on flow mesh
m	\equiv mass
κ_k	\equiv all elements connected to element k
ϕ	\equiv porosity
p_f	\equiv fluid pressure
\mathbf{u}	\equiv displacement
\mathbf{q}	\equiv volumetric flow rate
\mathbf{T}	\equiv Cauchy stress
ρ	\equiv density in the current configuration
\mathbf{x}	\equiv current position
\mathbf{w}	\equiv aperture, or gap vector

Introduction

The *SolidMechanics_LagrangianFEM* solver applies a Continuous Galerkin finite element method to solve the linear momentum balance equation. The primary variable is the displacement field which is discretized at the nodes.

Theory

Governing Equations

The *SolidMechanics_LagrangianFEM* solves the equations of motion as given by

$$T_{ij,j} + \rho(b_i - \ddot{x}_i) = 0,$$

which is a 3-dimensional expression for the well known expression of Newtons Second Law ($F = ma$). These equations of motion are discretized using the Finite Element Method, which leads to a discrete set of residual equations:

$$(R_{solid})_{ai} = \int_{\Gamma_t} \Phi_a t_i dA - \int_{\Omega} \Phi_{a,j} T_{ij} dV + \int_{\Omega} \Phi_a \rho (b_i - \Phi_b \ddot{x}_{ib}) dV = 0$$

Quasi-Static Time Integration

The Quasi-Static time integration option solves the equation of motion after removing the inertial term, which is expressed by

$$T_{ij,j} + \rho b_i = 0,$$

which is essentially a way to express the equation for static equilibrium ($\Sigma F = 0$). Thus, selection of the Quasi-Static option will yield a solution where the sum of all forces at a given node is equal to zero. The resulting finite element discretized set of residual equations are expressed as

$$(R_{solid})_{ai} = \int_{\Gamma_t} \Phi_a t_i dA - \int_{\Omega} \Phi_{a,j} T_{ij} dV + \int_{\Omega} \Phi_a \rho b_i dV = 0,$$

Taking the derivative of these residual equations wrt. the primary variable (displacement) yields

$$\frac{\partial (R_{solid}^e)_{ai}}{\partial u_{bj}} = - \int_{\Omega^e} \Phi_{a,k} \frac{\partial T_{ik}}{\partial u_{bj}} dV,$$

And finally, the expression for the residual equation and derivative are used to express a non-linear system of equations

$$\left(\frac{\partial (R_{solid}^e)_{ai}}{\partial u_{bj}} \right) \bigg|_{kiter}^{n+1} \left((u_{bj})_{kiter+1}^{n+1} - (u_{bj})_{kiter}^{n+1} \right) = -(R_{solid})_{ai} \big|_{kiter}^{n+1},$$

which are solved via the solver package.

Implicit Dynamics Time Integration (Newmark Method)

For implicit dynamic time integration, we use an implementation of the classical Newmark method. This update method can be posed in terms of a simple SDOF spring/dashpot/mass model. In the following, M represents the mass, C represent the damping of the dashpot, K represents the spring stiffness, and F represents some external load.

$$Ma^{n+1} + Cv^{n+1} + Ku^{n+1} = F_{n+1},$$

and a series of update equations for the velocity and displacement at a point:

$$\begin{aligned} u^{n+1} &= u^n + v^{n+1/2} \Delta t, \\ u^{n+1} &= u^n + \left(v^n + \frac{1}{2} [(1 - 2\beta)a^n + 2\beta a^{n+1}] \Delta t \right) \Delta t, \\ v^{n+1} &= v^n + [(1 - \gamma)a^n + \gamma a^{n+1}] \Delta t. \end{aligned}$$

As intermediate quantities we can form an estimate (predictor) for the end of step displacement and midstep velocity by assuming zero end-of-step acceleration.

$$\begin{aligned} \tilde{u}^{n+1} &= u^n + \left(v^n + \frac{1}{2} (1 - 2\beta) a^n \Delta t \right) \Delta t = u^n + \hat{u} \\ \tilde{v}^{n+1} &= v^n + (1 - \gamma) a^n \Delta t = v^n + \hat{v} \end{aligned}$$

This gives the end of step displacement and velocity in terms of the predictor with a correction for the end step acceleration.

$$\begin{aligned} u^{n+1} &= \tilde{u}^{n+1} + \beta a^{n+1} \Delta t^2 \\ v^{n+1} &= \tilde{v}^{n+1} + \gamma a^{n+1} \Delta t \end{aligned}$$

The acceleration and velocity may now be expressed in terms of displacement, and ultimately in terms of the incremental displacement.

$$\begin{aligned} a^{n+1} &= \frac{1}{\beta \Delta t^2} (u^{n+1} - \tilde{u}^{n+1}) = \frac{1}{\beta \Delta t^2} (\hat{u} - \hat{\tilde{u}}) \\ v^{n+1} &= \tilde{v}^{n+1} + \frac{\gamma}{\beta \Delta t} (u^{n+1} - \tilde{u}^{n+1}) = \tilde{v}^{n+1} + \frac{\gamma}{\beta \Delta t} (\hat{u} - \hat{\tilde{u}}) \end{aligned}$$

plugging these into equation of motion for the SDOF system gives:

$$M \left(\frac{1}{\beta \Delta t^2} (\hat{u} - \hat{\tilde{u}}) \right) + C \left(\tilde{v}^{n+1} + \frac{\gamma}{\beta \Delta t} (\hat{u} - \hat{\tilde{u}}) \right) + K u^{n+1} = F_{n+1}$$

Finally, we assume Rayleigh damping for the dashpot.

$$C = a_{mass} M + a_{stiff} K$$

Of course we know that we intend to model a system of equations with many DOF. Thus the representation for the mass, spring and dashpot can be replaced by our finite element discretized equation of motion. We may express the system in context of a nonlinear residual problem

$$\begin{aligned} (R_{solid}^e)_{ai} &= \int_{\Gamma_i^e} \Phi_a t_i dA \\ &\quad - \int_{\Omega^e} \Phi_{a,j} \left(T_{ij}^{n+1} + a_{stiff} \left(\frac{\partial T_{ij}^{n+1}}{\partial \hat{u}_{bk}} \right)_{elastic} \left(\tilde{v}_{bk}^{n+1} + \frac{\gamma}{\beta \Delta t} (\hat{u}_{bk} - \hat{\tilde{u}}_{bk}) \right) \right) dV \\ &\quad + \int_{\Omega^e} \Phi_a \rho \left(b_i - \Phi_b \left(a_{mass} \left(\tilde{v}_{bi}^{n+1} + \frac{\gamma}{\beta \Delta t} (\hat{u}_{bi} - \hat{\tilde{u}}_{bi}) \right) + \frac{1}{\beta \Delta t^2} (\hat{u}_{bi} - \hat{\tilde{u}}_{bi}) \right) \right) dV, \\ \frac{\partial (R_{solid}^e)_{ai}}{\partial \hat{u}_{bj}} &= - \int_{\Omega^e} \Phi_{a,k} \left(\frac{\partial T_{ik}^{n+1}}{\partial \hat{u}_{bj}} + a_{stiff} \frac{\gamma}{\beta \Delta t} \left(\frac{\partial T_{ik}^{n+1}}{\partial \hat{u}_{bj}} \right)_{elastic} \right) dV \\ &\quad - \left(\frac{\gamma a_{mass}}{\beta \Delta t} + \frac{1}{\beta \Delta t^2} \right) \int_{\Omega^e} \rho \Phi_a \Phi_c \frac{\partial \hat{u}_{ci}}{\partial \hat{u}_{bj}} dV. \end{aligned}$$

Again, the expression for the residual equation and derivative are used to express a non-linear system of equations

$$\left(\frac{\partial (R_{solid}^e)_{ai}}{\partial \hat{u}_{bj}} \right) \Big|_{kiter}^{n+1} \left((u_{bj})_{kiter+1}^{n+1} - (u_{bj})_{kiter}^{n+1} \right) = -(R_{solid}^e)_{ai} \Big|_{kiter}^{n+1},$$

which are solved via the solver package. Note that the derivatives involving u and \hat{u} are interchangeable, as are differences between the non-linear iterations.

Explicit Dynamics Time Integration (Special Implementation of Newmark Method with gamma=0.5, beta=0)

For the Newmark Method, if gamma=0.5, beta=0, and the inertial term contains a diagonalized “mass matrix”, the update equations may be carried out without the solution of a system of equations. In this case, the update equations simplify to a non-iterative update algorithm.

First the mid-step velocity and end-of-step displacements are calculated through the update equations

$$v^{n+1/2} = v^n + a^n \left(\frac{\Delta t}{2} \right), \text{ and}$$

$$u^{n+1} = u^n + v^{n+1/2} \Delta t.$$

Then the residual equation/s are calculated, and acceleration at the end-of-step is calculated via

$$\left(M + \frac{\Delta t}{2} C \right) a^{n+1} = F_{n+1} - C v^{n+1/2} - K u^{n+1}.$$

Note that the mass matrix must be diagonal, and damping term may not include the stiffness based damping coefficient for this method, otherwise the above equation will require a system solve. Finally, the end-of-step velocities are calculated from the end of step acceleration:

$$v^{n+1} = v^{n+1/2} + a^{n+1} \left(\frac{\Delta t}{2} \right).$$

Note that the velocities may be stored at the midstep, resulting one less kinematic update. This approach is typically referred to as the “Leapfrog” method. However, in GEOSX we do not offer this option since it can cause some confusion that results from the storage of state at different points in time.

Parameters

In the preceding XML block, The *SolidMechanics_LagrangianFEM* is specified by the title of the subblock of the *Solvers* block. The following attributes are supported in the input block for *SolidMechanics_LagrangianFEM*:

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRelationName	string	NOCONTACT	Name of contact relation to enforce constraints on fracture boundary.
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
massDamping	real64	0	Value of mass based damping coefficient.
maxNumResolves	integer	10	Value to indicate how many resolves may be executed after some other event is executed. For example, if a SurfaceGenerator is specified, it will be executed after the mechanics solve. However if a new surface is generated, then the mechanics solve must be executed again due to the change in topology.
name	string	required	A name is required for any non-unique nodes
newmarkBeta	real64	0.25	Value of β in the Newmark Method for Implicit Dynamic time integration option. This should be $\text{pow}(\text{newmarkGamma}+0.5, 2.0)/4.0$ unless you know what you are doing.
newmarkGamma	real64	0.5	Value of γ in the Newmark Method for Implicit Dynamic time integration option
solidMaterialNames	string_array	required	The name of the material that should be used in the constitutive updates
stiffnessDamping	real64	0	Value of stiffness based damping coefficient.

The following data are allocated and used by the solver:

Name	Type	Registered On	Description
maxForce	real64		The maximum force contribution in the problem domain.
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
Acceleration	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the current acceleration on the nodes. This array also is used to hold the summation of nodal forces resulting from the governing equations.
IncrementalDisplacement	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the incremental displacements for the current time step on the nodes.
Mass	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the mass on the nodes.
TotalDisplacement	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the total displacements on the nodes.
Velocity	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the current velocity on the nodes.
contact-Force	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the contact force.
external-Force	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the external forces on the nodes. This includes any boundary conditions as well as coupling forces such as hydraulic forces.
uhatTilde	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the incremental displacement predictors on the nodes.
velocityTilde	real64_array2ds	<i>structure: nodeManager</i>	An array that holds the velocity predictors on the nodes.
Linear-SolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
Nonlinear-SolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

Example

An example of a valid XML block is given here:

```
<Solvers>
  <SolidMechanics_LagrangianFEM
    name="lagsolve"
    strainTheory="1"
    cflFactor="0.25"
    discretization="FE1"
    targetRegions="{ Region2 }"
    solidMaterialNames="{ shale }"/>
</Solvers>
```

5.3.3 Singlephase Flow Solver

Introduction

Here, we describe the single-phase flow solver. The role of this solver is to implement the fully implicit finite-volume discretization (mainly, accumulation and source terms, boundary conditions) of the equations governing compressible single-phase flow in porous media. This solver can be combined with the SinglePhaseWell class which handles the discrete multi-segment well model and provides source/sink terms for the fluid flow solver.

Theory

Governing Equations

This is a cell-centered Finite Volume solver for compressible single-phase flow in porous media. Fluid pressure as the primary solution variable. Darcy's law is used to calculate fluid velocity from pressure gradient. The solver currently only supports Dirichlet-type boundary conditions (BC) applied on cells or faces and Neumann no-flow type BC.

The following mass balance equation is solved in the domain:

$$\frac{\partial}{\partial t}(\phi\rho) + \nabla \cdot (\rho\mathbf{u}) + q = 0,$$

where

$$\mathbf{u} = -\frac{1}{\mu}\mathbf{k}(\nabla p - \rho\mathbf{g})$$

and ϕ is porosity, ρ is fluid density, μ is fluid viscosity, \mathbf{k} is the permeability tensor, \mathbf{g} is the gravity vector, and q is the source function (currently not supported). The details on the computation of the density and the viscosity are given in *Compressible single phase fluid model*.

When the entire pore space is filled by a single phase, we can substitute the Darcy's law into the mass balance equation to obtain the single phase flow equation

$$\frac{\partial}{\partial t}(\phi\rho) - \nabla \cdot \frac{\rho\mathbf{k}}{\mu}(\nabla p - \gamma\nabla z) + q = 0,$$

with $\gamma\nabla z = \rho\mathbf{g}$.

Discretization

Space Discretization

Let $\Omega \subset \mathbb{R}^n$, $n = 1, 2, 3$ be an open set defining the computational domain. We consider Ω meshed by element such that $\Omega = \cup_i V_i$ and integrate the single phase flow equation, described above, over each element V_i :

$$\int_{V_i} \frac{\partial}{\partial t}(\phi\rho)dV - \int_{V_i} \nabla \cdot \frac{\rho\mathbf{k}}{\mu}(\nabla p - \gamma\nabla z)dV + \int_{V_i} qdV = 0.$$

Applying the divergence theorem to the second term leads to

$$\int_{V_i} \frac{\partial}{\partial t}(\phi\rho)_i - \oint_{S_i} \left(\frac{\rho\mathbf{k}}{\mu}(\nabla p - \gamma\nabla z) \right) \cdot \mathbf{n}dS + \int_{V_i} qdV = 0.$$

where S_i represents the surface area of the element V_i and \mathbf{n} is a outward unit vector normal to the surface.

For the flux term, the (static) transmissibility is currently computed with a Two-Point Flux Approximation (TPFA) as described in FiniteVolume.

The pressure-dependent mobility $\lambda = \frac{\rho}{\mu}$ at the interface is approximated using a first-order upwinding on the sign of the potential difference.

Time Discretization

Let $t_0 < t_1 < \dots < t_N = T$ be a grid discretization of the time interval $[t_0, T]$, $t_0, T \in \mathbb{R}^+$. We use the backward Euler (fully implicit) method to integrate the single phase flow equation between two grid points t_n and t_{n+1} , $n < N$ to obtain the residual equation:

$$\int_{V_i} \frac{(\phi\rho)_i^{n+1} - (\phi\rho)_i^n}{\Delta t} - \oint_{S_i} \left(\frac{\rho\mathbf{k}}{\mu}(\nabla p - \gamma\nabla z) \right)^{n+1} \cdot \mathbf{n}dS + \int_{V_i} q^{n+1}dV = 0$$

where $\Delta t = t_{n+1} - t_n$ is the time-step. The expression of this residual equation and its derivative are used to form a linear system, which is solved via the solver package.

Parameters

The solver is enabled by adding a `<SinglePhaseFVM>` node in the Solvers section. Like any solver, time stepping is driven by events, see [Event Management](#).

The following attributes are supported:

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object to use for this solver.
fluidNames	string_array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
inputFluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
permeabilityNames	string_array	required	Names of permeability constitutive models for each region.
solidNames	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

In particular:

- `discretization` must point to a Finite Volume flux approximation scheme defined in the Numerical Methods section of the input file (see [FiniteVolume](#))
- `fluidName` must point to a single phase fluid model defined in the Constitutive section of the input file (see [Constitutive Models](#))
- `solidName` must point to a solid mechanics model defined in the Constitutive section of the input file (see [Constitutive Models](#))
- `targetRegions` is used to specify the regions on which the solver is applied

Primary solution field label is `pressure`. Initial conditions must be prescribed on this field in every region, and boundary conditions must be prescribed on this field on cell or face sets of interest.

Example

```
<Solvers>
  <SinglePhaseFVM
    name="SinglePhaseFlow"
    logLevel="1"
    discretization="singlePhaseTPFA"
    fluidNames="{ water }"
    solidNames="{ rock }"
    permeabilityNames="{ rockPerm }"
    targetRegions="{ mainRegion }">
    <NonlinearSolverParameters
      newtonTol="1.0e-6"
      newtonMaxIter="8"/>
    <LinearSolverParameters
      solverType="gmres"
      krylovTol="1.0e-10"/>
  </SinglePhaseFVM>
</Solvers>
```

We refer the reader to [this page](#) for a complete tutorial illustrating the use of this solver.

5.3.4 Compositional Multiphase Flow Solver

Introduction

This flow solver is in charge of implementing the finite-volume discretization (mainly, accumulation and flux terms, boundary conditions) of the equations governing compositional multiphase flow in porous media. The present solver can be combined with the [Compositional Multiphase Well Solver](#) which handles the discrete multi-segment well model and provides source/sink terms for the fluid flow solver.

Below, we first review the set of [Governing Equations](#), followed by a discussion of the choice of [Primary Variables](#) used in the global variable formulation. Then we give an overview of the [Discretization](#) and, finally, we provide a list of the solver [Parameters](#) and an input [Example](#).

Theory

Governing Equations

Mass Conservation Equations

Mass conservation for component c is expressed as:

$$\phi \frac{\partial}{\partial t} \left(\sum_{\ell} \rho_{\ell} y_{c\ell} S_{\ell} \right) + \nabla \cdot \left(\sum_{\ell} \rho_{\ell} y_{c\ell} \mathbf{u}_{\ell} \right) - \sum_{\ell} \rho_{\ell} y_{c\ell} q_{\ell} = 0,$$

where ϕ is the porosity of the medium, S_{ℓ} is the saturation of phase ℓ , $y_{c\ell}$ is the mass fraction of component c in phase ℓ , ρ_{ℓ} is the phase density, and t is time. We note that the formulation currently implemented in GEOSX is isothermal.

Darcy's Law

Using the multiphase extension of Darcy's law, the phase velocity \mathbf{u}_ℓ is written as a function of the phase potential gradient $\nabla \Phi_\ell$:

$$\mathbf{u}_\ell := -\mathbf{k} \lambda_\ell \nabla \Phi_\ell = -\mathbf{k} \lambda_\ell (\nabla (p - P_{c,\ell}) - \rho_\ell g \nabla z).$$

In this equation, \mathbf{k} is the rock permeability, $\lambda_\ell = k_{r\ell} / \mu_\ell$ is the phase mobility, defined as the phase relative permeability divided by the phase viscosity, p is the reference pressure, $P_{c,\ell}$ is the the capillary pressure, g is the gravitational acceleration, and z is depth. The evaluation of the relative permeabilities, capillary pressures, and viscosities is reviewed in the section about [Constitutive Models](#).

Combining the mass conservation equations with Darcy's law yields a set of n_c equations written as:

$$\phi \frac{\partial}{\partial t} \left(\sum_\ell \rho_\ell y_{c\ell} S_\ell \right) - \nabla \cdot \mathbf{k} \left(\sum_\ell \rho_\ell y_{c\ell} \lambda_\ell \nabla \Phi_\ell \right) - \sum_\ell \rho_\ell y_{c\ell} q_\ell = 0.$$

Constraints and Thermodynamic Equilibrium

The volume constraint equation states that the pore space is always completely filled by the phases. The constraint can be expressed as:

$$\sum_\ell S_\ell = 1.$$

The system is closed by the following thermodynamic equilibrium constraints:

$$f_{c\ell} - f_{cm} = 0.$$

where $f_{c\ell}$ is the fugacity of component c in phase ℓ . The flash calculations performed to enforce the thermodynamical equilibrium are reviewed in the section about [Constitutive Models](#).

To summarize, the compositional multiphase flow solver assembles a set of $n_c + 1$ equations in each element, i.e., n_c mass conservation equations and one volume constraint equation. A separate module discussed in the [Constitutive Models](#) is responsible for the enforcement of the thermodynamic equilibrium at each nonlinear iteration.

Number of equations	Equation type
n_c	Mass conservation equations
1	Volume constraint

Primary Variables

The variable formulation implemented in GEOSX is a global variable formulation based on $n_c + 1$ primary variables, namely, one pressure, p , and n_c component densities, ρ_c . By default, we use molar component densities. A flag discussed in the section [Parameters](#) can be used to select mass component densities instead of molar component densities.

Number of primary variables	Variable type
1	Pressure
n_c	Component densities

Assembling the residual equations and calling the *Constitutive Models* requires computing the molar component fractions and saturations. This is done with the relationship:

$$z_c := \frac{\rho_c}{\rho_T},$$

where

$$\rho_T := \sum_c \rho_c.$$

These secondary variables are used as input to the flash calculations. After the flash calculations, the saturations are computed as:

$$S_\ell := \nu_\ell \frac{\rho_T}{\rho_\ell},$$

where ν_ℓ is the global mole fraction of phase ℓ and ρ_ℓ is the molar density of phase ℓ . These steps also involve computing the derivatives of the component fractions and saturations with respect to the pressure and component densities.

Discretization

Spatial Discretization

The governing equations are discretized using standard cell-centered finite-volume discretization.

In the approximation of the flux term at the interface between two control volumes, the calculation of the pressure stencil is general and will ultimately support a Multi-Point Flux Approximation (MPFA) approach. The current implementation of the transmissibility calculation is reviewed in the section about `/coreComponents/finiteVolume/docs/FiniteVolume`.

The approximation of the dynamic transport coefficients multiplying the discrete potential difference (e.g., the phase mobilities) is performed with a first-order phase-per-phase single-point upwinding based on the sign of the phase potential difference at the interface.

Temporal Discretization

The compositional multiphase solver uses a fully implicit (backward Euler) temporal discretization.

Solution Strategy

The nonlinear solution strategy is based on Newton's method. At each Newton iteration, the solver assembles a residual vector, R , collecting the n_c discrete mass conservation equations and the volume constraint for all the control volumes.

Parameters

The following attributes are supported:

Name	Type	De- fault	Description
allowLocalCompDensityChopping	integer	1	Flag indicating whether local (cell-wise) chopping of negative compositions is allowed
capPressureNames	string_array		Name of the capillary pressure constitutive model to use
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
computeCFLNumbers	integer	0	Flag indicating whether CFL numbers are computed or not
discretization	string	required	Name of discretization object to use for this solver.
fluidNames	string_array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
inputFluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
maxCompFractionChange	real64	1	Maximum (absolute) change in a component fraction between two Newton iterations
name	string	required	A name is required for any non-unique nodes
permeabilityNames	string_array	required	Names of permeability constitutive models for each region.
relPermNames	string_array	required	Name of the relative permeability constitutive model to use
solidNames	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
temperature	real64	required	Temperature
useMass	integer	0	Use mass formulation instead of molar
Linear-SolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-SolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

Example

```

<Solvers>
  <CompositionalMultiphaseFVM
    name="compflow"
    logLevel="1"
    discretization="fluidTPFA"
    targetRegions="{ Channel }"
    fluidNames="{ fluid }"
    solidNames="{ rock }"
    permeabilityNames="{ rockPerm }"
    relPermNames="{ relperm }"
    maxCompFractionChange="0.3"
    temperature="300">
    <NonlinearSolverParameters
      newtonTol="1.0e-10"
      newtonMaxIter="15"/>
    <LinearSolverParameters
      directParallel="0"/>
  </CompositionalMultiphaseFVM>
</Solvers>

```

We refer the reader to *Multiphase Flow* for a complete tutorial illustrating the use of this solver.

5.3.5 Compositional Multiphase Well Solver

Introduction

Here, we present a description of the well solvers. These solvers are designed to be coupled with the flow solvers. Their specific task is to implement the multi-segment well discretization using the fluid model used in the corresponding flow solver – i.e., either single-phase flow or compositional multiphase flow. In particular, the perforation terms computed by the well solvers are a source/sink term for the discrete reservoir equations assembled by the flow solvers.

In the present description, we focus on the compositional multiphase well solver. The structure of the single-phase well solver is analogous and will not be described here for brevity.

Theory

Here, we give an overview of the well formulation implemented in GEOSX. We review the set of *Discrete Equations*, and then we describe the *Primary variables* used the well solvers.

Discrete Equations

We assume that the well is discretized into segments denoted by the index i .

Mass Conservation

In well segment i , mass conservation for component c reads:

$$z_{c,(i-1,i)}^{upw} q_{(i-1,i)} - z_{c,(i,i+1)}^{upw} q_{(i,i+1)} + q_{c,i}^{perf} = 0,$$

where we have neglected the accumulation term. In the previous equation, $z_{c,(i,j)}^{upw}$ is the upwinded mass fraction of component c at the interface between segments i and j , $q_{(i,j)}$ is the total mass flux between segments i and j , and

$q_{c,i}^{perf}$ is the source/sink term generated by the perforations – i.e., the connections between the well and the reservoir. The upwinded mass fractions are computed as:

$$z_{c,(i,j)}^{upw} = \begin{cases} z_{c,i} & \text{if } q_{(i,j)} > 0 \\ z_{c,j} & \text{otherwise.} \end{cases}$$

The perforation terms are obtained with:

$$q_{c,i}^{perf} = \begin{cases} WI z_{c,i} \rho_{m,i} \lambda_T^{res} \Delta\Phi & \text{if the well is upstream (i.e., } \Delta\Phi > 0) \\ WI x_{c,\ell}^{res} \rho_{\ell}^{res} \lambda_{\ell}^{res} \Delta\Phi & \text{otherwise,} \end{cases}$$

where $\Delta\Phi = p_i - p^{res} + \rho_{m,i} g \Delta d_{i,perf}$ is the potential difference between the segment center and the reservoir center. In the expression of the potential difference, the mixture density is computed as $\rho_{m,i} = \sum_{\ell} S_{\ell,i} \rho_{\ell,i}$. The well index, WI , is currently an input of the simulator. The superscript *res* means that the variable is evaluated at the center of the reservoir element.

Volume Constraint Equation

As in the *Compositional Multiphase Flow Solver*, the system is closed with a volume constraint equation.

Pressure Relations

In the current implementation of the well solver, we assume a hydrostatic equilibrium:

$$p_{i+1} - p_i = \rho_{m,(i,i+1)} g \Delta d_{i,i+1},$$

where $\rho_{m,(i,i+1)}$ is the arithmetic average of the mixture densities evaluated in segments i and $i + 1$. Pressure drop components due to friction and acceleration are not implemented at the moment.

Pressure and Rate Controls

The well solver supports two types of control, namely, pressure control and rate control.

If pressure control is chosen, we add the following constraint for the pressure of the top segment of the well:

$$p_0 - p^{target} = 0$$

In this case, we check that at each iteration of the Newton solver, the rate at the top of the first segment is smaller than the maximum rate specified by the user. If this is not the case, we switch to rate control.

If rate control is used, we add the following constraint for the rate at the top of the first segment, denoted by $q_{(-1,0)}$:

$$q_{(-1,0)} - q^{target} = 0$$

If the pressure at the top segment becomes larger than the maximum pressure specified by the user, then we switch to pressure control.

To summarize, the compositional multiphase flow solver assembles a set of $n_c + 2$ equations, i.e., n_c mass conservation equations and 1 volume constraint equation in each segment, plus 1 pressure relation at the interface between a segment and the next segment in the direction of the well head. For the top segment, the pressure relation is replaced with the control equation.

Number of equations	Equation type
n_c	Mass conservation equations
1	Pressure relation or control equation
1	Volume constraint

Primary variables

The well variable formulation is the same as that of the *Compositional Multiphase Flow Solver*. In a well segment, in addition to the $n_c + 1$ primary variables of the *Compositional Multiphase Flow Solver*, namely, one pressure, p , and n_c component densities, ρ_c , we also treat the total mass flux at the interface with the next segment, denoted by q , as a primary variable.

Number of primary variables	Variable type
1	Pressure
1	Total mass flux at the interface with next segment
n_c	Component densities

Parameters

The following attributes are supported:

Name	Type	De- fault	Description
allowLocalCompDensityChopping	integer	1	Flag indicating whether local (cell-wise) chopping of negative compositions is allowed
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fluidNames	string array	required	Name of fluid constitutive object to use for this solver.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
maxCompFracChange	real64	1	Maximum (absolute) change in a component fraction between two Newton iterations
maxRelativePressureChange	real64	1	Maximum (relative) change in pressure between two Newton iterations (recommended with rate control)
name	string	required	A name is required for any non-unique nodes
relPermNames	string array	required	Names of relative permeability constitutive models to use
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
useMass	integer	0	Use mass formulation instead of molar
wellTemperature	real64	required	Temperature
Linear-SolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-SolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>
WellControls	node		<i>Element: WellControls</i>

Example

```
<CompositionalMultiphaseWell
  name="compositionalMultiphaseWell"
  logLevel="1"
  targetRegions="{ wellRegion1, wellRegion2, wellRegion3 }"
  fluidNames="{ fluid1 }"
```

(continues on next page)

(continued from previous page)

```

relPermNames="{ relperm }"
wellTemperature="297.15">
<WellControls
  name="wellControls1"
  type="producer"
  control="BHP"
  referenceElevation="0.5"
  targetBHP="4e6"
  targetPhaseRate="1e-3"
  targetPhaseName="oil"/>
<WellControls
  name="wellControls2"
  type="producer"
  control="phaseVolRate"
  referenceElevation="0.5"
  targetBHP="2e6"
  targetPhaseRate="2.5e-7"
  targetPhaseName="oil"/>
<WellControls
  name="wellControls3"
  type="injector"
  control="totalVolRate"
  referenceElevation="0.5"
  targetBHP="4e7"
  targetTotalRate="5e-7"
  injectionStream="{ 0.1, 0.1, 0.1, 0.7 }"/>
</CompositionalMultiphaseWell>

```

5.3.6 Poromechanics Solver

Introduction

This section describes the use of the poroelasticity models implemented in GEOSX.

Theory

Governing Equations

In our model, the geomechanics (elasticity) equation is expressed in terms of the total stress σ :

$$\nabla \sigma + \rho_b \mathbf{g} = 0$$

where it relates to effective stress σ' and pore pressure p through Biot's coefficient b :

$$\sigma = \sigma' - bp\mathbf{I}$$

The fluid mass conservation equation is expressed in terms of pore pressure and volumetric (mean) total stress:

$$\left(\frac{1}{M} + \frac{b^2}{K_{dr}} \right) \frac{\partial p}{\partial t} + \frac{b}{K_{dr}} \frac{\partial \sigma_v}{\partial t} + \nabla \cdot \mathbf{v}_f = f$$

where M is the Biot's modulus and K_{dr} is the drained bulk modulus.

Unlike the conventional reservoir model that uses Lagrange’s porosity, in the coupled geomechanics and flow model, Euler’s porosity ϕ is adopted so the porosity variation is derived as:

$$\partial\phi = \left(\frac{b - \phi}{K_s} \right) \partial p + (b - \phi) \partial \epsilon_v$$

where K_s is the bulk modulus of the solid grain and ϵ_v is the volumetric strain.

Parameters

The poroelasticity model is implemented as a main solver listed in `<Solvers>` block of the input XML file that calls both `SolidMechanicsLagrangianSSLE` and `SinglePhaseFlow` solvers. In the main solver, it requires the specification of `solidSolverName`, `fluidSolverName`, and `couplingTypeOption`.

The following attributes are supported:

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a <code>FiniteElement</code> should be specified. If this is a Finite Volume Method, the name of a <code>FiniteVolume</code> discretization should be specified.
fluid-Solver-Name	string	required	Name of the fluid mechanics solver to use in the poromechanics solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
porous-Material-Names	string array	required	The name of the material that should be used in the constitutive updates
solid-Solver-Name	string	required	Name of the solid mechanics solver to use in the poromechanics solver
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the <code>EventManager</code> .
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

- `couplingTypeOption`: defines the coupling scheme.

The solid constitutive model used here is `PoroLinearElasticIsotropic`, which derives from `ElasticIsotropic` and includes an additional parameter: Biot's coefficient. The fluid constitutive model is the same as `SinglePhaseFlow` solver. For the parameter setup of each individual solver, please refer to the guideline of the specific solver.

An example of a valid XML block for the constitutive model is given here:

```
<Constitutive>
  <PorousElasticIsotropic
    name="porousRock"
    solidModelName="skeleton"
    porosityModelName="skeletonPorosity"
    permeabilityModelName="skeletonPerm"/>

  <ElasticIsotropic
    name="skeleton"
    defaultDensity="0"
    defaultYoungModulus="1.0e4"
    defaultPoissonRatio="0.2"/>

  <CompressibleSinglePhaseFluid
    name="fluid"
    defaultDensity="1"
    defaultViscosity="1.0"
    referencePressure="0.0"
    referenceDensity="1"
    compressibility="0.0e0"
    referenceViscosity="1"
    viscosibility="0.0"/>

  <BiotPorosity
    name="skeletonPorosity"
    grainBulkModulus="1.0e27"
    defaultReferencePorosity="0.3"/>

  <ConstantPermeability
    name="skeletonPerm"
    permeabilityComponents="{ 1.0e-4, 1.0e-4, 1.0e-4 }"/>
</Constitutive>
```

Example

```
<SinglePhasePoromechanics
  name="PoroelasticitySolver"
  solidSolverName="LinearElasticitySolver"
  fluidSolverName="SinglePhaseFlowSolver"
  porousMaterialNames="{ porousRock }"
  logLevel="1"
  discretization="FE1"
  targetRegions="{ Domain }">
  <LinearSolverParameters
    directParallel="0"/>
</SinglePhasePoromechanics>

<SolidMechanicsLagrangianSSLE
  name="LinearElasticitySolver"
```

(continues on next page)

(continued from previous page)

```

timeIntegrationOption="QuasiStatic"
logLevel="1"
discretization="FE1"
targetRegions="{ Domain }"
solidMaterialNames="{ skeleton }"/>

<SinglePhaseFVM
  name="SinglePhaseFlowSolver"
  logLevel="1"
  discretization="singlePhaseTPFA"
  targetRegions="{ Domain }"
  fluidNames="{ fluid }"
  solidNames="{ porousRock }"
  permeabilityNames="{ skeletonPerm }"/>
</Solvers>

```

5.3.7 Proppant Transport Solver

Introduction

The *ProppantTransport* solver applies the finite volume method to solve the equations of proppant transport in hydraulic fractures. The behavior of proppant transport is described by a continuum formulation. Here we briefly outline the usage, governing equations and numerical implementation of the proppant transport model in GEOSX.

Theory

The following mass balance and constitutive equations are solved inside fractures,

Proppant-fluid Slurry Flow

$$\frac{\partial}{\partial t}(\rho_m) + \nabla \cdot (\rho_m \mathbf{u}_m) = 0,$$

where the proppant-fluid mixture velocity \mathbf{u}_m is approximated by the Darcy's law as,

$$\mathbf{u}_m = -\frac{K_f}{\mu_m}(\nabla p - \rho_m \mathbf{g}),$$

and p is pressure, ρ_m and μ_m are density and viscosity of the mixed fluid, respectively, and \mathbf{g} is the gravity vector. The fracture permeability K_f is determined based on fracture aperture a as

$$K_f = \frac{a^2}{12}$$

Proppant Transport

$$\frac{\partial}{\partial t}(c) + \nabla \cdot (c \mathbf{u}_p) = 0,$$

in which c and \mathbf{u}_p represent the volume fraction and velocity of the proppant particles.

Multi-component Fluid Transport

$$\frac{\partial}{\partial t}[\rho_i \omega_i (1 - c)] + \nabla \cdot [\rho_i \omega_i (1 - c) \mathbf{u}_f] = 0.$$

Here \mathbf{u}_f represents the carrying fluid velocity. ρ_i and ω_i denote the density and concentration of i -th component in fluid, respectively. The fluid density ρ_f can now be readily written as

$$\rho_f = \sum_{i=1}^{N_c} \rho_i \omega_i,$$

where N_c is the number of components in fluid. Similarly, the fluid viscosity μ_f can be calculated by the mass fraction weighted average of the component viscosities.

The density and velocity of the slurry fluid are further expressed as,

$$\rho_m = (1 - c)\rho_f + c\rho_p,$$

and

$$\rho_m \mathbf{u}_m = (1 - c)\rho_f \mathbf{u}_f + c\rho_p \mathbf{u}_p,$$

in which ρ_f and \mathbf{u}_f are the density and velocity of the carrying fluid, and ρ_p is the density of the proppant particles.

Proppant Slip Velocity

The proppant particle and carrying fluid velocities are related by the slip velocity \mathbf{u}_{slip} ,

$$\mathbf{u}_{slip} = \mathbf{u}_p - \mathbf{u}_f.$$

The slip velocity between the proppant and carrying fluid includes gravitational and collisional components, which take account of particle settling and collision effects, respectively.

The gravitational component of the slip velocity \mathbf{u}_{slipG} is written as a form as

$$\mathbf{u}_{slipG} = F(c)\mathbf{u}_{settling},$$

where $\mathbf{u}_{settling}$ is the settling velocity for a single particle, d_p is the particle diameter, and $F(c)$ is the correction factor to the particle settling velocity in order to account for hindered settling effects as a result of particle-particle interactions,

$$F(c) = e^{-\lambda_s c},$$

with the hindered settling coefficient λ_s as an empirical constant set to 5.9 by default (Barree & Conway, 1995).

The settling velocity for a single particle, $\mathbf{u}_{settling}$, is calculated based on the Stokes drag law by default,

$$\mathbf{u}_{settling} = (\rho_p - \rho_f) \frac{d_p^2}{18\mu_f} \mathbf{g}.$$

Single-particle settling under intermediate Reynolds-number and turbulent flow conditions can also be described respectively by the Allen's equation (Barree & Conway, 1995),

$$\mathbf{u}_{settling} = 0.2d_p^{1.18} \left[\frac{g(\rho_p - \rho_f)}{\rho_f} \right]^{0.72} \left(\frac{\rho_f}{\mu_f} \right)^{0.45} \mathbf{e},$$

and Newton's equation (Barree & Conway, 1995),

$$\mathbf{u}_{settling} = 1.74d_p^{0.5} \left[\frac{g(\rho_p - \rho_f)}{\rho_f} \right]^{0.5} \mathbf{e}.$$

\mathbf{e} is the unit gravity vector and d_p is the particle diameter.

The collisional component of the slip velocity is modeled by defining λ , the ratio of the particle velocity to the volume averaged mixture velocity as a function of the proppant concentration. From this the particle slip velocity in horizontal direction is related to the mixed fluid velocity by,

$$\mathbf{u}_{slipH} = \frac{\lambda - 1}{1 - c} \mathbf{v}_m$$

with \mathbf{v}_m denoting volume averaged mixture velocity. We use a simple expression of λ proposed by Barree & Conway (1995) to correct the particle slip velocity in horizontal direction,

$$\lambda = [\alpha - |c - c_{slip}|^\beta]$$

where α and β are empirical constants, c_{slip} is the volume fraction exhibiting the greatest particle slip. By default the model parameters are set to the values given in (Barree & Conway, 1995): $\alpha = 1.27$, $c_{slip} = 0.1$ and $\beta = 1.5$. This model can be extended to account for the transition to the particle pack as the proppant concentration approaches the jamming transition.

Proppant Bed Build-up and Load Transport

In addition to suspended particle flow the GEOSX has the option to model proppant settling into an immobile bed at the bottom of the fracture. As the proppant cannot settle further down the proppant bed starts to form and develop at the element that is either at the bottom of the fracture or has an underlying element already filled with particles. Such an “inter-facial” element is divided into proppant flow and immobile bed regions based on the proppant-pack height.

Although proppant becomes immobile fluid can continue to flow through the settled proppant pack. The pack permeability K is defined based on the Kozeny-Carmen relationship:

$$K = \frac{(sd_p)^2}{180} \frac{\phi^3}{(1 - \phi)^2}$$

and

$$\phi = 1 - c_s$$

where ϕ is the porosity of particle pack and c_s is the saturation or maximum fraction for proppant packing, s is the sphericity and d_p is the particle diameter.

The growth of the settled pack in an “inter-facial” element is controlled by the interplay between proppant gravitational settling and shear-force induced lifting as (Hu et al., 2018),

$$\frac{dH}{dt} = \frac{cu_{settling}F(c)}{c_s} - \frac{Q_{lift}}{Ac_s},$$

where H , t , c_s , Q_{lift} , and A represent the height of the proppant bed, time, saturation or maximum proppant concentration in the proppant bed, proppant-bed load (wash-out) flux, and cross-sectional area, respectively.

The rate of proppant bed load transport (or wash out) due to shear force is calculated by the correlation proposed by Wiberg and Smith (1989) and McClure (2018),

$$Q_{lift} = a \left(d_p \sqrt{\frac{gd_p(\rho_p - \rho_f)}{\rho_f}} \right) (9.64N_{sh}^{0.166})(N_{sh} - N_{sh,c})^{1.5}.$$

a is fracture aperture, and N_{sh} is the Shields number measuring the relative importance of the shear force to the gravitational force on a particle of sediment (Miller et al., 1977; Biot & Medlin, 1985; McClure, 2018) as

$$N_{sh} = \frac{\tau}{d_p g (\rho_p - \rho_f)},$$

and

$$\tau = 0.125 f \rho_f u_m^2$$

where τ is the shear stress acting on the top of the proppant bed and f is the Darcy friction coefficient. $N_{sh,c}$ is the critical Shields number for the onset of bed load transport.

Proppant Bridging and Screenout

Proppant bridging occurs when proppant particle size is close to or larger than fracture aperture. The aperture at which bridging occurs, h_b , is defined simply by

$$h_b = \lambda_b d_p,$$

in which λ_b is the bridging factor.

Slurry Fluid Viscosity

The viscosity of the bulk fluid, μ_m , is calculated as a function of proppant concentration as (Keck et al., 1992),

$$\mu_m = \mu_f \left[1 + 1.25 \left(\frac{c}{1 - c/c_s} \right) \right]^2.$$

Note that continued model development and improvement are underway and additional empirical correlations or functions will be added to support the above calculations.

Spatial Discretization

The above governing equations are discretized using a cell-centered two-point flux approximation (TPFA) finite volume method. We use an upwind scheme to approximate proppant and component transport across cell interfaces.

Solution Strategy

The discretized non-linear slurry flow and proppant/component transport equations at each time step are separately solved by the Newton-Raphson method. The coupling between them is achieved by a time-marching sequential (operator-splitting) solution approach.

Parameters

The solver is enabled by adding a `<ProppantTransport>` node and a `<SurfaceGenerator>` node in the Solvers section. Like any solver, time stepping is driven by events, see [Event Management](#).

The following attributes are supported:

Name	Type	Default	Description
bridgingFactor	real64	0	Bridging factor used for bridging/screen-out calculation
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
critical-Shields-Number	real64	0	Critical Shields number
discretization	string	required	Name of discretization object to use for this solver.
fluid-Names	string_array	required	Names of fluid constitutive models for each region.
frictionCoefficient	real64	0.03	Friction coefficient
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
input-FluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
max-Proppant-Concentration	real64	0.6	Maximum proppant concentration
name	string	required	A name is required for any non-unique nodes
permeability-Names	string_array	required	Names of permeability constitutive models for each region.
proppant-Density	real64	2500	Proppant density
proppantDiameter	real64	0.0004	Proppant diameter
proppant-Names	string_array	required	Name of proppant constitutive object to use for this solver.
solid-Names	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
updateProppant-Packing	integer	0	Flag that enables/disables proppant-packing update
Linear-Solver-	node	unique	<div>Chapter 5. User Guide</div> <div>Element: LinearSolverParameters</div>

In particular:

- `discretization` must point to a Finite Volume flux approximation scheme defined in the Numerical Methods section of the input file (see [FiniteVolume](#))
- `proppantName` must point to a particle fluid model defined in the Constitutive section of the input file (see [Constitutive Models](#))
- `fluidName` must point to a slurry fluid model defined in the Constitutive section of the input file (see [Constitutive Models](#))
- `solidName` must point to a solid mechanics model defined in the Constitutive section of the input file (see [Constitutive Models](#))
- `targetRegions` attribute is currently not supported, the solver is always applied to all regions.

Primary solution field labels are `proppantConcentration` and `pressure`. Initial conditions must be prescribed on these field in every region, and boundary conditions must be prescribed on these fields on cell or face sets of interest. For static (non-propagating) fracture problems, the fields `ruptureState` and `elementAperture` should be provided in the initial conditions.

In addition, the solver declares a scalar field named `referencePorosity` and a vector field named `permeability`, that contains principal values of the symmetric rank-2 permeability tensor (tensor axis are assumed aligned with the global coordinate system). These fields must be populated via [Element: FieldSpecification](#) section and `permeability` should be supplied as the value of `coefficientName` attribute of the flux approximation scheme used.

Example

First, we specify the proppant transport solver itself and apply it to the fracture region:

```
<ProppantTransport
  name="ProppantTransport"
  logLevel="1"
  discretization="singlePhaseTPFA"
  targetRegions="{ Fracture }"
  fluidNames="{ water }"
  proppantNames="{ sand }"
  solidNames="{ fractureFilling }"
  permeabilityNames="{ fracturePerm }">
  <NonlinearSolverParameters
    newtonTol="1.0e-8"
    newtonMaxIter="8"
    lineSearchAction="None"/>
  <LinearSolverParameters
    directParallel="0"/>
</ProppantTransport>
```

Then, we specify a compatible flow solver (currently a specialized `SinglePhaseProppantFVM` solver must be used):

```
<SinglePhaseProppantFVM
  name="SinglePhaseFVM"
  logLevel="1"
  discretization="singlePhaseTPFA"
  targetRegions="{ Fracture }"
  fluidNames="{ water }"
  solidNames="{ fractureFilling }"
```

(continues on next page)

(continued from previous page)

```

permeabilityNames="{ fracturePerm }">
  <NonlinearSolverParameters
    newtonTol="1.0e-8"
    newtonMaxIter="8"
    lineSearchAction="None"/>
  <LinearSolverParameters
    solverType="gmres"
    krylovTol="1.0e-12"/>
</SinglePhaseProppantFVM>

```

Finally, we couple them through a coupled solver that references the two above:

```

<FlowProppantTransport
  name="FlowProppantTransport"
  proppantSolverName="ProppantTransport"
  flowSolverName="SinglePhaseFVM"
  targetRegions="{ Fracture }"
  logLevel="1"/>

```

References

- R. D. Barree & M. W. Conway. “Experimental and numerical modeling of convective proppant transport”, JPT. Journal of petroleum technology, 47(3):216-222, 1995.
- M. A. Biot & W. L. Medlin. “Theory of Sand Transport in Thin Fluids”, Paper presented at the SPE Annual Technical Conference and Exhibition, Las Vegas, NV, 1985.
- X. Hu, K. Wu, X. Song, W. Yu, J. Tang, G. Li, & Z. Shen. “A new model for simulating particle transport in a low-viscosity fluid for fluid-driven fracturing”, AIChE J. 64 (9), 35423552, 2018.
- R. G. Keck, W. L. Nehmer, & G. S. Strumolo. “A new method for predicting friction pressures and rheology of proppant-laden fracturing fluids”, SPE Prod. Eng., 7(1):21-28, 1992.
- M. McClure. “Bed load proppant transport during slickwater hydraulic fracturing: insights from comparisons between published laboratory data and correlations for sediment and pipeline slurry transport”, J. Pet. Sci. Eng. 161 (2), 599610, 2018.
- M. C. Miller, I. N. McCave, & P. D. Komar. “Threshold of sediment motion under unidirectional currents”, Sedimentology 24 (4), 507527, 1977.
- P. L. Wiberg & J. D. Smith. “Model for calculating bed load transport of sediment”, J. Hydraul. Eng. 115 (1), 101123, 1989.

5.4 Constitutive Models

Constitutive models describe relations between various physical quantities. In a physics simulation they are used to model the response or state of material (solid, fluid, or a mixture) as a function of input variables. There are many types of constitutive models available in GEOSX. These models are grouped together based on their input/output interface.

5.4.1 Solid Models

A solid model governs the relationship between deformation and stress in a solid (or porous) material. GEOSX provides interfaces for both small and large deformation models, as well as specific implementations of a number of well known models.

Deformation Theories

Table of Contents

- *Deformation Theories*
 - *Introduction*
 - *Small Strain Models*
 - *Finite Deformation Models with Hypo-Materials*
 - *Finite Deformation Models with Hyper-Materials*

Introduction

The solid mechanics solvers in GEOSX work in a time-discrete setting, in which the system state at time t^n is fully known, and the goal of the solution procedure is to advance forward one timestep to $t^{n+1} = t^n + \Delta t$. As part of this process, calls to a solid model must be made to compute the updated stress σ^{n+1} resulting from incremental deformation over the timestep. History-dependent models may also need to compute updates to one or more internal state variables Q^{n+1} .

The exact nature of the incremental update will depend, however, on the kinematic assumptions made. Appropriate measures of deformation and stress depend on assumptions of *infinitesimal* or *finite* strain, as well as other factors like rate-dependence and material anisotropy.

This section briefly reviews three main classes of solid models in GEOSX, grouped by their kinematic assumptions. The presentation is deliberately brief, as much more extensive presentations can be found in almost any textbook on linear and nonlinear solid mechanics.

Small Strain Models

Let u denote the displacement field, and ∇u its gradient. In small strain theory, one assumes the displacement gradients $\nabla u \ll 1$. In this case, it is sufficient to use the linearized strain tensor

$$\epsilon = \frac{1}{2} (\nabla u + u \nabla)$$

as the deformation measure. Higher-order terms present in finite strain theories are neglected. For inelastic problems, this strain is additively decomposed into elastic and inelastic components as

$$\epsilon = \epsilon^e + \epsilon^i.$$

Inelastic strains can arise from a number of sources: plasticity, damage, etc. Most constitutive models (including nonlinear elastic and inelastic models) can then be generically expressed in rate form as

$$\dot{\sigma} = c : \dot{\epsilon}^e$$

where $\dot{\sigma}$ is the Cauchy stress rate and c is the tangent stiffness tensor. Observe that the stress rate is driven by the elastic component $\dot{\epsilon}^e$ of the strain rate.

In the time-discrete setting (as implemented in the code) the incremental constitutive update for stress is computed from a solid model update routine as

$$\sigma^{n+1} = \sigma(\Delta \epsilon, \Delta t, Q^n),$$

where $\Delta\epsilon = \epsilon^{n+1} - \epsilon^n$ is the incremental strain, Δt is the timestep size (important for rate-dependent models), and Q^n is a collection of material state variables (which may include the previous stress and strain).

For path and rate independent models, such as linear elasticity, a simpler constitutive update may be formulated in terms of the total strain:

$$\sigma^{n+1} = \sigma(\epsilon^{n+1}).$$

GEOSX will use this latter form in specific, highly-optimized solvers when we know in advance that a linear elastic model is being applied. The more general interface is the default, however, as it can accommodate a much wider range of constitutive behavior within a common interface.

When implicit timestepping is used, the solid models must also provide the stiffness tensor,

$$c^{n+1} = \frac{\partial \sigma^{n+1}}{\partial \epsilon^{n+1}},$$

in order to accurately linearize the governing equations. In many works, this fourth-order tensor is referred to as the algorithmic or consistent tangent, in the sense that it must be “consistent” with the discrete timestepping scheme being used (Simo and Hughes 1987). For inelastic models, it depends not only on the intrinsic material stiffness, but also the incremental nature of the loading process. The correct calculation of this stiffness can have a dramatic impact on the convergence rate of Newton-type solvers used in the implicit solid mechanics solvers.

Finite Deformation Models with Hypo-Materials

In the finite deformation regime, there are two broad classes of constitutive models frequently used:

- Hypo-elastic models (and inelastic extensions)
- Hyper-elastic models (and inelastic extensions)

Hypo-materials typically rely on a rate-form of the constitutive equations expressed in the spatial configuration. Let $v(x, t)$ denote the spatial velocity field. It can be decomposed into symmetric and anti-symmetric components as

$$d = \frac{1}{2} (\nabla v + v \nabla) \quad \text{and} \quad w = \frac{1}{2} (\nabla v - v \nabla),$$

where d is the rate of deformation tensor and w is the spin tensor. A hypo-material model can be written in rate form as

$$\dot{\tau} = c : d^e$$

where $\dot{\tau}$ is an **objective rate** of the Kirchhoff stress tensor, c is the tangent stiffness tensor, and d^e is the elastic component of the deformation rate. We see that the structure is similar to the rate form in the small strain regime, except the rate of Cauchy stress is replaced with an objective rate of Kirchhoff stress, and the linearized strain rate is replaced with the rate of deformation tensor.

The key difference separating most hypo-models is the choice of the objective stress rate. In GEOSX, we adopt the incrementally objective integration algorithm proposed by Hughes & Winget (1980). This method relies on the concept of an incrementally rotating frame of reference in order to preserve objectivity of the stress rate. In particular, the stress update sequence is

$$\begin{aligned} \Delta R &= (I - \frac{1}{2} \Delta t w)^{-1} (I + \frac{1}{2} \Delta t w) && \text{(compute incremental rotation),} \\ \bar{\tau}^n &= \Delta R \tau^n \Delta R^T && \text{(rotate previous stress),} \\ \tau^{n+1} &= \bar{\tau}^n + \Delta \tau && \text{(call constitutive model to update stress).} \end{aligned}$$

First, the previous timestep stress is rotated to reflect any rigid rotations occurring over the timestep. If the model has tensor-valued state variables besides stress, these must also be rotated. Then, a standard constitutive update routine

can be called, typically driven by the incremental strain $\Delta\epsilon = \Delta t d$. In fact, an identical update routine as used for small strain models can be re-used at this point.

Note: Hypo-models suffer from several well known deficiencies. Most notably, the energy dissipation in a closed loading cycle of a hypo-elastic material is not guaranteed to be zero, as one might desire from thermodynamic considerations.

Finite Deformation Models with Hyper-Materials

Hyper-elastic models (and inelastic extensions) attempt to correct the thermodynamic deficiencies of their hypo-elastic cousins. The constitutive update can be generically expressed at

$$S^{n+1} = S(\Delta\mathbf{F}, Q^n, \Delta t),$$

where S is the second Piola-Kirchoff stress and $\Delta\mathbf{F}$ is the incremental deformation gradient. Depending on the model, the deformation gradient can be converted to different deformation measures as needed. Similarly, different stress tensors can be recovered through appropriate push-forward and pull-back operations.

In a hyperelastic material, the elastic response is expressed in terms of a stored strain-energy function that serves as the potential for stress, e.g.

$$\mathbf{S} = \frac{\partial\psi(C)}{\partial C},$$

where ψ is the stored energy potential, and C is the right Cauchy-Green deformation tensor. This potential guarantees that the energy dissipated or gained in a closed elastic cycle is zero.

Voigt Notation

In GEOSX we express rank-two symmetric tensors using [Voigt notation](#). Stress tensors are represented as an “unrolled” six-component vector storing only the unique component values. For strain tensors, note that *engineering strains* are used such that the shear components of strain are multiplied by a factor of two. With this representation, the strain energy density is, conveniently, the inner product of the stress and strain vectors.

Voigt representation of common rank-2 symmetric tensors are:

$$\sigma = \begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix}, \mathbf{S} = \begin{bmatrix} S_{11} \\ S_{22} \\ S_{33} \\ S_{23} \\ S_{13} \\ S_{12} \end{bmatrix}, \epsilon = \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{23} \\ 2\epsilon_{13} \\ 2\epsilon_{12} \end{bmatrix}, \mathbf{D} = \begin{bmatrix} D_{11} \\ D_{22} \\ D_{33} \\ 2D_{23} \\ 2D_{13} \\ 2D_{12} \end{bmatrix}, \mathbf{E} = \begin{bmatrix} E_{11} \\ E_{22} \\ E_{33} \\ 2E_{23} \\ 2E_{13} \\ 2E_{12} \end{bmatrix}, \mathbf{B} = \begin{bmatrix} B_{11} \\ B_{22} \\ B_{33} \\ 2B_{23} \\ 2B_{13} \\ 2B_{12} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} C_{11} \\ C_{22} \\ C_{33} \\ 2C_{23} \\ 2C_{13} \\ 2C_{12} \end{bmatrix},$$

where σ is the Cauchy stress, \mathbf{S} is the second Piola-Kirchhoff stress, ϵ is the small strain tensor, \mathbf{D} is the rate of deformation tensor, \mathbf{E} is the Lagrangian finite strain tensor, \mathbf{B} is the left Cauchy-Green tensor, \mathbf{C} is the right Cauchy-Green deformation tensor.

Note: The factor of two in the shear components of strain (and strain-like) quantities is a frequent source of confusion, even for expert modelers. It can be particularly challenging to use in nuanced situations like stiffness tensor calculations or invariant decompositions. If you plan to implement new models within GEOSX, please pay extra attention to this detail. We also provide many common operations in centralized helper functions to avoid re-inventing the wheel.

Plasticity Notation

Table of Contents

- *Plasticity Notation*
 - *Overview*
 - *Two-Invariant Models*
 - *Three-Invariant Models*

Overview

According to the [theory of plasticity](#) in the small strain regime, the total strain ϵ can be additively split into elastic (ϵ^e) and plastic (ϵ^p) strains:

$$\epsilon = \epsilon^e + \epsilon^p.$$

The plastic strain tensor is obtained from the flow rule:

$$\dot{\epsilon}^p = \dot{\lambda} \frac{\partial g}{\partial \sigma},$$

in which $\dot{\lambda} \geq 0$ is the magnitude of plastic strain rate and g is the plastic potential. The elastic strain is related to Cauchy stress tensor in rate form as:

$$\dot{\sigma} = c^e : \dot{\epsilon}^e,$$

where c^e is the fourth order elastic stiffness tensor. The Cauchy stress tensor is related to the total strain as

$$\dot{\sigma} = c^{ep} : \dot{\epsilon},$$

in which c^{ep} is the fourth order elasto-plastic stiffness tensor.

Two-Invariant Models

Two-invariant plasticity models use the first invariant of the Cauchy stress tensor and the second invariant of the deviatoric stress tensor to describe the yield surface.

Here we use the following stress invariants to define the yield surface: the von Mises stress $q = \sqrt{3J_2} = \sqrt{3/2}\|s\|$ and mean normal stress $p = I_1/3$. Here, I_1 and J_2 are the first invariant of the stress tensor and second invariant of the deviatoric stress, defined as

$$I_1 = tr(\sigma)/3, \quad J_2 = \frac{1}{2}\|s\|^2, \quad s = \sigma - p\mathbf{1},$$

in which $\mathbf{1}$ is the identity tensor.

Similarly, we can define invariants of strain tensor, namely, volumetric strain ϵ_v and deviatoric strain ϵ_s .

$$\epsilon_v = tr(\epsilon), \quad \epsilon_s = \sqrt{\frac{2}{3}}\|e\|, \quad \text{where} \quad e = \epsilon - \frac{1}{3}\epsilon_v\mathbf{1}.$$

Stress and strain tensors can then be recomposed from the invariants as:

$$\sigma = p\mathbf{1} + \sqrt{\frac{2}{3}}q\hat{n}$$

$$\boldsymbol{\epsilon} = \frac{1}{3}\epsilon_v \mathbf{1} + \sqrt{\frac{3}{2}}\epsilon_s \hat{\mathbf{n}}$$

in which $\hat{\mathbf{n}} = \mathbf{e}/\|\mathbf{e}\|$.

The following two-invariant models are currently implemented in GEOSX:

- Drucker-Prager
- J2 plasticity
- Modified Cam-Clay
- Delft-Egg

Three-Invariant Models

Several three-invariant models are under active development, but are not yet available in develop. If you are interested in helping to add additional material models, please submit a feature request.

Triaxial Driver

Table of Contents

- *Triaxial Driver*
 - *Introduction*
 - *XML Structure*
 - *Test Modes*
 - *Output Format*
 - *Model Convergence*
 - *Unit Testing*

Introduction

When calibrating solid material parameters to experimental data, it can be a hassle to launch a full finite element simulation to mimic experimental loading conditions. Instead, GEOSX provides a `TriaxialDriver` allowing the user to run loading tests on a single material point. This makes it easy to understand the material response and fit it to lab data. The driver itself is launched like any other GEOSX simulation, but with a particular XML structure:

```
./bin/geosx -i myTest.xml
```

XML Structure

A typical XML file to run the triaxial driver will have the following key elements. We present the whole file first, before digging into the individual blocks.

```
src/coreComponents/unitTests/constitutiveTests/testTriaxial_sphinxExample.xml
```

The first thing to note is that the XML structure is identical to a standard GEOSX input deck. In fact, once the constitutive block is calibrated, one could start adding solver and discretization blocks to the same file to create a proper field simulation. This makes it easy to go back and forth between calibration and simulation.

The `TriaxialDriver` is added as a `Task`, a particular type of executable event often used for simple actions. It is added as a `SoloEvent` to the event queue. This leads to a trivial event queue, since all we do is launch the driver and then quit.

Internally, the triaxial driver uses a simple form of time-stepping to advance through the loading steps, allowing for both rate-dependent and rate-independent models to be tested. This timestepping is handled independently from the more complicated time-stepping pattern used by physics `Solvers` and coordinated by the `EventManager`. In particular, in the XML file above, the `maxTime` parameter in the `Events` block is an event manager control, controlling when/if certain events occur. Once launched, the triaxial driver internally determines its own max time and timestep size using a combination of the strain function's time coordinates and the requested number of loadsteps. It is therefore helpful to think of the driver as an instantaneous *event* (from the event manager's point of view), but one which has a separate, internal clock.

The key parameters for the `TriaxialDriver` are:

Name	Type	Default	Description
baseline	path	none	Baseline file
logLevel	integer	0	Log level
material	string	required	Solid material to test
mode	string	required	Test mode [triaxial, volumetric, oedometer]
name	string	required	A name is required for any non-unique nodes
output	string	none	Output file
steps	integer	required	Number of load steps to take
strainFunction	string	required	Function controlling strain loading (role depends on test mode)
stressFunction	string	required	Function controlling stress loading (role depends on test mode)

Note: GEOSX uses the *engineering* sign convention where compressive stresses and strains are *negative*. This is one of the most frequent issues users make when calibrating material parameters, as stress- and strain-like quantities often need to be negative to make physical sense. You may note in the XML above, for example, that `stressFunction` and `strainFunction` have negative values for a compressive test.

Test Modes

The most complicated part of the driver is understanding how the stress and strain functions are applied in different testing modes. The driver mimics laboratory core tests, with loading controlled in the axial and radial directions. These conditions may be either strain-controlled or stress-controlled. The following table describes the available test modes in detail:

mode	axial loading	radial loading	initial stress
triaxial	axial strain controlled with <code>strainFunction</code>	radial stress controlled with <code>stressFunction</code>	isotropic stress using <code>stressFunction(t=tmin)</code>
volumetric	axial strain controlled with <code>strainFunction</code>	radial strain = axial strain	isotropic stress using <code>stressFunction(t=tmin)</code>
oedometer	axial strain controlled with <code>strainFunction</code>	zero radial strain	isotropic stress using <code>stressFunction(t=tmin)</code>

To set the initial stress state, the `stressFunction` is evaluated at `t=tmin` (usually `t=0`, though conceivably a user

may put in a time-function with a non-zero starting point). This scalar value is used to set the material to an isotropic initial stress state. In the volumetric and oedometer tests, the remainder of the `stressFunction` time history is ignored, as they are strain-controlled tests. By setting the initial stress this way, it makes it easy to start the test from a well-defined confining pressure.

Output Format

The `output` key is used to identify a file to which the results of the simulation are written. If this key is omitted, or the user specifies `output="none"`, file output will be suppressed. The file is a simple ASCII format with a brief header followed by test data:

```
# column 1 = time
# column 2 = axial_strain
# column 3 = radial_strain_1
# column 4 = radial_strain_2
# column 5 = axial_stress
# column 6 = radial_stress_1
# column 7 = radial_stress_2
# column 8 = newton_iter
# column 9 = residual_norm
0.0000e+00 0.0000e+00 0.0000e+00 0.0000e+00 -1.0000e+00 -1.0000e+00 -1.0000e+00 0.
↪0000e+00 0.0000e+00
1.6000e-01 -1.6000e-04 4.0000e-05 4.0000e-05 -1.1200e+00 -1.0000e+00 -1.0000e+00 2.
↪0000e+00 0.0000e+00
3.2000e-01 -3.2000e-04 8.0000e-05 8.0000e-05 -1.2400e+00 -1.0000e+00 -1.0000e+00 2.
↪0000e+00 0.0000e+00
...
```

This file can be readily plotted using any number of plotting tools. Each row corresponds to one timestep of the driver, starting from initial conditions in the first row.

We note that the file contains two columns for radial strain and two columns for radial stress. For an isotropic material, the stresses and strains along the two radial axes will usually be identical. We choose to output this way, however, to accommodate both anisotropic materials and true-triaxial loading conditions. In these cases, the stresses and strains in the radial directions could potentially differ.

These columns can be added and subtracted to produce other quantities of interest, like mean stress or deviatoric stress. For example, we can plot the output produce stress / strain curves (in this case for a plastic rather than simple elastic material):

Fig. 5.1: Stress/strain behavior for a plastic material.

In this plot, we have reversed the sign convention to be consistent with typical experimental plots. Note also that the `strainFunction` includes two unloading cycles, allowing us to observe both plastic loading and elastic unloading.

Model Convergence

The last two columns of the output file contain information about the convergence behavior of the material driver. In `triaxial` mode, the mixed nature of the stress/strain control requires using a Newton solver to converge the solution. This last column reports the number of Newton iterations and final residual norm. Large values here would be indicative of the material model struggling (or failing) to converge. Convergence failures can result from several reasons, including:

1. Inappropriate material parameter settings

2. Overly large timesteps
3. Infeasible loading conditions (i.e. trying to load a material to a physically-unreachable stress point)
4. Poor model implementation

We generally spend a lot of time vetting the material model implementations (#4). When you first encounter a problem, it is therefore good to explore the other three scenarios first. If you find something unusual in the model implementation or are just really stuck, please submit an issue on our issue tracker so we can help resolve any bugs.

Unit Testing

The development team also uses the Triaxial Driver to perform unit testing on the various material models within GEOSX. The optional argument `baseline` can be used to point to a previous output file that has been validated (e.g. against analytical or experimental benchmarks). If such a file is specified, the driver will perform a loading run and then compare the new results against the baseline. In this way, any regressions in the material models can be quickly identified.

Developers of new models are encouraged to add their own baselines to `src/coreComponents/constitutive/unitTests`. Adding additional tests is straightforward:

1. Create a new xml file for your test in `src/coreComponents/constitutive/unitTests`. There are several examples in this directory already to use as a template. We suggest using the naming convention `testTriaxial_myTest.xml`, so that all triaxial tests will be grouped together alphabetically. Set the output file to `testTriaxial_myTest.txt`, and run your test. Validate the results however is appropriate.
2. This output file will now become your new baseline. Replace the `output` key with `baseline` so that the driver can read in your file as a baseline for comparison. Make sure there is no remaining `output` key, or set `output=none`, to suppress further file output. While you can certainly write a new output for debugging purposes, during our automated unit tests we prefer to suppress file output. Re-run the triaxial driver to confirm that the comparison test passes.
3. Modify `src/coreComponents/constitutive/unitTests/CMakeLists.txt` to enable your new test in the unit test suite. In particular, you will need to add your new XML file to the existing list in the `gtest_triaxial_xmls` variable:

```
set( gtest_triaxial_xmls
    testTriaxial_elasticIsotropic.xml
    testTriaxial_druckerPragerExtended.xml
    testTriaxial_myTest.xml
)
```

4. Run `make` in your build directory to make sure the CMake syntax is correct
5. Run `ctest -V -R Triax` to run the triaxial unit tests. Confirm your test is included and passes properly.

If you run into troubles, do not hesitate to contact the development team for help.

Model: Elastic Isotropic

Overview

This model may be used for solid materials with a linear elastic isotropic behavior. The relationship between stress and strain is given by [Hooke's Law](#), expressed as:

$$\sigma_{ij} = \lambda \epsilon_{kk} + 2\mu \epsilon_{ij},$$

where σ_{ij} is the ij component of the Cauchy stress tensor, ϵ_{ij} is the ij component of the strain tensor, λ is the first Lamé elastic constant, and μ is the elastic shear modulus.

Hooke's Law may also be expressed using [Voigt notation](#) for stress and strain vectors as:

$$\sigma = C \cdot \epsilon,$$

or,

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} 2\mu + \lambda & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & 2\mu + \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & 2\mu + \lambda & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{23} \\ 2\epsilon_{13} \\ 2\epsilon_{12} \end{bmatrix}.$$

Variations

For finite deformation solvers, the elastic isotropic model can be called within a hypo-elastic update routine. See [Finite Deformation Models with Hypo-Materials](#)

Parameters

The following attributes are supported. Note that any two elastic constants can be provided, and the other two will be internally calculated. The “default” keyword in front of certain properties indicates that this is the default value adopted for a region unless the user separately specifies a heterogeneous field via the `FieldSpecification` mechanism.

Name	Type	Default	Description
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultDensity	real64	required	Default Material Density
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
name	string	required	A name is required for any non-unique nodes

Example

A typical `Constitutive` block will look like:

```
<Constitutive>
  <ElasticIsotropic
    name="shale"
    defaultDensity="2700"
    defaultBulkModulus="60.0e6"
    defaultShearModulus="30.0e6" />
  </Constitutive>
```

Model: Elastic Isotropic Pressure Dependent

Overview

This model may be used for solid materials with a pressure-dependent elastic isotropic behavior. The relationship between stress and strain is given by a [hyperelastic law](#). The elastic constitutive equations for the volumetric and deviatoric stresses and strain are expressed as:

$$p = p_0 \exp\left(\frac{\epsilon_{v0} - \epsilon_v^e}{c_r}\right), \quad q = 3\mu\epsilon_s^e$$

where p and q are the volumetric and deviatoric components of the Cauchy stress tensor. ϵ_v^e and ϵ_s^e are the volumetric and deviatoric components of the strain tensor. ϵ_{v0} and p_0 are the initial volumetric strain and initial pressure. c_r denotes the elastic compressibility index, and μ is the elastic shear modulus. In this model, the shear modulus is constant and the bulk modulus, K , varies linearly with pressure as follows:

$$K = -\frac{p}{c_r}$$

Parameters

The following attributes are supported. Note that two elastic constants c_r and μ , as well as the initial volumetric strain and initial pressure need to be provided. The “default” keyword in front of certain properties indicates that this is the default value adopted for a region unless the user separately specifies a heterogeneous field via the `FieldSpecification` mechanism.

Example

A typical `Constitutive` block will look like:

```
<Constitutive>
  <ElasticIsotropicPressureDependent
    name="elasticPressure"
    defaultDensity="2700"
    defaultRefPressure="-1.0"
    defaultRefStrainVol="1"
    defaultRecompressionIndex="0.003"
    defaultShearModulus="200"/>
</Constitutive>
```

Model: Elastic Transverse-Isotropic

Overview

This model may be used for solid materials with a linear elastic, transverse-isotropic behavior. This is most readily expressed in Voight notation as

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{13} & 0 & 0 & 0 \\ C_{13} & C_{13} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & (C_{11} - C_{12})/2 \end{bmatrix} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{23} \\ 2\epsilon_{13} \\ 2\epsilon_{12} \end{bmatrix}.$$

This system contains five independent constants. These constants are calculated from the input parameters indicated below.

Parameters

The following attributes are supported. The “default” keyword in front of certain properties indicates that this is the default value adopted for a region unless the user separately specifies a heterogeneous field via the FieldSpecification mechanism.

Name	Type	Default	Description
defaultC11	real64	-1	Default Stiffness Parameter C11
defaultC13	real64	-1	Default Stiffness Parameter C13
defaultC33	real64	-1	Default Stiffness Parameter C33
defaultC44	real64	-1	Default Stiffness Parameter C44
defaultC66	real64	-1	Default Stiffness Parameter C66
defaultDensity	real64	required	Default Material Density
defaultPoissonRatioAxialTransverse	real64	-1	Default Axial-Transverse Poisson’s Ratio
defaultPoissonRatioTransverse	real64	-1	Default Transverse Poisson’s Ratio
defaultShearModulusAxialTransverse	real64	-1	Default Axial-Transverse Shear Modulus
defaultYoungModulusAxial	real64	-1	Default Axial Young’s Modulus
defaultYoungModulusTransverse	real64	-1	Default Transverse Young’s Modulus
name	string	required	A name is required for any non-unique nodes

Example

A typical Constitutive block will look like:

```
<Constitutive>
  <ElasticTransverseIsotropic
    name="shale"
    defaultDensity="2700"
    defaultPoissonRatioAxialTransverse="0.20"
    defaultPoissonRatioTransverse="0.30"
    defaultYoungModulusAxial="50.0e6"
    defaultYoungModulusTransverse="60.0e6"
    defaultShearModulusAxialTransverse="30.0e6" />
</Constitutive>
```

Model: Elastic Orthotropic

Overview

This model may be used for solid materials with a linear elastic, orthotropic behavior. This is most readily expressed in Voight notation as

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{13} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{12} & C_{22} & C_{23} & 0 & 0 & 0 \\ C_{13} & C_{23} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{66} \end{bmatrix} \begin{bmatrix} \epsilon_{11} \\ \epsilon_{22} \\ \epsilon_{33} \\ 2\epsilon_{23} \\ 2\epsilon_{13} \\ 2\epsilon_{12} \end{bmatrix}.$$

This system contains nine independent constants. These constants are calculated from the input parameters indicated below.

Parameters

The following attributes are supported. The “default” keyword in front of certain properties indicates that this is the default value adopted for a region unless the user separately specifies a heterogeneous field via the FieldSpecification mechanism.

Name	Type	Default	Description
defaultC11	real64	-1	Default C11 Component of Voigt Stiffness Tensor
defaultC12	real64	-1	Default C12 Component of Voigt Stiffness Tensor
defaultC13	real64	-1	Default C13 Component of Voigt Stiffness Tensor
defaultC22	real64	-1	Default C22 Component of Voigt Stiffness Tensor
defaultC23	real64	-1	Default C23 Component of Voigt Stiffness Tensor
defaultC33	real64	-1	Default C33 Component of Voigt Stiffness Tensor
defaultC44	real64	-1	Default C44 Component of Voigt Stiffness Tensor
defaultC55	real64	-1	Default C55 Component of Voigt Stiffness Tensor
defaultC66	real64	-1	Default C66 Component of Voigt Stiffness Tensor
defaultDensity	real64	required	Default Material Density
defaultE1	real64	-1	Default Young’s Modulus E1
defaultE2	real64	-1	Default Young’s Modulus E2
defaultE3	real64	-1	Default Young’s Modulus E3
defaultG12	real64	-1	Default Shear Modulus G12
defaultG13	real64	-1	Default Shear Modulus G13
defaultG23	real64	-1	Default Shear Modulus G23
defaultNu12	real64	-1	Default Poission’s Ratio Nu12
defaultNu13	real64	-1	Default Poission’s Ratio Nu13
defaultNu23	real64	-1	Default Poission’s Ratio Nu23
name	string	required	A name is required for any non-unique nodes

Example

A typical Constitutive block will look like:

```
<Constitutive>
  <ElasticOrthotropic
    name="shale"
    defaultDensity="2700"
    defaultNu12="0.20"
    defaultNu13="0.25"
    defaultNu23="0.30"
    defaultE1="40.0e6"
    defaultE2="50.0e6"
    defaultE3="60.0e6"
    defaultG12="20.0e6"
    defaultG13="30.0e6"
    defaultG23="40.0e6" />
  </Constitutive>
```

Model: Drucker-Prager

Overview

This model may be used to represent a solid material with plastic response to loading according to the [Drucker-Prager](#) yield criterion below:

$$f(p, q) = q + b p - a = 0.$$




Fig. 5.2: Mohr-Coulomb and Drucker-Prager yield surfaces in principal stress axes (Borja, 2002).

The material behavior is linear elastic (see [Model: Elastic Isotropic](#)) for $f < 0$, and plastic for $f = 0$. The two material parameters a and b are derived by approximating the Mohr-Coulomb surface with a cone. Figure 3 shows the Mohr-Coulomb yield surface and circumscribing Drucker-Prager surface in principal stress space. The Drucker-Prager yield surface has a circular cross-section in the deviatoric plane that passes through the tension or compression corners of the Mohr-Coulomb yield surface, as shown in the Figure 4. The material parameters a and b are derived as:

$$a = \frac{6 c \cos \phi}{3 \pm \sin \phi}, \quad b = \frac{6 \sin \phi}{3 \pm \sin \phi}$$

where plus signs are for circles passing through the tension corners, and minus signs are for circles passing through compression corners. Also, ϕ and c denote friction angle and cohesion, respectively, as defined by the Mohr-Coulomb failure envelope shown in Figure 5. In GEOSX, we use a compression corner fit (minus signs) to convert the user-specified friction angle and cohesion to a and b .

We consider a non-associative plastic potential to determine the direction of plastic flow.

$$g(p, q) = q + b' p,$$

where $b' \leq b$ is the dilatancy parameter. Setting $b' = b$ leads to associative flow rule, while for $b' < b$ non-associative flow is obtained. The parameter b' is related to [dilation angle](#) as:

$$b' = \frac{6 \sin \psi}{3 \pm \sin \psi},$$




Fig. 5.3: Mohr-Coulomb and Drucker-Prager yield surfaces on the deviatoric plane (Borja, 2013).




Fig. 5.4: The Mohr-Coulomb failure envelope (Borja, 2013).

where $\psi \leq \phi$ is the dilation angle. If $\psi > 0$, then the plastic flow is dilative. Again, we use a compression corner fit (minus sign).

A hardening rule is defined which determines how the yield surface will change as a result of plastic deformations. Here we use linear hardening for the cohesion parameter, a ,

$$\dot{a} = h \dot{\lambda},$$

where h is the hardening parameter. A positive hardening parameter will allow the cohesion to grow, shifting the cohesion intercept vertically on the q-axis. A negative hardening parameter will cause the cohesion to shrink, though negative cohesion values are not allowed. Once all cohesion has been lost, the cohesion will remain at zero, so the cone vertex is fixed at the origin. In either case, the friction and dilation angles remain constant. See the `_Extended-DruckerPrager` model for an alternative version of hardening behavior.

Parameters

The following attributes are supported:

Name	Type	Default	Description
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultCohesion	real64	0	Initial cohesion
defaultDensity	real64	required	Default Material Density
defaultDilationAngle	real64	30	Dilation angle (degrees)
defaultFrictionAngle	real64	30	Friction angle (degrees)
defaultHardeningRate	real64	0	Cohesion hardening/softening rate
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
name	string	required	A name is required for any non-unique nodes

Example

```
<Constitutive>
  <DruckerPrager name="drucker"
    defaultDensity="2700"
    defaultBulkModulus="1000.0"
    defaultShearModulus="1000.0"
    defaultFrictionAngle="30.0"
    defaultDilationAngle="20.0"
    defaultHardeningRate="0.0"
    defaultCohesion="10.0" />
</Constitutive>
```

Variant: J2 plasticity

J2 yield criterion can be obtained as a special case of the Drucker-Prager model by setting the friction and dilation angles to zero, i.e. $\phi = \psi = 0$.

Model: Extended Drucker-Prager

Overview

This model implements a more sophisticated version of the Drucker-Prager model (see [Model: Drucker-Prager](#)) allowing for both cohesion and friction hardening / softening. We implement the specific hardening model reported in [Liu et al. \(2020\)](#). The yield surface is given by

$$f(p, q) = q + b \left(p - \frac{a_i}{b_i} \right) = 0,$$

where b is the current yield surface slope, b_i is the initial slope, and a_i is the initial cohesion intercept in p - q space. The vertex of the Drucker-Prager cone is fixed at $p = a_i/b_i$. Let λ denote the accumulated plastic strain measure. The current yield surface slope is given by the hyperbolic relationship

$$b = b_i + \frac{\lambda}{m + \lambda} (b_r - b_i)$$

with m a parameter controlling the hardening rate. Here, b_r is the residual yield surface slope. If $b_r < b_i$, hardening behavior will be observed, while for $b_r > b_i$ softening behavior will occur.

In the resulting model, the yield surface begins at an initial position defined by the initial cohesion and friction angle. As plastic deformation occurs, the friction angle hardens (or softens) so that it asymptotically approaches a residual friction angle. The vertex of the cone remains fixed in p - q space, but the cohesion intercept evolves in tandem with the friction angle. See [Liu et al. \(2020\)](#) <<https://doi.org/10.1007/s00603-019-01992-5>> for complete details.

In order to allow for non-associative behavior, we define a “dilation ratio” parameter $\theta \in [0, 1]$ such that $b' = \theta b$, where b' is the slope of the plastic potential surface, while b is the slope of the yield surface. Choosing $\theta = 1$ leads to associative behavior, while $\theta = 0$ implies zero dilatancy.

Parameters

The supported attributes will be documented soon.

Example

```
<Constitutive>
  <ExtendedDruckerPrager
    name="edp"
    defaultDensity="2700"
    defaultBulkModulus="500"
    defaultShearModulus="300"
    defaultCohesion="0.0"
    defaultInitialFrictionAngle="15.0"
    defaultResidualFrictionAngle="23.0"
    defaultDilationRatio="1.0"
    defaultHardening="0.001"
  />
</Constitutive>
```

Model: Modified Cam-Clay

This model may be used to represent a solid material with plastic response to loading according to the [Modified Cam-Clay \(MCC\)](#) critical state model. The MCC yield function is defined as:

$$f = q^2 + M^2 [p(p - p_c)] = 0,$$

where p_c is the preconsolidation pressure, and M is the slope of the critical state line (CSL). M can be related to the critical state friction angle ϕ_{cs} as

$$M = \frac{6 \sin \phi_{cs}}{3 - \sin \phi_{cs}}.$$

Here f represents the yield surface, as shown in Figure 6.




Fig. 5.5: Cam-Clay and Modified Cam-Clay yield surfaces in p-q space (Borja, 2013).

Here we use a hyper-elastic constitutive law using the following elastic rate constitutive equation

$$\dot{p} = -\frac{p}{c_r} \dot{\epsilon}_v^e,$$

where $c_r > 0$ is the elastic compressibility index. The tangential elastic bulk modulus is $K = -\frac{p}{c_r}$ and varies linearly with pressure. We assume a constant shear modulus, and can write stress invariants p and q as

$$p = p_0 \exp\left(\frac{\epsilon_{v0} - \epsilon_v^e}{c_r}\right), \quad q = 3\mu\epsilon_s^e,$$

where p_0 is the reference pressure and ϵ_{v0} is the reference volumetric strain. The hardening law is derived from the linear relationship between logarithm of specific volume and logarithm of preconsolidation pressure, as show in Figure 7.

The hardening law describes evolution of the preconsolidation pressure p_c as

$$\dot{p}_c = -\frac{tr(\dot{\epsilon}^p)}{c_c - c_r} p_c,$$

where c_c is the virgin compressibility index and we have $0 < c_r < c_c$.

Parameters

The supported attributes will be documented soon.

Example




Fig. 5.6: Bilogarithmic hardening law derived from isotropic compression tests (Borja, 2013).

```
<Constitutive>
  <ModifiedCamClay name="mcc"
    defaultDensity="2700"
    defaultRefPressure="-1.0"
    defaultRefStrainVol="0"
    defaultShearModulus="200.0"
    defaultPreConsolidationPressure="-1.5"
    defaultCslSlope="1.2"
    defaultRecompressionIndex="0.002"
    defaultVirginCompressionIndex="0.003" />
</Constitutive>
```

Model: Delft Egg

The **Delft-Egg** plasticity model uses a generalization of the Modified Cam-Clay yield surface, defined as

$$f = q^2 - M^2 \left[\alpha^2 p \left(\frac{2\alpha}{\alpha+1} p_c - p \right) - \frac{\alpha^2(\alpha-1)}{\alpha+1} p_c^2 \right] = 0 \quad \left(\text{for } p_c > \frac{\alpha}{\alpha+1} \right)$$

$$f = q^2 - M^2 p \left(\frac{2\alpha}{\alpha+1} p_c - p \right) = 0 \quad \left(\text{for } p_c \leq \frac{\alpha}{\alpha+1} \right)$$

where $\alpha \geq 1$ is the shape parameter. For $\alpha = 1$, this model leads to a Modified Cam-Clay (MCC) type model with an ellipsoidal yield surface. For $\alpha > 1$, an egg-shaped yield surface is obtained. The additional parameter makes it easier to fit the cap behavior of a broader range of soils and rocks.

Because Delft-Egg is frequently used for hard rocks, GEOSX uses a linear model for the elastic response, rather than the hyper-elastic model used for MCC. This is a slight deviation from the original formulation proposed in the reference above. For reservoir applications, the ability to use a simpler linear model was a frequent user request.

Parameters

The supported attributes will be documented soon.

Example

```
<Constitutive>
  <DelftEgg
    name="DE"
    defaultDensity="2700"
    defaultBulkModulus="10.0e9"
    defaultShearModulus="6.0e9"
    defaultPreConsolidationPressure="-20.0e6"
    defaultShapeParameter="6.5"
    defaultCslSlope="1.2"
    defaultVirginCompressionIndex="0.005"
    defaultRecompressionIndex="0.001"/>
</Constitutive>
```

Damage Models

The damage models are in active development, and documentation will be added when they are ready for production release.

5.4.2 Fluid Models

These models provide density, viscosity, and composition relationships for single fluids and fluid mixtures.

Compressible single phase fluid model

Overview

This model represents a compressible single-phase fluid with constant compressibility and pressure-dependent viscosity. These assumptions are valid for slightly compressible fluids, such as water, and some types of oil with negligible amounts of dissolved gas.

Specifically, fluid density is computed as

$$\rho(p) = \rho_0 e^{c_\rho(p-p_0)}$$

where c_ρ is compressibility, p_0 is reference pressure, ρ_0 is density at reference pressure. Similarly,

$$\mu(p) = \mu_0 e^{c_\mu(p-p_0)}$$

where c_μ is viscosibility (viscosity compressibility), μ_0 is reference viscosity.

Either exponent may be approximated by linear (default) or quadratic terms of Taylor series expansion. Currently there is no temperature dependence in the model, although it may be added in future.

Parameters

The model is represented by `<CompressibleSinglePhaseFluid>` node in the input.

The following attributes are supported:

Name	Type	Default	Description
compressibility	real64	0	Fluid compressibility
defaultDensity	real64	required	Default value for density.
defaultViscosity	real64	required	Default value for viscosity.
densityModelType	geosx_constitutive_ExponentialApproximationType		Type of density model. Valid options: * exponential * linear * quadratic
name	string	required	A name is required for any non-unique nodes
referenceDensity	real64	1000	Reference fluid density
referencePressure	real64	0	Reference pressure
referenceViscosity	real64	0.001	Reference fluid viscosity
viscosibility	real64	0	Fluid viscosity exponential coefficient
viscosityModelType	geosx_constitutive_ExponentialApproximationType		Type of viscosity model. Valid options: * exponential * linear * quadratic

Example

```
<Constitutive>
  <CompressibleSinglePhaseFluid name="water"
    referencePressure="2.125e6"
    referenceDensity="1000"
    compressibility="1e-19"
    referenceViscosity="0.001"
    viscosibility="0.0"/>
</Constitutive>
```

Black-oil fluid model

Overview

In the black-oil model three pseudo-components, oil (o), gas (g) and water (w) are considered. These are assumed to be partitioned across three fluid phases, named liquid (l), vapor (v) and aqueous (a).

Phase behavior is characterized by the following quantities which are used to relate properties of the fluids in the reservoir to their properties at surface conditions.

- B_o : oil formation volume factor
- B_g : gas formation volume factor
- R_s : gas/oil ratio
- R_v : oil/gas ratio

By tables, that tabulate saturated and undersaturated oil and gas properties as functions of pressure and solution ratios.

Dead oil

In **dead-oil** each component occupies only one phase. Thus, the following partition matrix determines the components distribution within the three phases:

$$\begin{bmatrix} y_{gv} & y_{gl} & y_{ga} \\ y_{ov} & y_{ol} & y_{oa} \\ y_{wv} & y_{wl} & y_{wa} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and the phase densities are

$$\rho_l = \frac{\rho_o^{STC}}{B_o}$$

$$\rho_v = \frac{\rho_g^{STC}}{B_g}.$$

Live oil

The live oil fluid model make no assumptions about the partitioning of the hydrocarbon components and the following composition matrix can be used

$$\begin{bmatrix} y_{gv} & y_{gl} & y_{ga} \\ y_{ov} & y_{ol} & y_{oa} \\ y_{wv} & y_{wl} & y_{wa} \end{bmatrix} = \begin{bmatrix} \frac{\rho_g^{STC}}{\rho_g^{STC} + \rho_o^{STC} R_s} & \frac{\rho_g^{STC} R_s}{\rho_o^{STC} + \rho_g^{STC} R_s} & 0 \\ \frac{\rho_o^{STC} R_s}{\rho_g^{STC} + \rho_o^{STC} R_s} & \frac{\rho_o^{STC}}{\rho_o^{STC} + \rho_g^{STC} R_s} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

whereas the densities of the two hydrocarbon phases are

$$\rho_l = \frac{\rho_o^{STC} + \rho_g^{STC} R_s}{B_o}$$

$$\rho_v = \frac{\rho_g^{STC} + \rho_o^{STC} R_v}{B_g}$$

See [Petrowiki](#) for more information.

Parameters

Both types are represented by `<BlackOilFluid>` node in the input. Under the hood this is a wrapper around `PVTPackage` library, which is included as a submodule. In order to use the model, GEOSX must be built with `-DENABLE_PVTPACKAGE=ON` (default).

The following attributes are supported:

Name	Type	Default	Description
componentMolarWeight	real64_array	required	Component molar weights
componentNames	string_array	{}	List of component names
hydrocarbonFormationVolumeFactorTableNames	string_array	{}	<p>List of formation volume factor TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
hydrocarbonViscosityTableNames	string_array	{}	<p>List of viscosity TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases
surfaceDensities	real64_array	required	List of surface mass densities for each phase
tableFiles	path_array	{}	List of filenames with input PVT tables (one per phase)
waterCompressibility	real64	0	Water compressibility
waterFormationVolumeFactor	real64	0	Water formation volume factor
waterReferencePressure	real64	0	Water reference pressure
waterViscosity	real64	0	Water viscosity

Supported phase names are:

Value	Comment
oil	Oil phase
gas	Gas phase
water	Water phase

Example

```
<Constitutive>
  <BlackOilFluid name="fluid1"
    fluidType="LiveOil"
    phaseNames="{ oil, gas, water }"
    surfaceDensities="{ 800.0, 0.9907, 1022.0 }"
    componentMolarWeight="{ 114e-3, 16e-3, 18e-3 }"
    tableFiles="{ pvt_o.txt, pvt_g.txt, pvt_w.txt }"/>
</Constitutive>
```

Compositional multiphase fluid model

Overview

This model represents a full composition description of a multiphase multicomponent fluid. Phase behavior is modeled by an Equation of State (EOS) and partitioning of components into phases is computed based on instantaneous chemical equilibrium via a two- or three-phase flash. Each component (species) is characterized by molar weight and critical properties that serve as input parameters for the EOS. See [Petrowiki](#) for more information.

Parameters

The model represented by `<CompositionalMultiphaseFluid>` node in the input. Under the hood this is a wrapper around `PVTPackage` library, which is included as a submodule. In order to use the model, GEOSX must be built with `-DENABLE_PVTPACKAGE=ON` (default).

The following attributes are supported:

Name	Type	Default	Description
componentAcentricFactor	real64_array	required	Component acentric factors
componentBinaryCoeff	real64_array2d	{{0}}	Table of binary interaction coefficients
componentCriticalPressure	real64_array	required	Component critical pressures
componentCriticalTemperature	real64_array	required	Component critical temperatures
componentMolarWeight	real64_array	required	Component molar weights
componentNames	string_array	required	List of component names
componentVolumeShift	real64_array	{0}	Component volume shifts
equationsOfState	string_array	required	List of equation of state types for each phase
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases

Supported phase names are:

Value	Comment
oil	Oil phase
gas	Gas phase
water	Water phase

Supported Equation of State types:

Value	Comment
PR	Peng-Robinson EOS
SRK	Soave-Redlich-Kwong EOS

Example

```
<Constitutive>
  <CompositionalMultiphaseFluid name="fluid1"
    phaseNames="{ oil, gas }"
    equationsOfState="{ PR, PR }"
    componentNames="{ N2, C10, C20, H2O }"
    componentCriticalPressure="{ 34e5, 25.3e5, 14.6e5, ↵
↵220.5e5 }"
    componentCriticalTemperature="{ 126.2, 622.0, 782.0, ↵
↵647.0 }"
    componentAcentricFactor="{ 0.04, 0.443, 0.816, 0.344 }"
    ↵"
    componentMolarWeight="{ 28e-3, 134e-3, 275e-3, 18e-3 }"
    ↵"
    componentVolumeShift="{ 0, 0, 0, 0 }"
    componentBinaryCoeff="{ { 0, 0, 0, 0 },
                             { 0, 0, 0, 0 },
                             { 0, 0, 0, 0 },
                             { 0, 0, 0, 0 } }"/>
</Constitutive>
```

CO2-brine model

Summary

The CO2-brine model implemented in GEOSX includes two components (CO2 and H2O) that are transported by one or two fluid phases (the brine phase and the CO2 phase). We refer to the brine phase with the subscript ℓ and to the CO2 phase with the subscript g (although the CO2 phase can be in supercritical, liquid, or gas state). The water component is only present in the brine phase, while the CO2 component can be present in the CO2 phase as well as in the brine phase. Thus, considering the phase component fractions, $y_{c,p}$ (i.e., the fraction of the mass of phase p represented by component c) the following partition matrix determines the component distribution within the two phases:

$$\begin{bmatrix} y_{CO_2,g} & y_{CO_2,\ell} \\ 0 & 1 \end{bmatrix}$$

The update of the fluid properties is done in two steps:

- 1) The phase fractions (ν_p) and phase component fractions ($y_{c,p}$) are computed as a function of pressure (p), temperature (T), component fractions (z_c), and a constant salinity.

- 2) The phase densities (ρ_p) and phase viscosities (μ_p) are computed as a function of pressure, temperature, the updated phase component fractions, and a constant salinity.

Once the phase fractions, phase component fractions, phase densities, phase viscosities—and their derivatives with respect to pressure, temperature, and component fractions—have been computed, the *Compositional Multiphase Flow Solver* proceeds to the assembly of the accumulation and flux terms. Note that the current implementation of the flow solver is isothermal and that the derivatives with respect to temperature are therefore discarded.

The models that are used in steps 1) and 2) are reviewed in more details below.

Step 1: Computation of the phase fractions and phase component fractions (flash)

At initialization, GEOSX performs a preprocessing step to construct a two-dimensional table storing the values of CO2 solubility in brine as a function of pressure, temperature, and a constant salinity. The user can parameterize the construction of the table by specifying the salinity and by defining the pressure (p) and temperature (T) axis of the table in the form:

FlashModel	CO2Solubility	p_{min}	p_{max}	Δp	T_{min}	T_{max}	ΔT	Salinity
------------	---------------	-----------	-----------	------------	-----------	-----------	------------	----------

Note that the pressures are in Pascal, temperatures are in Kelvin, and the salinity is a molality (moles of NaCl per kg of brine). The temperature must be between 283.15 and 623.15 Kelvin. The table is populated using the model of Duan and Sun (2003). Specifically, we solve the following nonlinear CO2 equation of state (equation (A1) in Duan and Sun, 2003) for each pair (p, T) to obtain the reduced volume, V_r .

$$\frac{p_r V_r}{T_r} = 1 + \frac{a_1 + a_2/T_r^2 + a_3/T_r^3}{V_r} + \frac{a_4 + a_5/T_r^2 + a_6/T_r^3}{V_r^2} + \frac{a_7 + a_8/T_r^2 + a_9/T_r^3}{V_r^4} + \frac{a_{10} + a_{11}/T_r^2 + a_{12}/T_r^3}{V_r^5} + \frac{a_{13}}{T_r^3 V_r^2} \left(a_{14} + \frac{a_{15}}{V_r^2} \right) \exp\left(-\frac{a_{15}}{V_r^2}\right)$$

where $p_r = p/p_{crit}$ and $T_r = T/T_{crit}$ are respectively the reduced pressure and the reduced temperature. We refer the reader to Table (A1) in Duan and Sun (2003) for the definition of the coefficients a_i involved in the previous equation. Using the reduced volume, V_r , we compute the fugacity coefficient of CO2, $\ln_\phi(p, T)$, using equation (A6) of Duan and Sun (2003). To conclude this preprocessing step, we use the fugacity coefficient of CO2 to compute and store the solubility of CO2 in brine, s_{CO2} , using equation (6) of Duan and Sun (2003):

$$\ln \frac{x_{CO2} P}{s_{CO2}} = \frac{\Phi_{CO2}}{RT} - \ln_\phi(p, T) + \sum_c 2\lambda_c m + \sum_a 2\lambda_a m + \sum_{a,c} \zeta_{a,c} m^2$$

where Φ_{CO2} is the chemical potential of the CO2 component, R is the gas constant, and m is the salinity. The mole fraction of CO2 in the vapor phase, x_{CO2} , is computed with equation (4) of Duan and Sun (2003). Note that the first, third, fourth, and fifth terms in the equation written above are approximated using equation (7) of Duan and Sun (2003) as recommended by the authors.

During the simulation, Step 1 starts with a look-up in the precomputed table to get the CO2 solubility, s_{CO2} , as a function of pressure and temperature. Then, we compute the phase fractions as:

$$\nu_\ell = \frac{1 + s_{CO2}}{1 + z_{CO2}/(1 - z_{CO2})}$$

$$\nu_g = 1 - \nu_\ell$$

We conclude Step 1 by computing the phase component fractions as:

$$y_{CO2,\ell} = \frac{s_{CO2}}{1 + s_{CO2}}$$

$$y_{H2O,\ell} = 1 - y_{CO2,\ell}$$

$$y_{CO2,g} = 1$$

$$y_{H2O,g} = 0$$

Step 2: Computation of the phase densities and phase viscosities

CO2 phase density and viscosity

In GEOSX, the computation of the CO2 phase density and viscosity is entirely based on look-up in precomputed tables. The user defines the pressure (in Pascal) and temperature (in Kelvin) axis of the density table in the form:

DensityFun	SpanWagnerCO2Density	p_{min}	p_{max}	Δp	T_{min}	T_{max}	ΔT
------------	----------------------	-----------	-----------	------------	-----------	-----------	------------

This correlation is valid for pressures less than 8×10^8 Pascal and temperatures less than 1073.15 Kelvin. Using these parameters, GEOSX internally constructs a two-dimensional table storing the values of density as a function of pressure and temperature. This table is populated as explained in the work of Span and Wagner (1996) by solving the following nonlinear Helmholtz energy equation for each pair (p, T) to obtain the value of density, ρ_g :

$$\frac{p}{RT\rho_g} = 1 + \delta\phi_\delta^r(\delta, \tau)$$

where R is the gas constant, $\delta := \rho_g/\rho_{crit}$ is the reduced CO2 phase density, and $\tau := T_{crit}/T$ is the inverse of the reduced temperature. The definition of the residual part of the energy equation, denoted by ϕ_δ^r , can be found in equation (6.5), page 1544 of Span and Wagner (1996). The coefficients involved in the computation of ϕ_δ^r are listed in Table (31), page 1544 of Span and Wagner (1996). These calculations are done in a preprocessing step.

The pressure and temperature axis of the viscosity table can be parameterized in a similar fashion using the format:

ViscosityFun	FenghourCO2Viscosity	p_{min}	p_{max}	Δp	T_{min}	T_{max}	ΔT
--------------	----------------------	-----------	-----------	------------	-----------	-----------	------------

This correlation is valid for pressures less than 3×10^8 Pascal and temperatures less than 1493.15 Kelvin. This table is populated as explained in the work of Fenghour and Wakeham (1998) by computing the CO2 phase viscosity, μ_g , as follows:

$$\mu_g = \mu_0(T) + \mu_{excess}(\rho_g, T) + \mu_{crit}(\rho_g, T)$$

The “zero-density limit” viscosity, $\mu_0(T)$, is computed as a function of temperature using equations (3), (4), and (5), as well as Table (1) of Fenghour and Wakeham (1998). The excess viscosity, $\mu_{excess}(\rho_g, T)$, is computed as a function of temperature and CO2 phase density (computed as explained above) using equation (8) and Table (3) of Fenghour and Wakeham (1998). We currently neglect the critical viscosity, μ_{crit} . These calculations are done in a preprocessing step.

During the simulation, the update of CO2 phase density and viscosity is simply done with a look-up in the precomputed tables.

Brine density and viscosity

The computation of the brine density involves a tabulated correlation presented in Phillips et al. (1981). The user specifies the (constant) salinity and defines the pressure and temperature axis of the brine density table in the form:

DensityFun	BrineCO2Density	p_{min}	p_{max}	Δp	T_{min}	T_{max}	ΔT	Salinity
------------	-----------------	-----------	-----------	------------	-----------	-----------	------------	----------

The pressure must be in Pascal and must be less than 5×10^7 Pascal. The temperature must be in Kelvin and must be between 283.15 and 623.15 Kelvin. The salinity is a molality (moles of NaCl per kg of brine). Using these parameters,

GEOSX performs a preprocessing step to construct a two-dimensional table storing the brine density, $\rho_{\ell,table}$ for the specified salinity as a function of pressure and temperature using the expression:

$$\rho_{\ell,table} = A + Bx + Cx^2 + Dx^3$$

$$x = c_1 \exp(a_1 m) + c_2 \exp(a_2 T) + c_3 \exp(a_3 P)$$

We refer the reader to Phillips et al. (1981), equations (4) and (5), pages 14 and 15 for the definition of the coefficients involved in the previous equation. This concludes the preprocessing step.

Then, during the simulation, the brine density update proceeds in two steps. First, a table look-up is performed to retrieve the value of density, $\rho_{\ell,table}$. Then, in a second step, the density is modified using the method of Garcia (2001) to account for the presence of CO2 dissolved in brine as follows:

$$\rho_{\ell} = \rho_{\ell,table} + M_{CO2} c_{CO2} - c_{CO2} \rho_{\ell,table} V_{\phi}$$

where M_{CO2} is the molecular weight of CO2, c_{CO2} is the concentration of CO2 in brine, and V_{ϕ} is the apparent molar volume of dissolved CO2. The CO2 concentration in brine is obtained as:

$$c_{CO2} = \frac{y_{CO2,\ell} \rho_{\ell,table}}{M_{H2O}(1 - y_{CO2,\ell})}$$

where M_{H2O} is the molecular weight of water. The apparent molar volume of dissolved CO2 is computed as a function of temperature using the expression:

$$V_{\phi} = 37.51 - 9.585 \times 10^{-2} T + 8.740 \times 10^{-4} T^2 - 5.044 \times 10^{-7} T^3$$

The brine viscosity is controlled by a salinity parameter provided by the user in the form:

ViscosityFun	BrineViscosity	Salinity
--------------	----------------	----------

During the simulation, the brine viscosity is updated as a function of temperature using the analytical relationship of Phillips et al. (1981):

$$\mu_{\ell} = aT + b$$

where the coefficients a and b are defined as:

$$a = 0.00089 \times 0.000629(1.0 - \exp(-0.7m))$$

$$b = 0.00089(1.0 + 0.0816m + 0.0122m^2 + 0.000128m^3)$$

where m is the user-defined salinity (in moles of NaCl per kg of brine).

Parameters

The model is represented by <CO2BrineFluid> node in the input.

The following attributes are supported:

Name	Type	De- fault	Description
componentMolar- Weight	real64_array	{0}	Component molar weights
componentNames	string_array	{ }	List of component names
flashModelParaFile	path	re- quired	Name of the file defining the parameters of the flash model
name	string	re- quired	A name is required for any non-unique nodes
phaseNames	string_array	{ }	List of fluid phases
phasePVTParaFiles	path_array	re- quired	Names of the files defining the parameters of the viscosity and density models

Supported phase names are:

Value	Comment
gas	CO2 phase
water	Water phase

Supported component names are:

Value	Component
co2,CO2	CO2 component
water,liquid	Water component

Example

```
<Constitutive>
  <CO2BrineFluid
    name="fluid"
    phaseNames="{ gas, water }"
    componentNames="{ co2, water }"
    componentMolarWeight="{ 44e-3, 18e-3 }"
    phasePVTParaFiles="{ pvtgas.txt, pvtliquid.txt }"
    flashModelParaFile="co2flash.txt"/>
</Constitutive>
```

In the XML code listed above, “co2flash.txt” parameterizes the CO2 solubility table constructed in Step 1. The file “pvtgas.txt” parameterizes the CO2 phase density and viscosity tables constructed in Step 2, while the file “pvtliquid.txt” parameterizes the brine density and viscosity tables.

References

- Z. Duan and R. Sun, An improved model calculating CO2 solubility in pure water and aqueous NaCl solutions from 273 to 533 K and from 0 to 2000 bar., Chemical Geology, vol. 193.3-4, pp. 257-271, 2003.
- R. Span and W. Wagner, A new equation of state for carbon dioxide covering the fluid region from the triple-point temperature to 1100 K at pressure up to 800 MPa, J. Phys. Chem. Ref. Data, vol. 25, pp. 1509-1596, 1996.
- A. Fenghour and W. A. Wakeham, The viscosity of carbon dioxide, J. Phys. Chem. Ref. Data, vol. 27, pp. 31-44, 1998.
- S. L. Phillips et al., A technical databook for geothermal energy utilization, Lawrence Berkeley Laboratory report, 1981.
- J. E. Garcia, Density of aqueous solutions of CO2. No. LBNL-49023. Lawrence Berkeley National Laboratory, Berkeley, CA, 2001.

5.4.3 Relative Permeability Models

Brooks-Corey relative permeability model

Overview

The following paragraphs explain how the Brooks-Corey model is used to compute the phase relative permeabilities as a function of volume fraction with the expression:

$$k_{r\ell} = k_{r\ell,max} S_{\ell,scaled}^{\lambda_{\ell}},$$

where the scaled volume fraction of phase ℓ is computed as:

$$S_{\ell,scaled} = \frac{S_{\ell} - S_{\ell,min}}{1 - S_{w,min} - S_{o,min} - S_{g,min}}.$$

The minimum phase volume fractions $S_{\ell,min}$ are model parameters specified by the user.

Parameters

The relative permeability constitutive model is listed in <Constitutive> block of the input XML file. The relative permeability model must be assigned a unique name via name attribute. This name is used to assign the model to regions of the physical domain via a materialList attribute of the <ElementRegion> node.

The following attributes are supported:

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases
phaseRelPermExponent	real64_array	{1}	MinimumRel perm power law exponent for each phase
phaseRelPermMaxValue	real64_array	{0}	Maximum rel perm value for each phase

Below are some comments on the model parameters.

- `phaseNames` - The number of phases can be either 2 or 3. The capillary pressure model assumes that oil is always present. Supported phase names are:

Value	Phase
oil	Oil phase
gas	Gas phase
water	Water phase

- `phaseMinVolFraction` - The list of minimum volume fractions $S_{\ell,min}$ for each phase is specified in the same order as in `phaseNames`. Below this volume fraction, the phase is assumed to be immobile.
- `phaseRelPermExponent` - The list of exponents λ_{ℓ} for each phase is specified in the same order as in `phaseNames`.
- `phaseMaxValue` - The list of maximum values $k_{r\ell,max}$ for each phase is specified in the same order as in `phaseNames`.

Example

```

<Constitutive>
...
  <BrooksCoreyRelativePermeability name="relPerm"
    phaseNames="{oil, water}"
    phaseMinVolumeFraction="{0.02, 0.015}"
    phaseRelPermExponent="{2, 2.5}"
    phaseRelPermMaxValue="{0.8, 1.0}"/>
...
</Constitutive>

```

Three-phase relative permeability model

Overview

For the simulation of three-phase flow in porous media, it is common to use a specific treatment (i.e., different from the typical two-phase procedure) to evaluate the oil relative permeability. Specifically, the three-phase oil relative permeability is obtained by interpolation of oil-water and oil-gas experimental data measured independently in two-phase displacements.

Let $k_{rw,wo}$ and $k_{ro,wo}$ be the water-oil two-phase relative permeabilities for the water phase and the oil phase, respectively. Let $k_{rg,go}$ and $k_{ro,go}$ be the oil-gas two-phase relative permeabilities for the gas phase and the oil phase, respectively. In the current implementation, the two-phase relative permeability data is computed analytically using the *Brooks-Corey relative permeability model*.

The water and gas three-phase relative permeabilities are simply given by two-phase data and only depend on S_w and S_g , respectively. That is,

$$k_{rw,wog}(S_w) = k_{rw,wo}(S_w),$$

$$k_{rg,wog}(S_g) = k_{rg,go}(S_g).$$

The oil three-phase relative permeability is obtained using a variant of the saturation-weighted interpolation procedure initially proposed by [Baker](#). Specifically, we compute:

$$k_{ro,wog}(S_w, S_g) = \frac{(S_w - S_{w,min})k_{ro,wo}(S_w) + S_g k_{rg,go}(S_g)}{(S_w - S_{w,min}) + S_g}.$$

This procedure provides a simple but effective formula avoiding the problems associated with the other interpolation methods (negative values).

Parameters

The relative permeability constitutive model is listed in `<Constitutive>` block of the input XML file. The relative permeability model must be assigned a unique name via `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute of the `<ElementRegion>` node.

The following attributes are supported:

Name	Type	Default	Description
gasOilRelPermExponent	real64_array	{1}	Rel perm power law exponent for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasExp, oilExp }”, in that order
gasOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasMax, oilMax }”, in that order
name	string	required	A name is required for any non-unique nodes
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases
waterOilRelPermExponent	real64_array	{1}	Rel perm power law exponent for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterExp, oilExp }”, in that order
waterOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterMax, oilMax }”, in that order

Below are some comments on the model parameters.

- `phaseNames` - The number of phases should be 3. Supported phase names are:

Value	Phase
oil	Oil phase
gas	Gas phase
water	Water phase

- `phaseMinVolFraction` - The list of minimum volume fractions $S_{\ell,min}$ for each phase is specified in the same order as in `phaseNames`. Below this volume fraction, the phase is assumed to be immobile.
- `waterOilRelPermExponent` - The list of exponents $\lambda_{\ell,wo}$ for the two-phase water-oil relative permeability data, with the water exponent first and the oil exponent next. These exponents are then used to compute $k_{r\ell,wo}$ in the *Brooks-Corey relative permeability model*.
- `waterOilRelPermMaxValue` - The list of maximum values $k_{r\ell,wo,max}$ for the two-phase water-oil relative permeability data, with the water max value first and the oil max value next. These exponents are then used to compute $k_{r\ell,wo}$ in the *Brooks-Corey relative permeability model*.
- `gasOilRelPermExponent` - The list of exponents $\lambda_{\ell,go}$ for the two-phase gas-oil relative permeability data, with the gas exponent first and the oil exponent next. These exponents are then used to compute $k_{r\ell,go}$ in the *Brooks-Corey relative permeability model*.
- `gasOilRelPermMaxValue` - The list of maximum values $k_{r\ell,go,max}$ for the two-phase gas-oil relative permeability data, with the gas max value first and the oil max value next. These exponents are then used to compute $k_{r\ell,go}$ in the *Brooks-Corey relative permeability model*.

Example

```
<Constitutive>
...
<BrooksCoreyBakerRelativePermeability name="relperm"
    phaseNames="{oil, gas, water}"
    phaseMinVolumeFraction="{0.05, 0.05, 0.05}"
    waterOilRelPermExponent="{2.5, 1.5}"
    waterOilRelPermMaxValue="{0.8, 0.9}"
    gasOilRelPermExponent="{3, 3}"
    gasOilRelPermMaxValue="{0.4, 0.9}"/>
...
</Constitutive>
```

5.4.4 Capillary Pressure Models

Brooks-Corey capillary pressure model

Overview

In GEOSX, the oil-phase pressure is assumed to be the primary pressure. The following paragraphs explain how the Brooks-Corey capillary pressure model is used to compute the water-phase and gas-phase pressures as:

$$p_w = p_o - P_{c,w}(S_w),$$

and

$$p_g = p_o + P_{c,g}(S_g).$$

In the Brooks-Corey model, the water-phase capillary pressure is computed as a function of the water-phase volume fraction with the following expression:

$$P_{c,w}(S_w) = p_{e,w} S_{w,scaled}^{-1/\lambda_w},$$

where the scaled water-phase volume fraction is computed as:

$$S_{w,scaled} = \frac{S_w - S_{w,min}}{1 - S_{w,min} - S_{o,min} - S_{g,min}}.$$

The gas capillary pressure is computed analogously.

Parameters

The capillary pressure constitutive model is listed in `<Constitutive>` block of the input XML file. The capillary pressure model must be assigned a unique name via `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute of the `<ElementRegion>` node.

The following attributes are supported:

Name	Type	De- fault	Description
capPressureEp- silon	real64	1e- 06	Wetting-phase saturation at which the max cap. pressure is attained; used to avoid infinite cap. pressure values for saturations close to zero
name	string	re- quired	A name is required for any non-unique nodes
phaseCapPres- sureExponentInv	real64_array {3}		Inverse of capillary power law exponent for each phase
phaseEntryPres- sure	real64_array {4}		Entry pressure value for each phase
phaseMinVol- umeFraction	real64_array {4}		Minimum volume fraction value for each phase
phaseNames	string_array {4}	re- quired	List of fluid phases

Below are some comments on the model parameters:

- `phaseNames` - The number of phases can be either 2 or 3. The names entered for this attribute should match the phase names specified in the relative permeability block, either in *Brooks-Corey relative permeability model* or in *Three-phase relative permeability model*. The capillary pressure model assumes that oil is always present. Supported phase names are:

Value	Phase
oil	Oil phase
gas	Gas phase
water	Water phase

- `phaseMinVolFraction` - The list of minimum volume fractions $S_{\ell,min}$ for each phase is specified in the same order as in `phaseNames`. Below this volume fraction, the phase is assumed to be immobile. The values entered for this attribute have to match those of the same attribute in the relative permeability block.
- `phaseCapPressureExponentInv` - The list of exponents λ_ℓ for each phase is specified in the same order as in `phaseNames`. The parameter corresponding to the oil phase is currently not used.
- `phaseEntryPressure` - The list of entry pressures $p_{e,\ell}$ for each phase is specified in the same order as in `phaseNames`. The parameter corresponding to the oil phase is currently not used.

- `capPressureEpsilon` - This parameter is used for both the water-phase and gas-phase capillary pressure. To avoid extremely large, or infinite, capillary pressure values, we set $P_{c,w}(S_w) := P_{c,w}(\epsilon)$ whenever $S_w < \epsilon$. The gas-phase capillary pressure is treated analogously.

Example

```
<Constitutive>
...
<BrooksCoreyCapillaryPressure name="capPressure"
    phaseNames="{oil, gas}"
    phaseMinVolumeFraction="{0.01, 0.015}"
    phaseCapPressureExponentInv="{0, 6}"
    phaseEntryPressure="{0, 1e8}"
    capPressureEpsilon="1e-8"/>
...
</Constitutive>
```

Van Genuchten capillary pressure model

Overview

In GEOSX, the oil-phase pressure is assumed to be the primary pressure. The following paragraphs explain how the Van Genuchten capillary pressure model is used to compute the water-phase and gas-phase pressures as:

$$p_w = p_o - P_{c,w}(S_w),$$

and

$$p_g = p_o + P_{c,g}(S_g),$$

The Van Genuchten model computes the water-phase capillary pressure as a function of the water-phase volume fraction as:

$$P_c(S_w) = \alpha_w (S_{w,scaled}^{-1/m_w} - 1)^{(1-m_w)/2},$$

where the scaled water-phase volume fraction is computed as:

$$S_{w,scaled} = \frac{S_w - S_{w,min}}{1 - S_{w,min} - S_{o,min} - S_{g,min}}.$$

The gas-phase capillary pressure is computed analogously.

Parameters

The capillary pressure constitutive model is listed in `<Constitutive>` block of the input XML file. The capillary pressure model must be assigned a unique name via `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute of the `<ElementRegion>` node.

The following attributes are supported:

Name	Type	Default	Description
capPressureEpsilon	real64	1e-06	Saturation at which the extremum capillary pressure is attained; used to avoid infinite capillary pressure values for saturations close to 0 and 1
name	string	required	A name is required for any non-unique nodes
phaseCapPressureExponentInv	real64_array	{0.5}	Inverse of capillary power law exponent for each phase
phaseCapPressureMultiplier	real64_array	{1}	Entry pressure value for each phase
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases

Below are some comments on the model parameters:

- `phaseNames` - The number of phases can be either 2 or 3. The phase names entered for this attribute should match the phase names specified in the relative permeability block, either in *Brooks-Corey relative permeability model* or in *Three-phase relative permeability model*. The capillary model assumes that oil is always present. Supported phase names are:

Value	Phase
oil	Oil phase
gas	Gas phase
water	Water phase

- `phaseMinVolFraction` - The list of minimum volume fractions $S_{\ell,min}$ for each phase is specified in the same order as in `phaseNames`. Below this volume fraction, the phase is assumed to be immobile. The values entered for this attribute have to match those of the same attribute in the relative permeability block.
- `phaseCapPressureExponentInv` - The list of exponents m_{ℓ} for each phase is specified in the same order as in `phaseNames`. The parameter corresponding to the oil phase is not used.
- `phaseCapPressureMultiplier` - The list of multipliers α_{ℓ} for each phase is specified in the same order as in `phaseNames`. The parameter corresponding to the oil phase is not used.
- `capPressureEpsilon` - The parameter ϵ . This parameter is used for both the water-phase and gas-phase capillary pressure. To avoid extremely large, or infinite, capillary pressure values, we set $P_{c,w}(S_w) := P_{c,w}(\epsilon)$ whenever $S_w < \epsilon$. The gas-phase capillary pressure is treated analogously.

Example

```
<Constitutive>
...
  <VanGenuchtenCapillaryPressure name="capPressure"
    phaseNames="{water, oil}"
    phaseMinVolumeFraction="{0.1, 0.015}"
    phaseCapPressureExponentInv="{0.55, 0}"
    phaseCapPressureMultiplier="1e6 0"
    capPressureEpsilon="1e-7"/>
...
</Constitutive>
```

5.4.5 Porosity models

Pressure dependent porosity

Overview

This model assumes a simple exponential law for the porosity as function of pressure, i.e.

$$\phi = \phi_{ref} \exp(c \cdot (p - p_{ref}))$$

where ϕ_{ref} is the reference porosity at reference pressure, p_{ref} , p is the pressure and, c is the compressibility.

Parameters

The following attributes are supported:

Name	Type	Default	Description
compressibility	real64	required	Solid compressibility
defaultReferencePorosity	real64	required	Default value of the reference porosity
name	string	required	A name is required for any non-unique nodes
referencePressure	real64	required	Reference pressure for solid compressibility

Example

```
<Constitutive>
...
  <PressurePorosity name="rockPorosity"
    referencePressure="1.0e27"
    defaultReferencePorosity="0.3"
    compressibility="1.0e-9"/>
...
</Constitutive>
```

Biot Porosity Model

Overview

According to the poroelasticity theory, the porosity (pore volume), ϕ , can be computed as

$$\phi = \phi_{ref} + \alpha(\epsilon_v - \epsilon_{v,ref}) + (p - p_{ref})/N.$$

Here, ϕ_{ref} is the porosity at a reference state with pressure p_{ref} and volumetric strain $\epsilon_{v,ref}$. Additionally, α is the Biot coefficient, ϵ_v is the volumetric strain, p is the fluid pressure and $N = \frac{K_s}{\alpha - \phi_{ref}}$, where K_s is the grain bulk modulus.

Parameters

The Biot Porosity Model can be called in the <Constitutive> block of the input XML file. This porosity model must be assigned a unique name via the `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute in the <ElementRegions> block.

The following attributes are supported:

Name	Type	Default	Description
defaultReferencePorosity	real64	required	Default value of the reference porosity
grainBulkModulus	real64	required	Grain bulk modulus
name	string	required	A name is required for any non-unique nodes

Example

```
<Constitutive>
...
  <BiotPorosity name="rockPorosity"
    grainBulkModulus="1.0e27"
    defaultReferencePorosity="0.3"/>
...
</Constitutive>
```

5.4.6 Permeability models

Constant Permeability Model

Overview

This model is used to define a diagonal permeability tensor that does not depend on any primary variable.

Parameters

The following attributes are supported:

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityComponents	RTensor	required	xx, yy and zz components of a diagonal permeability tensor.

Example

```
<Constitutive>
...
  <ConstantPermeability name="matrixPerm"
    permeabilityComponents="{1.0e-12, 1.0e-12, 1.0e-12}"/>
...
</Constitutive>
```

Kozeny-Carman Permeability Model

Overview

In the Kozeny-Carman model (see [ref](#)), the permeability of a porous medium is governed by several key parameters, including porosity, grain size, and grain shape:

$$k = \frac{(s_\epsilon D_p)^2 \phi^3}{150(1 - \phi)^2}$$

where s_ϵ is the sphericity of the particles, D_p is the particle diameter, ϕ is the porosity of the porous medium.

Parameters

The Kozeny-Carman Permeability Model can be called in the `<Constitutive>` block of the input XML file. This permeability model must be assigned a unique name via the `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute in the `<ElementRegions>` block.

The following attributes are supported:

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
particleDiameter	real64	required	Diameter of the spherical particles.
sphericity	real64	required	Sphericity of the particles.

Example

```
<Constitutive>
...
  <CarmanKozenyPermeability name="matrixPerm"
    particleDiameter="0.0002"
    sphericity="1.0"/>
...
</Constitutive>
```

Parallel Plates Permeability Model

Overview

The parallel plates permeability model defines the relationship between the hydraulic fracture aperture and its corresponding permeability following the classic lubrication model ([Witherspoon et al.](#)). In this model, the two fracture walls are assumed to be smooth and parallel to each other and separated by a uniform aperture.

$$k = \frac{a^3}{12}$$

where a denotes the hydraulic fracture aperture.

Remark: k , dimensionally, is not a permeability (as it is expressed in m^3).

Parameters

The Parallel Plates Permeability Model can be called in the `<Constitutive>` block of the input XML file. This permeability model must be assigned a unique name via the `name` attribute. This name is used to assign the model to regions of the physical domain via a `materialList` attribute in the `<ElementRegions>` block.

The following attributes are supported:

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

5.4.7 Porous Solids

Overview

Simulation of fluid flow in porous media and of poromechanics, requires to define, along with fluid properties, the hydrodynamical properties of the solid matrix. Thus, for porous media flow and poromechanical simulation in GEOSX, two types of composite constitutive models can be defined to specify the characteristics of a porous material: (1) a *CompressibleSolid* model, used for flow-only simulations and which assumes that all poromechanical effects can be represented by the pressure dependency of the porosity; (2) a *PorousSolid* model which, instead, allows to couple any solid model with a *BiotPorosity* model and to include permeability's dependence on the mechanical response.

Both these composite models require the names of the solid, porosity and permeability models that, combined, define the porous material. The following sections outline how these models can be defined in the Constitutive block of the xml input files and which type of submodels they allow for.

CompressibleSolid

This composite constitutive model requires to define a *NullModel* as solid model (since no mechanical properties are used), a *PressurePorosity* model and any type of *Permeability* model.

To define this composite model the keyword *CompressibleSolid* has to be appended to the name of the permeability model of choice, as shown in the following example for the *ConstantPermeability* model.

```
<Constitutive>
  <CompressibleSolidConstantPermeability name="porousRock"
    solidModelName="nullSolid"
    porosityModelName="rockPorosity"
    permeabilityModelName="rockPermeability"/>

  <NullModel name="nullSolid"/>

  <PressurePorosity name="rockPorosity"
    referencePressure="1.0e27"
    defaultReferencePorosity="0.3"
    compressibility="1.0e-9"/>

  <ConstantPermeability name="rockPermeability"
    permeabilityComponents="{ 1.0e-4, 1.0e-4, 1.0e-4 }"/>
</Constitutive>
```

PorousSolid

To run poromechanical problems, the total stress is decomposed into an “effective stress” (driven by mechanical deformations) and a pore fluid pressure component, following the [Biot theory of poroelasticity](#). For single-phase flow, or multiphase problems with no capillarity, this decomposition reads

$$\sigma_{ij} = \sigma'_{ij} - bp\delta_{ij}$$

where σ_{ij} is the ij component of the total stress tensor, σ'_{ij} is the ij component of the effective (Cauchy) stress tensor, b is Biot’s coefficient, p is fluid pressure, and δ is the Kronecker delta.

The *PorousSolid* models simply append the keyword Porous in front of the solid model they contain, e.g., PorousElasticIsotropic, PorousDruckerPrager, and so on. Additionally, they require to define a *BiotPorosity* model and a *ConstantPermeability* model. For example, a Poroelastic material with a certain permeability can be defined as

```
<Constitutive>
  <PorousElasticIsotropic name="porousRock"
    porosityModelName="rockPorosity"
    solidModelName="rockSkeleton"
    permeabilityModelName="rockPermeability"/>

  <ElasticIsotropic name="rockSkeleton"
    defaultDensity="0"
    defaultYoungModulus="1.0e4"
    defaultPoissonRatio="0.2"/>

  <BiotPorosity name="rockPorosity"
    grainBulkModulus="1.0e27"
    defaultReferencePorosity="0.3"/>

  <ConstantPermeability name="rockPermeability"
    permeabilityComponents="{ 1.0e-4, 1.0e-4, 1.0e-4 }"/>
</Constitutive>
```

Note that any of the previously described solid models is used by the *PorousSolid* model to compute the effective stress, leading to either poro-elastic, poro-plastic, or poro-damage behavior depending on the specific model chosen.

In an input XML file, constitutive models are listed in the <Constitutive> block. Each parameterized model has its own XML tag, and each must be assigned a unique name via the name attribute. Names are used to assign models to regions of the physical domain via the materialList attribute of the <ElementRegion> node (see [Element: ElementRegions](#)). In some cases, physics solvers must also be assigned specific constitutive models to use (see [Physics Solvers](#)).

A typical <Constitutive> and <ElementRegions> block will look like:

```
<Problem>

  <Constitutive>

    <!-- Define a compressible, single-phase fluid called "water"-->
    <CompressibleSinglePhaseFluid
      name="water"
      referencePressure="2.125e6"
      referenceDensity="1000"
      compressibility="1e-19"
      referenceViscosity="0.001"
      viscosibility="0.0"/>
```

(continues on next page)

(continued from previous page)

```

</Constitutive>

<ElementRegions>

  <!--Add water to the material list for region 1-->
  <ElementRegion
    name="region1"
    cellBlocks="cellBlock1"
    materialList="water"/>

</ElementRegions>

... remainder of problem definition here ...

</Problem>

```

5.5 Event Management

The goal of the GEOSX event manager is to be flexible with regards to event type, application order, and method of triggering. The event manager is configured via the `Event` block in an input .xml file, i.e.:

```

<Events maxTime="1.0e-2">
  <PeriodicEvent name="event_a"
    target="/path/to/event"
    forceDt="1" />
  <HaltEvent name="event_b"
    target="/path/to/halt_target"
    maxRunTime="1e6" />
</Events>

```

5.5.1 Event Execution Rules

The `EventManager` will repeatedly iterate through a list of candidate events specified via the `Events` block **in the order they are defined in the xml**. When certain user-defined criteria are met, they will trigger and perform a task. The `simulation cycle` denotes the number of times the primary event loop has completed, `time` denotes the simulation time at the beginning of the loop, and `dt` denotes the global timestep during the loop.

During each cycle, the `EventManager` will do the following:

1. Loop through each event and obtain its timestep request by considering:
 - a. The maximum `dt` specified via the target's `GetTimestepRequest` method
 - b. The time remaining until user-defined points (e.g. application start/stop times)
 - c. Any timestep overrides (e.g. user-defined maximum `dt`)
 - d. The timestep request for any of its children
2. Set the cycle `dt` to the smallest value requested by any event
3. Loop through each event and:
 - a. Calculate the event `forecast`, which is defined as the expected number of cycles until the event is expected to execute.

- b. `if (forecast == 1)` the event will signal its target to prepare to execute. This is useful for preparing time-consuming I/O operations.
 - c. `if (forecast <= 0)` the event will call the `Execute` method on its target object
4. Check to see if the `EventManager` exit criteria have been met

After exiting the main event loop, the `EventManager` will call the `Cleanup` method for each of its children (to produce final plots, etc.). Note: if the code is resuming from a restart file, the `EventManager` will pick up exactly where it left off in the execution loop.

5.5.2 Event Manager Configuration

Event

The children of the `Event` block define the events that may execute during a simulation. These may be of type `HaltEvent`, `PeriodicEvent`, or `SoloEvent`. The exit criteria for the global event loop are defined by the attributes `maxTime` and `maxCycle` (which by default are set to their max values). If the optional `logLevel` flag is set, the `EventManager` will report additional information with regards to timestep requests and event forecasts for its children.

Name	Type	Default	Description
<code>logLevel</code>	integer	0	Log level
<code>maxCycle</code>	integer	2147483647	Maximum simulation cycle for the global event loop.
<code>maxTime</code>	real64	1.79769e+308	Maximum simulation time for the global event loop.
<code>HaltEvent</code>	node		<i>Element: <code>HaltEvent</code></i>
<code>PeriodicEvent</code>	node		<i>Element: <code>PeriodicEvent</code></i>
<code>SoloEvent</code>	node		<i>Element: <code>SoloEvent</code></i>

PeriodicEvent

This is the most common type of event used in GEOSX. As its name suggests, it will execute periodically during a simulation. It can be triggered based upon a user-defined `cycleFrequency` or `timeFrequency`.

If `cycleFrequency` is specified, the event will attempt to execute every `X` cycles. Note: the default behavior for a `PeriodicEvent` is to execute every cycle. The event forecast for this case is given by: `forecast = cycleFrequency - (cycle - lastCycle)`.

If `timeFrequency` is specified, the event will attempt to execute every `X` seconds (this will override any cycle-dependent behavior). By default, the event will attempt to modify its timestep requests to respect the `timeFrequency` (this can be turned off by specifying `targetExactTimestep="0"`). The event forecast for this case is given by: `if (dt > 0), forecast = (timeFrequency - (time - lastTime)) / dt, otherwise forecast=max`

By default, a `PeriodicEvent` will execute throughout the entire simulation. This can be restricted by specifying the `beginTime` and/or `endTime` attributes. Note: if either of these values are set, then the event will modify its timestep requests so that a cycle will occur at these times (this can be turned off by specifying `targetExactStartStop="0"`).

The timestep request event is typically determined via its target. However, this value can be overridden by setting the `forceDt` or `maxEventDt` attributes.

Name	Type	De- fault	Description
begin-Time	real64	0	Start time of this event.
cycleFrequency	integer	1	Event application frequency (cycle, default)
endTime	real64	1e+100	End time of this event.
finalDt-Stretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the endTime exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
function	string		Name of an optional function to evaluate when the time/cycle criteria are met. If the result is greater than the specified eventThreshold, the function will continue to execute.
logLevel	integer	0	Log level
max-EventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
name	string	required	A name is required for any non-unique nodes
object	string		If the optional function requires an object as an input, specify its path here.
set	string		If the optional function is applied to an object, specify the setname to evaluate (default = everything).
stat	integer	0	If the optional function is applied to an object, specify the statistic to compare to the eventThreshold. The current options include: min, avg, and max.
target	string		Name of the object to be executed when the event criteria are met.
targetExactStart-Stop	integer	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
targetExact-Timestep	integer	1	If this option is set, the event will reduce its timestep requests to match the specified timeFrequency perfectly: $dt_request = \min(dt_request, t_last + time_frequency - time)$.
threshold	real64	0	If the optional function is used, the event will execute if the value returned by the function exceeds this threshold.
timeFrequency	real64	-1	Event application frequency (time). Note: if this value is specified, it will override any cycle-based behavior.
HaltEvent	node		<i>Element: HaltEvent</i>
PeriodicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

SoloEvent

This type of event will execute once once the event loop reaches a certain cycle (targetCycle) or time (targetTime). Similar to the PeriodicEvent type, this event will modify its timestep requests so that a cycle occurs at the exact time requested (this can be turned off by specifying targetExactTimestep="0"). The forecast calculations follow an similar approach to the PeriodicEvent type.

Name	Type	De- fault	Description
beginTime	real64	0	Start time of this event.
endTime	real64	1e+100	End time of this event.
finalDt-Stretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the end-Time exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
logLevel	in- te- ger	0	Log level
max-EventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
name	string	re- quired	A name is required for any non-unique nodes
target	string		Name of the object to be executed when the event criteria are met.
targetCy- cle	in- te- ger	-1	Targeted cycle to execute the event.
targetEx- actStart- Stop	in- te- ger	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
targe- tExact- Timestep	in- te- ger	1	If this option is set, the event will reduce its timestep requests to match the specified execution time exactly: $dt_request = \min(dt_request, t_target - time)$.
targetTime	real64	-1	Targeted time to execute the event.
HaltEvent	node		<i>Element: HaltEvent</i>
Peri- odicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

HaltEvent

This event type is designed to track the wall clock. When the time exceeds the value specified via maxRunTime, the event will trigger and set a flag that instructs the main EventManager loop to cleanly exit at the end of the current cycle. The event for cast for this event type is given by: $forecast = (maxRuntime - (currentTime - startTime)) / realDt$

Name	Type	De- fault	Description
beginTime	real64	0	Start time of this event.
endTime	real64	1e+100	End time of this event.
finalDtStretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the endTime exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
logLevel	integer	0	Log level
maxEventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
maxRuntime	real64	required	The maximum allowable runtime for the job.
name	string	required	A name is required for any non-unique nodes
target	string		Name of the object to be executed when the event criteria are met.
targetExact-StartStop	integer	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
HaltEvent	node		<i>Element: HaltEvent</i>
PeriodicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

5.5.3 Other Event Features

Event Progress Indicator

Because the event manager allows the user to specify the order of events, it could introduce ambiguity into the timestamps of output files. To resolve this, we pass two arguments to the target's Execute method:

1. eventCounter (integer) - the application index for the event (or sub-event)
2. eventProgress (real64) - the percent completion of the event loop, paying attention to events whose targets are associated with physics (from the start of the event, indicated via target->GetTimestepBehavior())

For example, consider the following Events block:

```
<Events maxTime="1.0e-2">
  <PeriodicEvent name="outputs"
    timeFrequency="1e-6"
    targetExactTimestep="0"
    target="/Outputs/siloOutput">
  <PeriodicEvent name="solverApplications_a"
    forceDt="1.0e-5"
    target="/Solvers/lagsolve" />
  <PeriodicEvent name="solverApplications_b"
    target="/Solvers/otherSolver" />
  <PeriodicEvent name="restarts"
    timeFrequency="5.0e-4"
    targetExactTimestep="0"
    target="/Outputs/restartOutput"/>
</Events>
```

In this case, the events solverApplications_a and solverApplications_b point target physics events. The eventCounter,

eventProgress pairs will be: outputs (0, 0.0), solverApplications_a (1, 0.0), solverApplications_b (2, 0.5), and restarts (3, 1.0). These values are supplied to the target events via their Execute methods for use. For example, for the name of a silo output file will have the format: “%s_%06d%02d” % (name, cycle, eventCounter), and the time listed in the file will be $\text{time} = \text{time} + \text{dt} * \text{eventProgress}$

Nested Events

The event manager allows its child events to be nested. If this feature is used, then the manager follows the basic execution rules, with the following exception: When its criteria are met, an event will first execute its (optional) target. It will then estimate the forecast for its own sub-events, and execute them following the same rules as in the main loop. For example:

```
<Events maxTime="1.0e-2">
  <PeriodicEvent name="event_a"
    target="/path/to/target_a" />

  <PeriodicEvent name="event_b"
    timeFrequency="100">

    <PeriodicEvent name="subevent_b_1"
      target="/path/to/target_b_1"/>

    <PeriodicEvent name="subevent_b_2"
      target="/path/to/target_b_2"/>

  </PeriodicEvent/>
</Events>
```

In this example, event_a will trigger during every cycle and call the Execute method on the object located at /path/to/target_a. Because it is time-driven, event_b will execute every 100 s. When this occurs, it will execute it will execute its own target (if it were defined), and then execute subevent_b_1 and subevent_b_2 in order. Note: these are both cycle-driven events which, by default would occur every cycle. However, they will not execute until each of their parents, grandparents, etc. execution criteria are met as well.

5.6 Tasks Manager

The GEOSX tasks manager allows a user to specify tasks to be executed. These tasks are compatible targets for the *Event Management*.

The tasks manager is configured via the Tasks block in an input .xml file, i.e.:

```
<Tasks>
  <PackCollection name="historyCollection" objectPath="nodeManager" fieldName=
    "Velocity" />
</Tasks>
```

5.6.1 Tasks Manager Configuration

Task

The children of the Tasks block define different Tasks to be triggered by events specified in the *Event Management* during the execution of the simulation. At present the only supported task is the PackCollection used to collect time history data for output by a TimeHistory output.

Name	Type	Default	Description
PackCollection	node		<i>Element: PackCollection</i>
TriaxialDriver	node		<i>Element: TriaxialDriver</i>

PackCollection

The `PackCollection` Task is used to collect time history information from fields. Either the entire field or specified named sets of indices in the field can be collected.

Name	Type	Default	Description
fieldName	string	required	The name of the (packable) field associated with the specified object to retrieve data from
name	string	required	A name is required for any non-unique nodes
objectPath	string	required	The name of the object from which to retrieve field values.
onlyOn-SetChange	localIndex	0	Whether or not to only collect when the collected sets of indices change in any way.
setNames	string_array	{ }	The set(s) for which to retrieve data.

Note: The time history information collected via this task is buffered internally until it is output by a linked `TimeHistory` Output.

Triggering the Tasks

Tasks can be triggered using the *Event Management*. Recurring tasks should use a `<PeriodicEvent>` and one-time tasks should use a `<SoloEvent>`:

```
<PeriodicEvent name="historyCollectEvent"
  timeFrequency="1.0"
  targetExactTimeset="1"
  target="/Tasks/historyCollection" />
```

The keyword `target` has to match the name of a Task specified as a child of the `<Tasks>` block.

5.7 Functions

Functions are the primary avenue for specifying values that change in space, time, or any other dimension. These are specified in the `Functions` block, and may be referenced by name throughout the rest of the .xml file. For example:

```
<Functions>
  <TableFunction name="q"
    inputVarNames="time"
    coordinates="0 60 1000"
    values="0 1 1" />
</Functions>
<FieldSpecifications>
  <SourceFlux name="sourceTerm"
```

(continues on next page)

(continued from previous page)

```
objectPath="ElementRegions/Region1/block1"  
scale="0.001"  
functionName="q"  
setNames="{source}"/>  
</FieldSpecifications>
```

5.7.1 Function Inputs and Application

The inputs to each function type are specified via the `inputVarName` attribute. These can either be the name of an array (e.g. “Pressure”) or the special keyword “time” (time at the beginning of the current cycle). If any of the input variables are vectors (e.g. “referencePosition”), the components will be given as function arguments in order.

In the .xml file, functions are referenced by name. Depending upon the application, the functions may be applied in one of three ways:

1. Single application: The function is applied to get a single scalar value. For example, this could define the flow rate applied via a BC at a given time.
2. Group application: The function is applied to a (user-specified) `ManagedGroup` of size `N`. When called, it will iterate over the `inputVarNames` list and build function inputs from the group’s wrappers. The resulting value will be a wrapper of size `N`. For example, this could be used to apply a user-defined constitutive relationship or specify a complicated boundary condition.
3. Statistics mode: The function is applied in the same manner as the group application, except that the function will return an array that contains the minimum, average, and maximum values of the results.

5.7.2 Function Types

There are three types of functions available for use: `TableFunction`, `SymbolicFunction`, and `CompositeFunction`. **Note: the symbolic and composite function types are currently only available for x86-64 systems.**

TableFunction

A table function uses a set of pre-computed values defined at points on a structured grid to represent an arbitrary-dimensional function. Typically, the axes of the table will represent time and/or spatial dimensions; however, these can be applied to represent phase diagrams, etc.

Name	Type	Default	Description
coordinateFiles	path_array	{}	List of coordinate file names for ND Table
coordinates	real64_array	{0}	Coordinates inputs for 1D tables
inputVarNames	string_array	{}	Name of fields are input to function.
interpolation	geosx_TableFunction_InterpolationType		Interpolation method. Valid options: * linear * nearest * upper * lower
name	string	required	A name is required for any non-unique nodes
values	real64_array	{0}	Values for 1D tables
voxelFile	path		Voxel file name for ND Table

1D Table

For 1D tables, the function may be defined using the coordinates and values attributes. These represent the location of the grid nodes (ordered from smallest to largest) and the function values at those points, respectively. For example, the following function defines a simple ramp function with a rise-time of 60 seconds:

```
<TableFunction name="q"
  inputVarNames="time"
  coordinates="0 60 1000"
  values="0 1 1" />
```

ND Table

For ND tables, the grid coordinates and values may be defined using a set of .csv files. The `coordinateFiles` attribute specifies the file names that define the coordinates for each axis. The values in each coordinate file must be comma-delimited and ordered from smallest to largest. The dimensionality of the table is defined by the number of coordinate files (`coordinateFiles="x.csv"` would indicate a 1D table, `coordinateFiles="x.csv y.csv z.csv t.csv"` would indicate a 4D table, etc.). The `voxelFile` attribute specifies name of the file that defines the value of the function at each point along the grid. These values must be comma-delimited (line-breaks are allowed) and be specified in Fortran order, i.e., column-major order (where the index of the first dimension changes the fastest, and the index of the last dimension changes slowest).

The following would define a simple 2D function $c = a + 2*b$:

```
<TableFunction name="c"
  inputVarNames="a b"
  coordinateFiles="a.csv b.csv"
  voxelFile="c.csv" />
```

- a.csv: "0, 1"

- b.csv: “0, 0.5, 1”
- c.csv: “0, 1, 1, 2, 2, 3”

Interpolation Methods

There are four interpolation methods available for table functions. Within the table axes, these will return a value:

- linear: using piecewise-linear interpolation
- upper: equal to the value of the next table vertex
- nearest: equal to the value of the nearest table vertex
- lower: equal to the value of the previous table vertex

Outside of the table axes, these functions will return the edge-values. The following figure illustrates how each of these methods work along a single dimension, given identical table values:




Table Generation Example

The following is an example of how to generate the above tables in Python:

```
import numpy as np

# Define table axes
a = np.array([0.0, 1.0])
b = np.array([0.0, 0.5, 1.0])

# Generate table values (note: the indexing argument is important)
A, B = np.meshgrid(a, b, indexing='ij')
C = A + 2.0*B

# Write axes, value files
np.savetxt('a.csv', a, fmt='%1.2f', delimiter=',')
np.savetxt('b.csv', b, fmt='%1.2f', delimiter=',')
values = np.reshape(C, (-1), order='F')
np.savetxt('c.csv', values, fmt='%1.2f', delimiter=',')
```

SymbolicFunction

This function leverages the symbolic expression library mathpresso to define and evaluate functions. These functions are processed using an x86-64 JIT compiler, so are nearly as efficient as natively compiled C++ expressions.

Name	Type	De- fault	Description
expression	string	re- quired	Symbolic math expression
inputVar- Names	string_array	{ }	Name of fields are input to function.
name	string	re- quired	A name is required for any non-unique nodes
variable- Names	string_array	re- quired	List of variables in expression. The order must match the evaluate argument

The `variableNames` attribute defines a set of single-character names for the inputs to the symbolic function. There should be a definition for each scalar input and for each component of a vector input. For example if `inputVarName="time, ReferencePosition"`, then `variableNames="t, x, y, z"`. The `expression` attribute defines the symbolic expression to be executed. Aside from the following exceptions, the syntax mirrors python:

- The function string cannot contain any spaces
- The power operator is specified using the C-style expression (e.g. `pow(x,3)` instead of `x**3`)

The following would define a simple 2D function `c = a + 2*b`:

```
<SymbolicFunction name="c"
    inputVarNames="a b"
    variableNames="x y"
    expression="x+(2*y)"/>
```

CompositeFunction

This function is derived from the symbolic function. However, instead of using the time or object as inputs, it is used to combine the outputs of other functions using a symbolic expression.

Name	Type	De- fault	Description
expression	string		Composite math expression
function-Names	string_array	{ }	List of source functions. The order must match the variableNames argument.
inputVar-Names	string_array	{ }	Name of fields are input to function.
name	string	re- quired	A name is required for any non-unique nodes
variable-Names	string_array	{ }	List of variables in expression

The `functionNames` attribute defines the set of input functions to use (these may be of any type, and may each have any number of inputs). The `variableNames` attribute defines a set of single-character names for each function. The `expression` attribute defines the symbolic expression, and follows the same rules as above. The `inputVarNames` attribute is ignored for this function type.

The following would define a simple 1D table function $f(t) = 1 + t$, a 3D symbolic function $g(x, y, z) = x^2 + y^2 + z^2$, and a 4D composite function $h = \sin(f(t)) + g(x, y, z)$:

```
<Functions>
  <TableFunction name="f"
    inputVarNames="time"
    coordinates="0 1000"
    values="1 1001" />

  <SymbolicFunction name="g"
    inputVarNames="ReferencePosition"
    variableNames="x y z"
    expression="pow(x,2)+pow(y,2)+pow(z,2)"/>

  <CompositeFunction name="h"
    inputVarNames="ignored"
    functionNames="f g"
    variableNames="x y"
    expression="sin(x)+y"/>
</Events>
```

5.8 Linear Solvers

5.8.1 Introduction

Any physics solver relying on standard finite element and finite volume techniques requires the solution of algebraic linear systems, which are obtained upon linearization and discretization of the governing equations, of the form:

$$Ax = b$$

with a A a square sparse matrix, x the solution vector, and b the right-hand side. For example, in a classical linear elastostatics problem A is the stiffness matrix, and x and b are the displacement and nodal force vectors, respectively.

This solution stage represents the most computationally expensive portion of a typical simulation. Solution algorithms generally belong to two families of methods: direct methods and iterative methods. In GEOSX both options are made available wrapping around well-established open-source linear algebra libraries, namely [HYPRE](#), [PETSc](#), [SuperLU](#), and [Trilinos](#).

5.8.2 Direct methods

The major advantages are their reliability, robustness, and ease of use. However, they have large memory requirements and exhibit poor scalability. Direct methods should be used in a prototyping stage, for example when developing a new formulation or algorithm, when the dimension of the problem, namely the size of matrix A , is small. Irrespective of the selected direct solver implementation, three stages can be identified:

- (1) **Setup Stage:** the matrix is first analyzed and then factorized
- (2) **Solve Stage:** the solution to the linear systems involving the factorized matrix is computed
- (3) **Finalize Stage:** the systems involving the factorized matrix have been solved and the direct solver lifetime ends

The default option in GEOSX relies on [SuperLU](#), a general purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations, that is called taking advantage of the interface provided in [HYPRE](#).

5.8.3 Iterative methods

As the problem size (number of computational cells) increases, global iterative solution strategies are the method of choice—typically nonsymmetric Krylov solvers. Because of the possible poor conditioning of A , preconditioning is essential to solve such systems efficiently. *“Preconditioning is simply a means of transforming the original linear system into one which has the same solution, but which is likely to be easier to solve with an iterative solver”* [Saad (2003)].

The design of a robust and efficient preconditioner is based on a trade-off between two competing objectives:

- **Robustness:** reducing the number of iterations needed by the preconditioned solver to achieve convergence;
- **Efficiency:** limiting the time required to construct and apply the preconditioner.

Assuming a preconditioning matrix M is available, three standard approaches are used to apply the preconditioner:

- (1) **Left preconditioning:** the preconditioned system is $M^{-1}Ax = M^{-1}b$
- (2) **Right preconditioning:** the preconditioned system is $AM^{-1}y = b$, with $x = M^{-1}y$
- (3) **Split preconditioning:** the preconditioned system is $M_L^{-1}AM_R^{-1}y = M_L^{-1}b$, with $x = M_R^{-1}y$

5.8.4 Summary

The following table summarizes the available input parameters for the linear solver.

Name	Type	Default	Description
amgAggressiveCoarseningLevel	integer	0	AMG number levels for aggressive coarsening Available options are: TODO
amgCoarseSolver	geosx_LinearSolverParameters	direct AMG_CoarseType	AMG coarsest level solver/smoother type. Available options are: default\ jacobi\ ll
amgCoarseningType	string	HMIS	AMG coarsening algorithm Available options are: TODO
amgInterpolationType	integer	6	AMG interpolation algorithm Available options are: TODO
amgNullSpaceType	geosx_LinearSolverParameters	constantModes	AMG near null space approximation. Available options are: constantModes\ rigidBodyMode
amgNumFunctions	integer	1	AMG number of functions Available options are: TODO
amgNumSweeps	integer	2	AMG smoother sweeps
amgSmootherType	geosx_LinearSolverParameters	gs AMG_SmootherType	AMG smoother type. Available options are: default\ jacobi\ ll
amgThreshold	real64	0	AMG strength-of-connection threshold
directCheckResidual	integer	0	Whether to check the linear system solution residual
directColPerm	geosx_LinearSolverParameters	none Direct_ColPerm	How to permute the columns. Available options are: none\ MMD_AtA\ col
directEquil	integer	1	Whether to scale the rows and columns of the matrix
directIterRef	integer	1	Whether to perform iterative refinement
directParallel	integer	1	Whether to use a parallel solver (instead of a serial one)
5.8. Linear Solvers			223
directReplTinyPivot	integer	1	Whether to replace tiny pivots by sqrt(epsilon)*norm(A)

5.8.5 Preconditioner descriptions

This section provides a brief description of the available preconditioners.

- **None:** no preconditioning is used, i.e., $M^{-1} = I$.
- **Jacobi:** diagonal scaling preconditioning, with $M^{-1} = D^{-1}$, with D the matrix diagonal. Further details can be found in:
 - [HYPRE documentation](#),
 - [PETSc documentation](#),
 - [Trilinos documentation](#).
- **ILUK:** incomplete LU factorization with fill level k of the original matrix: $M^{-1} = U^{-1}L^{-1}$. Further details can be found in:
 - [HYPRE documentation](#),
 - [PETSc documentation](#),
 - [Trilinos documentation](#).
- **ILUT:** a dual threshold incomplete LU factorization: $M^{-1} = U^{-1}L^{-1}$. Further details can be found in:
 - [HYPRE documentation](#),
 - not yet available through PETSc interface,
 - [Trilinos documentation](#).
- **ICC:** incomplete Cholesky factorization of a symmetric positive definite matrix: $M^{-1} = L^{-T}L^{-1}$. Further details can be found in:
 - not yet available through *hypre* interface,
 - [PETSc documentation](#),
 - [Trilinos documentation](#).
- **AMG:** algebraic multigrid (can be classical or aggregation-based according to the specific package). Further details can be found in:
 - [HYPRE documentation](#),
 - [PETSc documentation](#),
 - [Trilinos documentation](#).
- **MGR:** multigrid reduction. Available through *hypre* interface only. Specific documentation coming soon. Further details can be found in [MGR documentation](#).
- **Block:** custom preconditioner designed for a 2 x 2 block matrix.

5.8.6 HYPRE MGR Preconditioner

MGR stands for multigrid reduction, a multigrid method that uses the interpolation, restriction operators, and the Galerkin triple product, to reduce a linear system to a smaller one, similar to a Schur complement approach. As such, it is designed to target block linear systems resulting from discretizations of multiphysics problems. GEOSX uses MGR through an implementation in [HYPRE](#). More information regarding MGR can be found [here](#). Currently, MGR strategies are implemented for hydraulic fracturing, poroelastic, compositional flow with and without wells. More multiphysics solvers with MGR will be enabled in the future.

To use MGR for a specific block system, several components need to be specified.

- (1) **The number of reduction levels and the coarse points (corresponding to fields) for each level.** For example, for single-phase hydraulic fracturing, there are two fields, i.e. displacement and fluid pressure, a two-level MGR strategy can be used with the fluid pressure being the coarse degrees of freedom.
- (2) **Interpolation/restriction operators and the coarse-grid computation strategy.** A simple but effective strategy is to use Jacobi diagonalization for interpolation and injection for restriction. For most cases, a Galerkin coarse grid strategy can be used, but for special cases such as poroelastic, a non-Galerkin approach is preferable.
- (3) **Global smoother.** Depending on the problem, a global relaxation step could be beneficial. Some options include ILU(k), (block) Jacobi, (block) Gauss-Seidel.
- (4) **Solvers for F-relaxation and coarse-grid correction.** These solvers should be chosen carefully for MGR to be effective. The choice of these solvers should correspond to the properties of the blocks specified by the C- and F-points. For example, if the A_{FF} block is hyperbolic, a Jacobi smoother is sufficient while for an elliptic operator an AMG V-cycle might be required. For the single-phase hydraulic fracturing case, an AMG V-cycle is needed for both F-relaxation and coarse-grid correction.

Note that these are only general guidelines for designing a working MGR recipe. For complicated multiphysics problems, experimentation with different numbers of levels, choices of C- and F-points, and smoothers/solvers, etc., is typically needed to find the best strategy. Currently, these options are only available to developers. We are working on exposing these functionalities to the users in future releases.

5.8.7 Block preconditioner

This framework allows the user to design a block preconditioner for a 2×2 block matrix. The key component is the Schur complement $S = A_{11} - A_{10}\tilde{A}_{00}^{-1}A_{01}$ computation, that requires an approximation of the leading block. Currently, available options for \tilde{A}_{00}^{-1} are:

- diagonal with diagonal values (essentially, a Jacobi preconditioner);
- diagonal with row sums as values (e.g., used for CPR-like preconditioners).

Once the Schur complement is computed, to properly define the block preconditioner we need:

- the preconditioner for A_{00} (any of the above listed single-matrix preconditioner);
- the preconditioner for S (any of the above listed single-matrix preconditioner);
- the application strategy. This can be:
 - diagonal: none of the coupling terms is used;
 - upper triangular: only the upper triangular coupling term is used;
 - lower-upper triangular: both coupling terms are used.

Moreover, a block scaling is available. Feasible options are:

- none: keep the original scaling;
- Frobenius norm: equilibrate Frobenius norm of the diagonal blocks;
- user provided.

5.9 Numerical Methods

This section describes the specification of numerical methods used by solvers.

Name	Type	Default	Description
FiniteElements	node	unique	<i>Element: FiniteElements</i>
FiniteVolume	node	unique	<i>Element: FiniteVolume</i>

5.9.1 Finite Element Discretization

We are currently refactoring the finite element infrastructure, and will update the documentation soon to reflect the new structure.

5.9.2 Finite Volume Discretization

Two different finite-volume discretizations are available to simulate single-phase flow in GEOSX, namely, a standard cell-centered TPFA approach, and a hybrid finite-volume scheme relying on both cell-centered and face-centered degrees of freedom. The key difference between these two approaches is the computation of the flux, as detailed below.

Standard cell-centered TPFA FVM

This is the standard scheme implemented in the *SinglePhaseFVM* flow solver. It only uses cell-centered degrees of freedom and implements a Two-Point Flux Approximation (TPFA) for the computation of the flux. The numerical flux is obtained using the following expression for the mass flux between cells K and L :

$$F_{KL} = \Upsilon_{KL} \frac{\rho^{upw}}{\mu^{upw}} (p_K - p_L - \rho^{avg} g(d_K - d_L)),$$

where p_K is the pressure of cell K , d_K is the depth of cell K , and Υ_{KL} is the standard TPFA transmissibility coefficient at the interface. The fluid density, ρ^{upw} , and the fluid viscosity, μ^{upw} , are upwinded using the sign of the potential difference at the interface.

This is currently the only available discretization in the *Compositional Multiphase Flow Solver*.

Hybrid FVM

This discretization scheme overcomes the limitations of the standard TPFA on non K-orthogonal meshes. The hybrid finite-volume scheme—equivalent to the well-known hybrid Mimetic Finite Difference (MFD) scheme—remains consistent with the pressure equation even when the mesh does not satisfy the K-orthogonality condition. This numerical scheme is currently implemented in the *SinglePhaseHybridFVM* solver.

The hybrid FVM scheme uses both cell-centered and face-centered pressure degrees of freedom. The one-sided face flux, $F_{K,f}$, at face f of cell K is computed as:

$$F_{K,f} = \frac{\rho^{upw}}{\mu^{upw}} \tilde{F}_{K,f},$$

where $\tilde{F}_{K,f}$ reads:

$$\tilde{F}_{K,f} = \sum_{f'} \Upsilon_{ff'} (p_K - \pi_f - \rho_K g(d_K - d_f)).$$

In the previous equation, p_K is the cell-centered pressure, π_f is the face-centered pressure, d_K is the depth of cell K , and d_f is the depth of face f . The fluid density, ρ^{upw} , and the fluid viscosity, μ^{upw} , are upwinded using the sign of $\tilde{F}_{K,f}$. The local transmissibility Υ of size $n_{local\ faces} \times n_{local\ faces}$ satisfies:

$$NK = \Upsilon C$$

Above, N is a matrix of size $n_{local\ faces} \times 3$ storing the normal vectors to each face in this cell, C is a matrix of size $n_{local\ faces} \times 3$ storing the vectors from the cell center to the face centers, and K is the permeability tensor. The local transmissibility matrix, Υ , is currently computed using the quasi-TPFA approach described in Chapter 6 of this [book](#). The scheme reduces to the TPFA discretization on K-orthogonal meshes but remains consistent when the mesh does not satisfy this property. The mass flux $F_{K,f}$ written above is then added to the mass conservation equation of cell K .

In addition to the mass conservation equations, the hybrid FVM involves algebraic constraints at each mesh face to enforce mass conservation. For a given interior face f between two neighboring cells K and L , the algebraic constraint reads:

$$\tilde{F}_{K,f} + \tilde{F}_{L,f} = 0.$$

We obtain a numerical scheme with n_{cells} cell-centered degrees of freedom and n_{faces} face-centered pressure degrees of freedom. The system involves n_{cells} mass conservation equations and n_{faces} face-based constraints. The linear systems can be efficiently solved using the MultiGrid Reduction (MGR) preconditioner implemented in the Hypr linear algebra package.

The implementation of the hybrid FVM scheme for *Compositional Multiphase Flow Solver* is in progress.

5.10 Parallel Partitioning

Parallel GEOSX simulations involves multiple `partitions` and there are `ghost` objects in each partition. Users need to understand these concepts to effectively design models and visualize results.

5.10.1 Partition and ghosting : simple examples

A model, or more strictly, a computational domain, is stored in a distributed fashion among many processors. In the following simple example, the computational domain has 10 cells and the simulation involves two processors. The first processor, “partition 0” (“0” is called the “rank” of the processor) owns the first five cells (0 to 4) while “partition 1” owns 5 to 9. When the whole domain is divided into partitions, each partition will number the cells and other objects such as nodes, faces, and edges in the partition. Therefore, in both partitions, the cells IDs start from zero. Element 0 in partition 1 is cell 5 (or the sixth cell) of the original domain. In parallel computing, each partition does not only need to know information about its own cells, but it also needs to know information about some cells owned by the neighbor partitions if these cells are directly connected to objects in this partition. For example, cell 0 in partition 1 (i.e. cell 5 in the original whole domain) is connected to cell 4 in partition 0. Therefore, partition 0 will keep a copy of this cell (including the data associated with this cell) which is synchronized with the corresponding information in the partition that actually owns this cell.

In summary, a partition owns a number of cells and other objects (e.g. faces) and also keeps copies of objects from neighbor partitions. Partitioning is automatically handled by GEOSX once the user specifies how the domain should be divided.

The following figure show the partitioning of a simple mesh. Real nodes appear as solid red circles in the owning partition and ghost nodes are shown as hollow circles.

This concept of ghosting and communications between owned cells and ghost cells can also be applied to the other types of elements in GEOSX (Faces, Edges, Nodes). The next figure summarizes the way nodes, edges, faces and cells are ghosted.

5.10.2 Specifying partitioning pattern

Cartesian partitioning

In the command line to run GEOSX, the user can specify the partitioning pattern by adding the following switches:

- `-x`, `--x-partitions` - Number of partitions in the x-direction
- `-y`, `--y-partitions` - Number of partitions in the y-direction
- `-z`, `--z-partitions` - Number of partitions in the z-direction

Graph-based partitioning

The Graph-based partitioning is used only when importing external meshes using the `PAMELAMeshGenerator` (see [Tutorial 3: Regions and Property Specifications](#) section for more details using external meshes). While importing the mesh, PAMELA computes the graph of connectivity between all the volume elements of the mesh. The partitioning is then done using the `METIS` library. The graph is not weighted so the expected result is as mesh divided in `n` parts, with `n` being the number of MPI ranks used for simulation containing a similar amount of cells.

5.10.3 Ghost ranks

Each object (node, edge, face, or cell) has a `ghost_rank` attribute, stored in the `ghostRank` field. If a object does not appear in any other partition as a ghost, its ghost rank is a large negative number, `-2.14e9` in a typical system. If a object is real (owned by the current partition) but exists in other partitions as ghosts, its ghost rank is `-1`. The ghost rank of a ghost object is the rank of the partition that owns the corresponding real object.

5.10.4 Considerations for visualization

In VisIt, a partition is called a `domain`. The ID of a domain is the rank of the corresponding partition in GEOSX plus one. VisIt would display all elements/objects regardless if they are real or ghosts. As information about a ghost is synchronized with the real object, VisIt just overlaying the same images on top of each other. The user would not perceive the overlapping between partitions unless the models are shown as semi-transparent entities. Note that if ghosts are not hidden, results from a `query` operation, such as summation of variable values, would be wrong due to double-counting. Therefore, it is a good practice or habit to hide ghost objects using `ghostRank` as a filter.

If the visualization method involves interpolation, such as interpolating a zonal field into a nodal field or generating contours, the interpretation near partition boundaries is not accurate.

5.11 Outputs

This section describes how outputs are handled by GEOSX

The outputs are defined in a `<Outputs>` XML block.

There are three available formats to output the results of a simulation: `SILO`, `VTK`, and Time History output into simple dataset `HDF5` files which are consumable by post-processing scripts..

5.11.1 Defining an output

SILO Output

The SILO output is defined through the `<Silo>` XML node (subnode of `<Outputs>` XML block) as shown here:

```
<Outputs>
  <Silo name="siloOutput" />
</Outputs>
```

The parameter options are listed in the following table:

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.
plotFileRoot	string	plot	(no description available)
plotLevel	integer	1	(no description available)
writeCellElementMesh	integer	1	(no description available)
writeEdgeMesh	integer	0	(no description available)
writeFEMFaces	integer	0	(no description available)
writeFaceElementMesh	integer	1	(no description available)

VTK Output

The VTK output is defined through the `<VTK>` XML node (subnode of `<Outputs>` XML block) as shown here:

```
<Outputs>
  <VTK name="vtkOutput" />
</Outputs>
```

The parameter options are listed in the following table:

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	re- quired	A name is required for any non-unique nodes
parallelThreads	inte- ger	1	Number of plot files.
plotFileRoot	string	VTK	Name of the root file for this output.
plotLevel	inte- ger	1	Level detail plot. Only fields with lower of equal plot level will be output.
writeBinary- Data	inte- ger	1	Output the data in binary format
writeFEMFaces	inte- ger	0	(no description available)

TimeHistory Output

The TimeHistory output is defined through the `<TimeHistory>` XML node (subnode of `<Outputs>` XML block) as shown here:

```
<Outputs>
  <TimeHistory name="timeHistoryOutput" sources="{/Tasks/collectionTask}" filename=
    ↪ "timeHistory" />
</Outputs>
```

The parameter options are listed in the following table:

Name	Type	Default	Description
childDirectory	string		Child directory path
filename	string	TimeHistory	The filename to which to write time history output.
format	string	hdf	The output file format for time history output.
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.
sources	string_array	required	A list of collectors from which to collect and output time history information.

In order to properly collect and output time history information the following steps must be accomplished:

1. Specify one or more collection tasks using the *Tasks Manager*.
2. Specify a *TimeHistory Output* using the collection task(s) as source(s).
3. Specify an event in the *Event Management* to trigger the collection task(s).
4. Specify an event in the *Event Management* to trigger the output.

Note: Currently if the collection and output events are triggered at the same simulation time, the one specified first will also trigger first. Thus in order to output time history for the current time in this case, always specify the time history collection events prior to the time history output events.

5.11.2 Triggering the outputs

The outputs can be triggered using the *Event Management*. It is recommended to use a `<PeriodicEvent>` to output results with a defined frequency:

```
<PeriodicEvent name="outputs"
  timeFrequency="5000.0"
  targetExactTimestep="1"
  target="/Outputs/siloOutput" />
```

The keyword `target` has to match with the name of the `<Silo>`, `<VTK>`, or `<TimeHistory>` node.

5.11.3 Visualisation of the outputs

We suggest the use of [VisIT](#), [Paraview](#), and [Matplotlib](#) to visualize the outputs.

Visualizing Silo outputs with VisIT

If the `<Silo>` XML node was defined, GEOSX writes the results in a folder called `siloFiles`.

In VisIT :

1. File > Open file...
2. On the right panel, browse to the `silosFiles` folder.
3. On the left panel, select the file(s) you want to visualize. Usually, one file is written according the frequency defined in the `timeFrequency` keyword of the Event that has triggered the output.
4. To load fields, use the “Add” button and browse to the fields you want to plot.
5. To plot fields, use the “Draw” button.

Please consult the [VisIT](#) documentation for further explanations on its usage.

Visualizing VTK outputs with Paraview

If the `<VTK>` XML node was defined, GEOSX writes a folder and a `.pvd` file named after the string defined in `name` keyword.

The `.pvd` file contains references to the `.pvtu` files. One `.pvtu` file is output according the frequency defined in the `timeFrequency` keyword of the Event that has triggered the output.

One `.pvtu` contains references to `.vtu` files. There is as much `.vtu` file as there were MPI processes used for the computation.

All these files can be opened with paraview. To have the whole results for every output time steps, you can open the `.pvd` file.

Visualizing TimeHistory outputs with Matplotlib

If the `<TimeHistory>` XML node was defined, GEOSX writes a file named after the string defined in the `filename` keyword and formatted as specified by the string defined in the `format` keyword (only `HDF5` is currently supported).

The TimeHistory file contains the collected time history information from each specified time history collector. This information includes datasets for the time itself, any metadata sets describing index association with specified collection sets, and the time history information itself.

It is recommended to use [Matplotlib](#) and format-specific accessors (like `H5PY` for `HDF5`) to access and easily plot the time history data.

5.12 Advanced XML Features (geosx_xml_tools)

`geosx_xml_tools` is a python module that enables advanced xml features in GEOSX (parameters, units, symbolic math, etc.), and is used to format xml files.

5.12.1 Setup

The `geosx_xml_tools` package can be installed using the following command within the GEOSX build directory:

```
make geosx_xml_tools
```

During the installation step, two console scripts will be created: `preprocess_xml` and `format_xml`. These will be located within the GEOSX build/bin directory. Additional things to consider:

- The above make command will create a new virtual python environment to install the package. By default, the source python environment will be the the version used to run the *config-build.py* command. The version of python can be changed by specifying the *PYTHON_POST_EXECUTABLE* variable. To build the environment, the *virtualenv* package must be installed in the source distribution.
- The *geosx_xml_tools* package depends on the *lxml* package. If *lxml* is missing from the parent environment, the install script will attempt to fetch an appropriate version from the internet.
- The package may also be manually installed within an existing python distribution (this required administrative privileges) via pip: *pip install src/coreComponents/python/modules/geosx_xml_tools_package*. In this case, the console scripts will be located in the python/bin directory

5.12.2 Usage

geosx_xml_tools can be used via the command-line or can be imported into a python script.

Command-line xml formatting

The following command will update the formatting of an existing xml file in-place:

```
format_xml input_file.xml
```

To update the formatting for all xml files located within the GEOSX repository, you can run the following command within the GEOSX build directory:

```
make geosx_format_all_xml_files
```

Command-line xml preprocessing

The following command will read an xml file, process any advanced xml features located within it, and write a new file that can be read by GEOSX:

```
preprocess_xml input_file.xml
```

The script returns the (randomly generated) name of the new xml file. Optional arguments for this script include:

- *-o/-output*: The desired name for the output file (otherwise, it is randomly generated)
- *-s/-schema*: The location of a schema to validate the final .xml file
- *-v/-verbose*: Increase module verbosity

For convenience, this script can be embedded within a call to GEOSX:

```
srn -n 16 geosx -i 'preprocess_xml input_file.xml' -x 4 -z 4
```

Script-based Example

The *geosx_xml_tools* module can also be called from within a python script. For example:

```
from geosx_xml_tools import xml_processor

initial_filename = 'input.xml'
new_filename = output_name = xml_processor.process(initial_filename)
```

5.12.3 Advanced XML Features

The xml preprocessor in `geosx_xml_tools` is designed to take a raw input file, and generate a new file that can be directly read by GEOSX. The syntax for the advanced XML format is given below. During the processing the order of operations are:

- 1) Merging any included XML files into the root structure
- 2) Substituting in any parameters
- 3) Evaluating unit strings
- 4) Evaluating symbolic math
- 5) Error checking and validation

Including Child XML Files

XML inputs can point to included children (these children can then include grandchildren and so on). During processing, these are recursively inserted into the root XML structure by the following scheme:

- Merge two objects if:
 - At the root level an object with the matching tag exists.
 - If the “name” attribute is present and an object with the matching tag and name exists.
 - Any preexisting attributes are overwritten by the donor.
- Otherwise append the XML structure with the target.

```
<Included>
  <File name='/path/to/included_a.xml' />
  <File name='/path/to/included_b.xml' />
</Included>
```

Parameters

Parameters are a convenient way to build a configurable and human-readable input XML. They are defined via a block in the XML structure. Parameter names may only include upper/lower case letters and underscores (to avoid conflicts with symbolic math). Parameters may have any value:

- Path to a file
- Numbers
- A symbolic expression
- Other parameters
- Etc.

They can be used in any field within in the XML file (except in Includes) as follows:

- `x_par` (preferred)
- `$x_par`
- `$.x_par`
- `$.x_par$`

For Example:

```
<Parameters>
  <Parameter
    name="mu"
    value="0.005"/>
  <Parameter
    name="table_root"
    value="/path/to/table/root"/>
</Parameters>

<Constitutive>
  <CompressibleSinglePhaseFluid
    name="water"
    defaultDensity="1000"
    defaultViscosity="$mu$"
    referencePressure="0.0"
    referenceDensity="1000"
    compressibility="5e-10"
    referenceViscosity="$mu$"
    viscosibility="0.0"/>
</Constitutive>

<Functions>
  <TableFunction
    name="flow_rate"
    inputVarNames="{time}"
    coordinateFiles="{table_root$/time_flow.geos}"
    voxelFile="table_root$/flow.geos"
    interpolation="linear"/>
</Functions>
```

Units

By default, input values are specified using SI units. In some cases, it is useful to override this behavior by explicitly specifying the units of the input. These are specified by appending a valid number with a unit definition in square braces. The unit manager supports most common units and SI prefixes, using both long- and abbreviated names (e.g.: c, centi, k, kilo, etc.). Units may include predefined composite units (dyne, N, etc.) or may be built up from sub-units using a python syntax (e.g.: [N], [kg*m/s**2]). Any (or no) amount of whitespace is allowed between the number and the unit bracket.

Examples:

```
<Parameters>
  <Parameter name='a' value='2[m] '/>
  <Parameter name='b' value='1.2 [cm] '/>
  <Parameter name='c' value='1.23e4 [bbl/day] '/>
  <Parameter name='d' value='1.23E-4 [km**2] '/>
</Parameters>
```

Symbolic Math

Input XML files can also include symbolic mathematical expressions. These are placed within pairs of backticks (`), and use a python syntax. Parameters and units are evaluated before symbolic expressions. Note: symbolic expressions are sanitized before by removing any residual alpha characters, but this can be relaxed if more complicated function are needed. Also, while symbolic expressions are allowed within parameters, errors may occur if these are used in a way that leads to nested symbolic expressions.

Examples:

```
<Parameters>
  <Parameter name='a' value='2[m]' />
  <Parameter name='b' value='1.2 [cm]' />
  <Parameter name='c' value='1.23e4 [bbl/day]' />
  <Parameter name='d' value='1.23E-4 [km**2]' />
</Parameters>
<Geometry>
  <Box
    name='perf'
    xMin='`$a$ - 0.2*$b$`, -1e6, -1e6'
    xMax='`$c$**2 / $d$`, 1e6, 1e6' />
  </Box>
</Geometry>
```

Validation

Unmatched special characters (\$, [, ', etc.) mean that parameters, units, or symbolic math were not specified correctly. If the code detects these, it will throw an error. The XML is validated against the input schema to check if all of the required fields are present, and that input parameters match their expected types.

5.13 pygeosx — GEOSX in Python

GEOSX can be manipulated and executed through a Python script.

High-level control of GEOSX is managed through the top-level `pygeosx` functions, like `initialize` and `run`. GEOSX's data can be manipulated by getting `pylvarray` views of `LvArray` objects living in GEOSX's data repository. These `pylvarray` views are fetched by calling `Wrapper.value()` after getting a `Wrapper` out of the data repository.

Warning: The `pygeosx` module provides plenty of opportunities to crash Python. See the Segmentation Faults section below.

Only Python 3 is supported.

5.13.1 Module Functions

`pygeosx.initialize(rank, args)`

Initialize GEOSX for the first time, with a rank and command-line arguments.

This function should only be called **once**. To reinitialize, use the `reinit` function.

Generally the rank is obtained from the `mpi4py` module and the arguments are obtained from `sys.argv`.

Returns a `Group` representing the `ProblemManager` instance.

`pygeosx.reinit(args)`

Reinitialize GEOSX with a new set of command-line arguments.

Returns a `Group` representing the `ProblemManager` instance.

`pygeosx.apply_initial_conditions()`

Apply the initial conditions.

`pygeosx.finalize()`

Finalize GEOSX. After this no calls into pygeosx or to MPI are allowed.

`pygeosx.run()`

Enter the GEOSX event loop.

Runs until hitting a breakpoint defined in the input deck, or until the simulation is complete.

Returns one of the state constants defined below.

5.13.2 GEOSX State

`pygeosx.UNINITIALIZED`

`pygeosx.INITIALIZED`

`pygeosx.READY_TO_RUN`

This state indicates that GEOSX still has time steps left to run.

`pygeosx.COMPLETED`

This state indicates that GEOSX has completed the current simulation.

5.13.3 Module Classes

class `pygeosx.Group`

Python interface to `geosx::dataRepository::Group`.

Used to get access to other groups, and ultimately to get wrappers and convert them into Python views of C++ objects.

groups ()

Return a list of the subgroups.

wrappers ()

Return a list of the wrappers.

get_group (*path*)

get_group (*path*, *default*)

Return the `Group` at the relative path *path*; *default* is optional. If no group exists and *default* is not given, raise a `ValueError`; otherwise return *default*.

get_wrapper (*path*)

get_wrapper (*path*, *default*)

Return the `Wrapper` at the relative path *path*; *default* is optional. If no `Wrapper` exists and *default* is not given, raise a `ValueError`; otherwise return *default*.

register (*callback*)

Register a callback on the physics solver.

The callback should take two arguments: the `CRSMatrix` and the array.

Raise `TypeError` if the group is not the `Physics` solver.

class `pygeosx.Wrapper`

Python interface to `geosx::dataRepository::WrapperBase`.

Wraps a generic C++ object. Use `repr` to get a description of the type.

value()

Return a view of the wrapped value, or `None` if it cannot be exported to Python.

A breakdown of the possible return types:

- Instance of a `pylvarray` class If the wrapped type is one of the `LvArray` types that have a Python wrapper type.
- 1D `numpy.ndarray` If the wrapped type is a numeric constant. The returned array is a shallow copy and has a single entry.
- `str` If the wrapped type is a `std::string` this returns a copy of the string.
- list of `str` If the wrapped type is a `LvArray::Array< std::string, 1, ... >` or a `std::vector< std::string >`. This is a copy.
- `None` If the wrapped type is not covered by any of the above.

5.13.4 Segmentation Faults

Improper use of this module and associated programs can easily cause Python to crash. There are two main causes of crashes. Both can be avoided by following some general guidelines.

Stale Numpy Views

The `pylvarray` classes (which may be returned from `Wrapper.value()`) provide various ways to get Numpy views of their data. However, those views are only valid as long as the `LvArray` object's buffer is not reallocated. The buffer may be reallocated by invoking methods (the ones that require the `pylvarray.RESIZEABLE` permission) or by calls into `pygeosx`. It is strongly recommended that you do not keep Numpy views of `LvArray` objects around after calls to `pygeosx`.

```
my_array = pygeosx.get_wrapper("path").value()
view = my_array.to_numpy()
my_array.resize(1000)
print(view) # segfault
```

Destroyed LvArray C++ objects

As mentioned earlier, the classes defined in this module cannot be created in Python; `pygeosx` must create an `LvArray` object in C++, then create a `pylvarray` view of it. However, the Python view will only be valid as long as the underlying `LvArray` C++ object is kept around. If that is destroyed, the Python object will be left holding an invalid pointer and subsequent attempts to use the Python object will cause undefined behavior. Unfortunately, `pygeosx` may destroy `LvArray` objects without warning. It is therefore strongly recommended that you do not keep `pylvarray` objects around after calls to `pygeosx`. The following code snippet, for instance, could segfault:

```
my_array = pygeosx.get_wrapper("path").value()
pygeosx.run()
view = my_array.to_numpy() # segfault
```

5.14 Indices and tables

- `genindex`
- `modindex`

- [search](#)

Welcome to the GEOSX developer guide.

6.1 Contributing

6.1.1 Code style

Introduction

GEOSX is written in standard c++14. In general, target platforms are:

- Linux
- Mac OS X

Currently, our CI/CD system tests on these platforms:

- Ubuntu 18.04, with gcc 8.0 and clang 8.0.0 + cuda10.1.243
- Centos 7.6.1810, with gcc 8.3.1 + cuda10.1.243
- Centos 7.7, with clang 9.0.0
- Mac OS X, with xcode 11.2

Naming Conventions

File Names

- File names should be [PascalCase](#).
- C++ header files are always named with a file extension of *.hpp.
- C++ header implementation files, which contain templated or inline function definitions, are always named *Helpers.hpp.

- C++ source files are always named with a file extension of *.cpp.
- C++ class declarations and definitions are contained files with identical names, except for the extensions.
- C++ free function headers and source files are declared/defined in files with identical names, except for the extension.

For example, a class named “Foo” may be declared in a file named “Foo.hpp”, with inline/templated functions defined in “FooHelpers.hpp”, with the source implementation contained in Foo.cpp.

Warning: There should not be identical filenames that only differ by case. Some filesystems are not case-sensitive, and worse, some filesystems such as MacOSX are case-preserving but not case sensitive.

Function Names

Function and member function names should be `camelCase`.

Variable Names

Variables should be `camelCase`.

Member Names

Member data should be `camelCase` prefix with “m_” (i.e. `double m_dataVariable;`)

Class/Struct Names

Please use `PascalCase` for typenames (i.e. classes)

```
class MyClass;
```

```
class MyClass
{
    double m_doubleDataMember;
    int m_integerDataMember;
}
```

Alias/Typedef Names

Alias and typedefs should be the case of the underlying type that they alias. If no clear format is apparent, as is the case with `double`, then use `camelCase`

Namespace Names

Namespaces names are all lower `camel case`.

Example

One example of would be a for a class named “Foo”, the declaration would be in a header file named “Foo.hpp”

```
/*
 * Foo.hpp
 */

namespace bar
{

class Foo
{
public:
    Foo();
private:
    double m_myDouble;
}
}
```

and a source file named “Foo.cpp”

```
/*
 * Foo.cpp
 */
namespace bar
{
    Foo::Foo() :
        m_myDouble(0.0)
    {
        // some constructor stuff
    }
}
```

Const Keyword

1. All functions and accessors should be declared as “const” functions unless modification to the class is required.
2. In the case of accessors, both a “const” and “non-const” version should be provided.
3. The const keyword should be placed in the location read by the compiler, which is right to left.

The following examples are provided:

```
int a=0; // regular int
int const b = 0; // const int
int * const c = &a; // const pointer to non const int
int const * const d = &b; // const pointer to const int
int & e = a; // reference to int
int const & f = b; // reference to const int
```

Code Format

GEOSX applies a variant of the [BSD/Allman Style](#). Key points to the GEOSX style are:

1. Opening braces (i.e. “{”) go on the next line of any control statement, and are not indented from the control statement.

2. NO TABS. Only spaces. In case it isn't clear ... NO TABS!
3. 2-space indentation

```
for( int i=0 ; i<10 ; ++i )
{
    std::cout << "blah" << std::endl;
}
```

4. Try to stay under 100 character line lengths. To achieve this apply these rules in order
5. Align function declaration/definitions/calls on argument list
6. Break up return type and function definition on new line
7. Break up scope resolution operators

```
void
SolidMechanics_LagrangianFEM::
TimeStepExplicit( real64 const& time_n,
                  real64 const& dt,
                  const int cycleNumber,
                  DomainPartition * const domain )
{
    code here
}
```

As part of the continuous integration testing, this GEOSX code style is enforced via the uncrustify tool. While quite extensive, uncrustify does not enforce every example of the preferred code style. In cases where uncrustify is unable to enforce code style, it will ignore formatting rules. In these cases it is acceptable to proceed with pull requests, as there is no logical recourse.

Header Guards

Header guard names should consist of the name *GEOSX*, followed by the component name (e.g. *dataRepository*), and finally the name of the header file. All characters in the macro should be capitalized.

6.1.2 Git Workflow

The GEOSX project is hosted on github [here](#). For instructions on how to clone and build GEOSX, please refer to the *Quick Start Guide*. Consider consulting <https://try.github.io/> for practical references on how to use git.

Git Credentials

Those who want to contribute to GEOSX should setup SSH keys for authentication, and connect to github through SSH as discussed in [this article](#). Before going further, you should [test your ssh connection](#). If it fails (perhaps because of your institution's proxy), you may consider the [personal access token option](#) as an alternative.

Downloading the Code

Once you have created an `ssh-key` and you have added it to your *Github* account you can download the code through SSH. The following steps clone the repository into `your_geosx_dir`:


```
git clone git@github.com:GEOSX/GEOSX.git your_geosx_dir
cd your_geosx_dir
git lfs install
git submodule init
git submodule update
```

If all goes well, you should have a complete copy of the GEOSX source at this point. The most common errors people encounter here have to do with Github not recognizing their authentication settings.

Branching Model

The branching model used in GEOSX is a modified [Gitflow](#) approach, with some modifications to the merging strategy, and the treatment of release branches, and hotfix branches.

In GEOSX, there are two main branches, `release` and `develop`. The `develop` branch serves as the main branch for the development of new features. The `release` branch serves as the “stable release” branch. The remaining branch types are described in the following subsections.

Note: The early commits in GEOSX (up to version 0.2) used a pure [Gitflow](#) approach for merging feature branches into `develop`. This was done without cleaning the commit history in each feature branch prior to the merge into `develop`, resulting in an overly verbose history. Furthermore, as would be expected, having many active feature branches resulted in a fairly wide (spaghetti) history. At some point in the development process, we chose to switch primarily to a squash-merge approach which results in a linear `develop` history. While this fixes the spaghetti history, we do potentially lose important commit history during the development process. Options for merging are discussed in the following sections.

Feature Branches

New developments (new features or modifications to features) are branched off of `develop` into a feature branch. The naming of feature branches should follow `feature/[developer]/[branch-description]` if you expect that only a single developer will contribute to the branch, or `feature/[branch-description]` if you expect it will be a collaborative effort. For example, if a developer named `neo` were to add or modify a code feature expecting that they would be the only contributor, they would create a branch using the following commands to create the local branch and push it to the remote repository:

```
git checkout -b feature/neo/freeYourMind
git push -u origin feature/neo/freeYourMind
```

However if the branch is a collaborative branch amongst many developers, the appropriate commands would be:

```
git checkout -b feature/freeYourMind
git push -u origin feature/freeYourMind
```

When feature branches are ready to be merged into `develop`, a Pull Request should be created to perform the review and merging process.

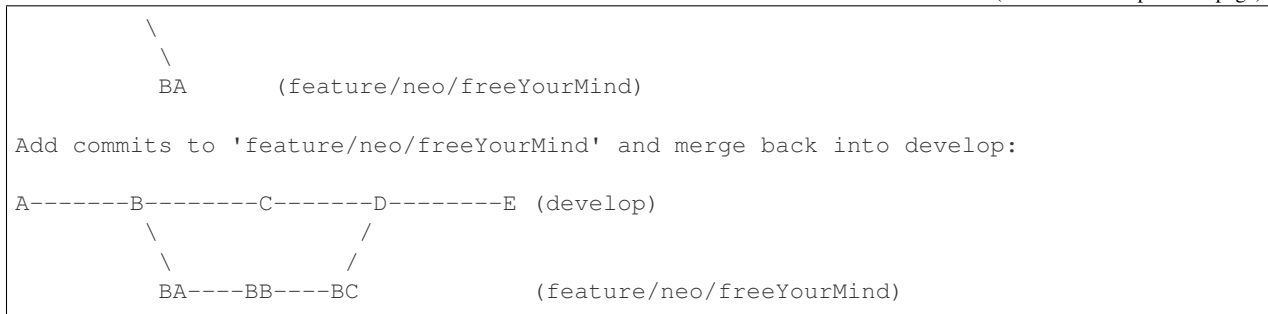
An example lifecycle diagram for a feature branch:

```
create new feature branch:
git checkout -b feature/neo/freeYourMind

A-----B-----C (develop)
```

(continues on next page)

(continued from previous page)



See below for details about [Submitting a Pull Request](#).

Bugfix Branches


Bugfix branches are used to fix bugs that are present in the `develop` branch. A similar naming convention to that of the `feature` branches is used, replacing “feature” with “bugfix” (i.e. `bugfix/neo/squashAgentSmith`). Typically, bugfix branches are completed by a single contributor, but just as with the `feature` branches, a collaborative effort may be required resulting in dropping the developer name from the branch name.

When bugfix branches are ready to be merged into `develop`, a Pull Request should be created to perform the review and merging process. See below for details about [Submitting a Pull Request](#).

Release Candidate Branches

When `develop` has progressed to a point where we would like to create a new release, we will create a release candidate branch with the name consisting of `release_major.minor.x` number, where the `x` represents the sequence of patch tags that will be applied to the branch. For instance if we were releasing version `1.2.0`, we would name the branch `release_1.2.x`. Once the release candidate is ready, it is merged back into `develop`. Then the `develop` branch is merged into the `release` branch and tagged. From that point the `release` branch exists to provide a basis for maintaining a stable release version of the code. Note that the absence of hotfix branches, the history for `release` and `develop` would be identical.


An example lifecycle diagram for a release candidate branch:



Hotfix Branches

A hotfix branch fixes a bug in the `release` branch. It uses the same naming convention as a bugfix branch. The main difference with a bugfix branch is that the primary target branch is the `release` branch instead of `develop`. As a soft policy, merging a hotfix into a release branch should result in a patch increment for the release sequence of tags. So if a hotfix was merged into `release` with a most recent tag of `1.2.1`, the merged commit would be tagged with `1.2.2`. Finally, at some point prior to the next major/minor release, the `release` branch should be merged back into `develop` to incorporate any hotfix changes into `develop`.

An example lifecycle diagram for hotfix branches:



Documentation Branches

A docs branch is focused on writing and improving the documentation for GEOSX. The use of the docs branch name root applies to both sphinx documentation and doxygen documentation. The docs branch follows the same naming conventions as described in the [Feature Branches](#) section. The html produced by a documentation branch should be proofread using sphinx/doxygen prior to merging into develop.

Keeping Your Branch Current

Over the course of a long development effort in a single feature branch, a developer may need to either merge develop into their feature branch, or rebase their feature branch on develop. We do not have a mandate on how you keep your branch current, but we do have guidelines on the branch history when merging your branch into develop. Typically, merging develop into your branch is the easiest approach, but will lead to a complex relationship with develop with multiple interactions... which can lead to a confusing history. Conversely, rebasing your branch onto develop is more difficult, but will lead to a linear history within the branch. For a complex history, we will perform a squash merge into develop, thereby the work from the branch will appear as a single commit in develop. For clean branch histories where the individual commits are meaningful and should be preserved, we have the option to perform a merge commit in with the PR is merged into develop, with the addition of a merge commit, thus maintaining the commit history.

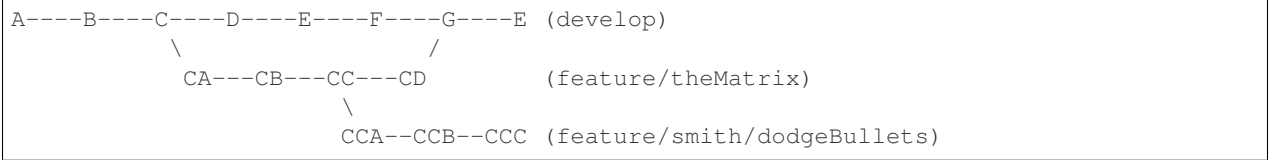
Branching off of a Branch

During the development processes, sometimes it is appropriate to create a branch off of a branch. For instance, if there is a large collaborative development effort on the branch feature/theMatrix, and a developer would like to add a self-contained and easily reviewable contribution to that effort, he/she should create a branch as follows:

```
git checkout feature/theMatrix
git checkout -b feature/smith/dodgeBullets
git push -u origin feature/smith/dodgeBullets
```

If feature/smith/dodgeBullets is intended to be merged into feature/theMatrix, and the commit history of feature/theMatrix is not changed via git rebase, then the process of merging the changes back into feature/theMatrix is fairly standard.

However, if feature/theMatrix is merged into develop via a squash merge, and then smith would like to merge feature/smith/dodgeBullets into develop, there is a substantial problem due to the diverged history of the branches. Specifically, feature/smith/dodgeBullets branched off a commit in feature/theMatrix that does not exist in develop (because it was squash-merged). For simplicity, let us assume that the commit hash that feature/smith/dodgeBullets originated from is CC, and that there were commits CA, CB, CC, CD in feature/theMatrix. When feature/theMatrix was squash-merged, all of the changes appear in develop as commit G. To further complicate the situation, perhaps a complex PR was merged after G, resulting in E on develop. The situation is illustrated by:




In order to successfully merge `feature/smith/dodgeBullets` into `develop`, all commits present in `feature/smith/dodgeBullets` after `CC` must be included, while discarding `CA`, `CB`, which exist in `feature/smith/dodgeBullets` as part of its history, but not in `develop`.

One “solution” is to perform a `git rebase --onto` of `feature/smith/dodgeBullets` onto `develop`. Specifically, we would like to rebase `CCA`, `CCB`, `CCC` onto `G`, and proceed with our development of `feature/smith/dodgeBullets`. This would look like:

```

git checkout develop
git pull
git checkout feature/smith/dodgeBullets
git rebase -onto G CC
  
```

As should be apparent, we have specified the starting point as `G`, and the point at which we replay the commits in `feature/smith/dodgeBullets` as all commits AFTER `CC`. The result is:



Now you may proceed with standard methods for keeping `feature/smith/dodgeBullets` current with `develop`.

Submitting a Pull Request

Once you have created your branch and pushed changes to Github, you can create a [Pull Request](#) on Github. The PR creates a central place to review and discuss the ongoing work on the branch. Creating a pull request early in the development process is preferred as it allows for developers to collaborate on the branch more readily.

Note: When initially creating a pull request (PR) on GitHub, always create it as a *draft* PR while work is ongoing and the PR is not ready for testing, review, and merge consideration.

When you create the initial draft PR, please ensure that you apply appropriate labels. Applying labels allows other developers to more quickly filter the live PRs and access those that are relevant to them. Always add the *new* label upon PR creation, as well as to the appropriate *type*, *priority*, and *effort* labels. In addition, please also add any appropriate *flags*.

Note: If your branch and PR will resolve any open issues, be sure to *link* them to the PR to ensure they are appropriately resolved once the PR is merged. In order to *link* the issue to the PR for automatic resolution, you must use one of the keywords followed by the issue number (e.g. resolves #1020) in either the main description of the PR, or a commit message. Entries in PR comments that are not the main description or a commit message will be ignored, and the issue will not be automatically closed. A complete list of keywords are:

- close
- closes
- closed

- fix
- fixes
- fixed
- resolve
- resolves
- resolved

For more details, see the [Github Documentation](#).

Once you are satisfied with your work on the branch, you may promote the PR out of draft status, which will allow our integrated testing suite to execute on the PR branch to ensure all tests are passing prior to merging.

Once the tests are passing – or in some cases immediately – add the *flag: ready for review* label to the PR, and be sure to tag any relevant developers to review the PR. The PR *must* be approved by reviewers in order to be merged.

Note that whenever a pull request is merged into `develop`, commits are either *squashed*, or preserved depending on the cleanliness of the history.

Keeping Submodules Current

Whenever you switch between branches locally, pull changes from `origin` and/or merge from the relevant branches, it is important to update the submodules to move the head to the proper commit.

```
git submodule update --recursive
```

You may also wish to modify your *git pull* behavior to update your submodules recursively for you in one command, though you forfeit some control granularity to do so. The method for accomplishing this varies between git versions, but as of git 2.15 you should be able to globally configure git to accomplish this via:

```
git config --global submodule.recurse true
```

In some cases, code changes will require to rebaseline the `Integrated Tests`. If that is the case, you will need to modify the `integrated tests submodule`. Instructions on how to modify a submodule are presented in the following section.

Working on the Submodules

Sometimes it may be necessary to modify one of the submodules. In order to do so, you need to create a pull request on the submodule repository. The following steps can be followed in order to do so.

Move to the folder of the `submodule` that you intend to modify.

```
cd submodule-folder
```

Currently the `submodule` is in detached head mode, so you first need to move to the main branch (either `develop` or `master`) on the submodule repository, pull the latest changes, and then create a new branch.

```
git checkout <main-branch>
git pull
git checkout -b <branch-name>
```

You can perform some work on this branch, *add* and *commit* the changes and then push the newly created branch to the submodule repository on which you can eventually create a pull request using the same process discussed above in *Submitting a Pull Request*.

```
git push --set-upstream origin <branch-name>
```

Resolving Submodule Changes in Primary Branch PRs

When you conduct work on a submodule during work on a primary GEOSX branch with an open PR, the merging procedure requires that the submodule referenced by the GEOSX PR branch be consistent with the submodule in the main branch of the project. This is checked and enforced via TravisCI.

Thus, in order to merge a PR that includes modifications to submodules, the various PRs for each repository should be staged and finalized, to the point they are all ready to be merged, with higher-level PRs in the merge hierarchy having the correct submodule references for the current main branch for their repository.

Starting from the bottom of the submodule hierarchy, the PRs are resolved, after which the higher-level PRs with reference to a resolved PR must update their submodule references to point to the new main branch of the submodule with the just-resolved PR merged. After any required automated tests pass, the higher-level PRs can then be merged.

The name of the main branch of each submodule is presented in the table below.

Submodule	Main branch
blt	develop
LvArray	develop
integratedTests	develop
hdf5_interface	master
PAMELA	master
PVTPackage	master

6.1.3 Using Docker

6.1.4 Sphinx Documentation

Generating the documentation

- To generate the documentation files, you will need to install Sphinx using

```
sudo apt install python-sphinx
```

Then you can generate the documentation files with the following command

```
cd GEOSX/build-your-platform-release
make geosx_docs
```

- That will create a new folder

```
GEOSX/build-your-platform-release/html/docs/sphinx
```

which contains all the html files generated.

Documenting the code

The documentation is generated from restructured text files (`.rst`). Most files can be found in `src/docs/sphinx`. Files which are specific to parts of the code, like those describing a specific class, can instead be found in `docs` subdirectory in the folder containing the source code.

Information about how to write `rst` files can be found [here](#).

6.1.5 Doxygen Documentation

Developer documentation of code is provided in the form of Doxygen-style comment blocks. [Doxygen](#) is a tool for generating html/xml/latex documentation for C++ code from specially marked code comments. Having concise but high quality documentation of public APIs helps both users and developers of these APIs. We use Doxygen and Ctest to enforce documentation coverage. If Doxygen produces any warnings, your pull request will fail CI checks! See [Git Workflow](#) for more on pull requests and CI.

Accessing

There are two ways to access Doxygen documentation.

Build locally

Prior to configuring a GEOSX build, have Doxygen installed:

```
sudo apt install doxygen
```

Note: Eventually, doxygen (version 1.8.13) is provided within the *thirdPartyLibs* repository.

Configure GEOSX and go the build directory:

```
cd GEOSX/build-your-platform-release
```

Build doxygen docs only:

```
make geosx_doxygen
```

Or build all docs:

```
make geosx_docs
```

Open in browser:

```
google-chrome html/doxygen_output/html/index.html
```

On readthedocs

Go to [GEOSX documentation](#), select the version of interest, and follow the Doxygen link at the left-hand-side.

Guidelines

What to document

The following entities declared in project header files within *geosx* namespace require documentation:

- all classes and structs, including public nested ones
- global functions, variables and type aliases
- public and protected member functions, variables and type aliases in classes
- preprocessor macros

Exceptions are made for:

- overrides of virtual functions in derived types
- implementation details nested in namespace *internal*
- template specializations in some cases

How to document

The following rules and conventions are used. Some are stricter than others.

1. We use `@`-syntax for all Doxygen commands (e.g. `@brief` instead of `\brief`).
2. Entities such as type aliases and member variables that typically only require a brief description, can have a single-line documentation starting with `///`.
 - `@brief` is not required for single-line comments.
3. Entities such as classes and functions that typically require either detailed explanation or parameter documentation, are documented with multiline comment blocks.
 - `@brief` is required for comment blocks.
4. Brief and detailed descriptions should be complete sentences (i.e. start with a capital letter and end with a dot).
5. Prefer concise wording in `@brief`, e.g. “Does X.” instead of “This is a function that does X.”
6. All functions parameters and return values must be explicitly documented via `@param` and `@return`.
 - An exception to this rule seem to be copy/move constructor/assignment, where parameter documentation can be omitted.
7. Add `[in]` and `[out]` tags to function parameters, as appropriate.
8. Function and template parameter descriptions are not full sentences (i.e. not capitalized nor end with a dot).
9. For hierarchies with virtual inheritance, document base virtual interfaces rather than overriding implementations.
10. Documented functions cannot use `GEOSX_UNUSED_ARG()` in their declarations.
11. For empty virtual base implementations that use `GEOSX_UNUSED_ARG(x)` to remove compiler warnings, use one of two options:
 - move empty definition away (e.g. out of class body) and keep `GEOSX_UNUSED_ARG(x)` in definition only;
 - put `GEOSX_UNUSED_VAR(x)` into the inline empty body.
12. For large classes, logically group functions using member groups via `///@{` and `///@}` and give them group names and descriptions (if needed) via a `@name` comment block. Typical groups may include:

- constructors/destructor/assignment operators;
 - getter/setter type functions;
 - overridable virtual functions;
 - any other logically coherent groups (functions related to the same aspect of class behavior).
13. In-header implementation details (e.g. template helpers) often shouldn't appear in user documentation. Wrap these into *internal* namespace.
 14. Use `/// @cond DO_NOT_DOCUMENT` and `/// @endcond` tags to denote a section of public API that should not be documented for some reason. This should be used rarely and selectively. An example is in-class helper structs that must be public but that user should not refer to explicitly.

Example

```

/// This is a documented macro
#define USEFUL_MACRO

/**
 * @brief Short description.
 * @tparam T type of input value
 * @param[in] x input value explanation
 * @return return value explanation
 *
 * Detailed description goes here.
 *
 * @note A note warning users of something unexpected.
 */
template<typename T>
int Foo( T const & x );

/**
 * @brief Class for showing Doxygen.
 * @tparam T type of value the class operates on
 *
 * This class does nothing useful except show how to use Doxygen.
 */
template<typename T>
class Bar
{
public:

    /// A documented member type alias.
    using size_type = typename std::vector<T>::size_type;

    /**
     * @name Constructors/destructors.
     */
    //@{

    /**
     * @brief A documented constructor.
     * @param value to initialize the object
     */
    explicit Bar( T t );

```

(continues on next page)

(continued from previous page)

```

/**
 * @brief A deleted, but still documented copy constructor.
 * @param an optionally documented parameter
 */
Bar( Bar const & source ) = delete;

/**
 * @brief A defaulted, but still documented move constructor.
 * @param an optionally documented parameter
 */
Bar( Bar const & source ) = default;

/**
 * @brief A documented desctructor.
 * virtual ~Bar() = default;
 */

///<}

/**
 * @name Getters for stored value.
 */
///<{

/**
 * @brief A documented public member function.
 * @return a reference to contained value
 */
T & getValue();

/**
 * @copydoc getValue()
 */
T const & getValue() const;

///<}

protected:

/**
 * @brief A documented protected pure virtual function.
 * @param[in] x the input value
 * @param[out] y the output value
 *
 * Some detailed explanation for users and implementers.
 */
virtual void doSomethingOverridable( int const x, T & y ) = 0;

/// @cond DO_NOT_DOCUMENT
// Some stuff we don't want showing up in Doxygen
struct BarHelper
{
};
/// @endcond

private:

/// An optionally documented (not enforced) private member.

```

(continues on next page)

(continued from previous page)

```

    T m_value;

};

```

Current Doxygen

[Link to Doxygen](#)

6.1.6 Unit Testing

Unit testing is integral to the GEOSX development process. While not all components naturally lend themselves to unit testing (for example a physics solver) every effort should be made to write comprehensive quality unit tests.

Each sub-directory in `coreComponents` should have a `unitTests` directory containing the test sources. Each test consists of a `cpp` file whose name begins with `test` followed by a name to describe the test. Please read over the `LvArray` unit test documentation as it gives an intro to the Google Test framework and a set of best practices.

GEOSX Specific Recommendations

An informative example is `testSinglePhaseBaseKernels` which tests the single phase flow mobility and accumulation kernels on a variety of inputs.

```

TEST( SinglePhaseBaseKernels, mobility )
{
    int constexpr NTEST = 3;

    real64 const dens[NTEST]      = { 800.0, 1000.0, 1500.0 };
    real64 const dDens_dPres[NTEST] = { 1e-5, 1e-10, 0.0 };
    real64 const visc[NTEST]      = { 5.0, 2.0, 1.0 };
    real64 const dVisc_dPres[NTEST] = { 1e-7, 0.0, 0.0 };

    for( int i = 0; i < NTEST; ++i )
    {
        SCOPED_TRACE( "Input # " + std::to_string( i ) );

        real64 mob;
        real64 dMob_dPres;

        MobilityKernel::compute( dens[i], dDens_dPres[i], visc[i], dVisc_dPres[i], mob,
↪dMob_dPres );

        // compute etalon
        real64 const mob_et = dens[i] / visc[i];
        real64 const dMob_dPres_et = mob_et * (dDens_dPres[i] / dens[i] - dVisc_dPres[i] /
↪ visc[i]);

        EXPECT_DOUBLE_EQ( mob, mob_et );
        EXPECT_DOUBLE_EQ( dMob_dPres, dMob_dPres_et );
    }
}

```

[Source: `coreComponents/physicsSolvers/fluidFlow/unitTests/testSinglePhaseBaseKernels.cpp`]

What makes this such a good test is that it depends on very little other than kernels themselves. There is no need to involve the data repository or parse an XML file. Sometimes however this is not possible, or at least not without a significant duplication of code. In this case it is better to embed the XML file into the test source as a string instead of creating a separate XML file and passing it to the test as a command line argument or hard coding the path. One example of this is `testLaplaceFEM` which tests the laplacian solver. The embedded XML is shown below.

```
char const * xmlInput =
"<Problem>\n"
"  <Solvers gravityVector=\"{ 0.0, 0.0, -9.81 }\">\n"
"    <CompositionalMultiphaseFVM name=\"compflow\">\n"
"      logLevel=\"0\">\n"
"        discretization=\"fluidTPFA\">\n"
"          targetRegions=\"{Region2}\">\n"
"            fluidNames=\"{fluid1}\">\n"
"              solidNames=\"{rock}\">\n"
"                permeabilityNames=\"{rockPerm}\">\n"
"                  relPermNames=\"{relperm}\">\n"
"                    capPressureNames=\"{cappressure}\">\n"
"                      temperature=\"297.15\">\n"
"                        useMass=\"1\">\n"
"                          \n"
"        <NonlinearSolverParameters newtonTol=\"1.0e-6\">\n"
"          newtonMaxIter=\"2\">\n"
"        <LinearSolverParameters solverType=\"gmres\">\n"
"          krylovTol=\"1.0e-10\">\n"
"        </CompositionalMultiphaseFVM>\n"
"      </Solvers>\n"
"    <Mesh>\n"
"      <InternalMesh name=\"mesh1\">\n"
"        elementTypes=\"{C3D8}\">\n"
"          xCoords=\"{0, 3}\">\n"
"            yCoords=\"{0, 1}\">\n"
"              zCoords=\"{0, 1}\">\n"
"                nx=\"{3}\">\n"
"                  ny=\"{1}\">\n"
"                    nz=\"{1}\">\n"
"                      cellBlockNames=\"{cb1}\">\n"
"        </InternalMesh>\n"
"      </Mesh>\n"
"    <Geometry>\n"
"      <Box name=\"source\" xMin=\"{-0.01, -0.01, -0.01}\" xMax=\"{ 1.01, 1.01, 1.01}\">\n"
"    </Box>\n"
"      <Box name=\"sink\" xMin=\"{ 1.99, -0.01, -0.01}\" xMax=\"{ 3.01, 1.01, 1.01}\">\n"
"    </Box>\n"
"      </Geometry>\n"
"    <NumericalMethods>\n"
"      <FiniteVolume>\n"
"        <TwoPointFluxApproximation name=\"fluidTPFA\">\n"
"          fieldName=\"pressure\">\n"
"            coefficientName=\"permeability\">\n"
"              coefficientModelNames=\"{rockPerm}\">\n"
"        </TwoPointFluxApproximation>\n"
"      </FiniteVolume>\n"
"    </NumericalMethods>\n"
"    <ElementRegions>\n"
"      <CellElementRegion name=\"Region2\" cellBlocks=\"{cb1}\" materialList=\"\n"
"    {fluid1, rock, relperm, cappressure, rockPerm, rockPorosity, nullSolid}\">\n"
"    </CellElementRegion>\n"
"    </ElementRegions>\n"
"    <Constitutive>\n"
"      <CompositionalMultiphaseFluid name=\"fluid1\">
```

(continues on next page)

(continued from previous page)

```

"           phaseNames="{oil, gas}"\n"
"           equationsOfState="{PR, PR}"\n"
"           componentNames="{N2, C10, C20, H2O}"\n"
"           componentCriticalPressure="{34e5, 25.3e5, 14.
↪6e5, 220.5e5}"\n"
"           componentCriticalTemperature="{126.2, 622.0,
↪782.0, 647.0}"\n"
"           componentAcentricFactor="{0.04, 0.443, 0.816, 0.
↪344}"\n"
"           componentMolarWeight="{28e-3, 134e-3, 275e-3,
↪18e-3}"\n"
"           componentVolumeShift="{0, 0, 0, 0}"\n"
"           componentBinaryCoeff="{ {0, 0, 0, 0},\n"
"                                   {0, 0, 0, 0},\n"
"                                   {0, 0, 0, 0},\n"
"                                   {0, 0, 0, 0} }"/>\n"
"       <CompressibleSolidConstantPermeability name="rock"\n"
"           solidModelName="nullSolid"\n"
"           porosityModelName="rockPorosity"\n"
"           permeabilityModelName="rockPerm"/>\n"
"       <NullModel name="nullSolid"/> \n"
"       <PressurePorosity name="rockPorosity"\n"
"           defaultReferencePorosity="0.05"\n"
"           referencePressure = "0.0"\n"
"           compressibility="1.0e-9"/>\n"
"       <BrooksCoreyRelativePermeability name="relperm"\n"
"           phaseNames="{oil, gas}"\n"
"           phaseMinVolumeFraction="{0.1, 0.15}"\n"
"           phaseRelPermExponent="{2.0, 2.0}"\n"
"           phaseRelPermMaxValue="{0.8, 0.9}"/>\n"
"       <BrooksCoreyCapillaryPressure name="cappressure"\n"
"           phaseNames="{oil, gas}"\n"
"           phaseMinVolumeFraction="{0.2, 0.05}"\n"
"           phaseCapPressureExponentInv="{4.25, 3.5}"\n"
"           phaseEntryPressure="{0., 1e8}"\n"
"           capPressureEpsilon="0.0"/> \n"
"       <ConstantPermeability name="rockPerm"\n"
"           permeabilityComponents="{2.0e-16, 2.0e-16, 2.0e-16}"/> \n
↪"
"   </Constitutive>\n"
"   <FieldSpecifications>\n"
"       <FieldSpecification name="initialPressure"\n"
"           initialCondition="1"\n"
"           setNames="{all}"\n"
"           objectPath="ElementRegions/Region2/cb1"\n"
"           fieldName="pressure"\n"
"           functionName="initialPressureFunc"\n"
"           scale="5e6"/>\n"
"       <FieldSpecification name="initialComposition_N2"\n"
"           initialCondition="1"\n"
"           setNames="{all}"\n"
"           objectPath="ElementRegions/Region2/cb1"\n"
"           fieldName="globalCompFraction"\n"
"           component="0"\n"
"           scale="0.099"/>\n"
"       <FieldSpecification name="initialComposition_C10"\n"
"           initialCondition="1"\n"

```

(continues on next page)

(continued from previous page)

```

"           setNames="{all}"\n"
"           objectPath="ElementRegions/Region2/cb1"\n"
"           fieldName="globalCompFraction"\n"
"           component="1"\n"
"           scale="0.3"/>\n"
"       <FieldSpecification name="initialComposition_C20"\n"
"           initialCondition="1"\n"
"           setNames="{all}"\n"
"           objectPath="ElementRegions/Region2/cb1"\n"
"           fieldName="globalCompFraction"\n"
"           component="2"\n"
"           scale="0.6"/>\n"
"       <FieldSpecification name="initialComposition_H2O"\n"
"           initialCondition="1"\n"
"           setNames="{all}"\n"
"           objectPath="ElementRegions/Region2/cb1"\n"
"           fieldName="globalCompFraction"\n"
"           component="3"\n"
"           scale="0.001"/>\n"
"   </FieldSpecifications>\n"
"   <Functions>\n"
"       <TableFunction name="initialPressureFunc"\n"
"           inputVarNames="{elementCenter}"\n"
"           coordinates="{0.0, 3.0}"\n"
"           values="{1.0, 0.5}">\n"
"   </Functions>"
"</Problem>";

```

[Source: *coreComponents/physicsSolvers/fluidFlow/unitTests/testCompMultiphaseFlow.cpp*]

MPI

Often times it makes sense to write a unit test that is meant to be run with multiple MPI ranks. This can be accomplished by simply adding the `NUM_MPI_TASKS` parameter to `blt_add_test` in the CMake file. For example

```

blt_add_test( NAME testWithMPI
              COMMAND testWithMPI
              NUM_MPI_TASKS ${NUMBER_OF_MPI_TASKS} )

```

With this addition make test or calling `ctest` directly will run `testWithMPI` via something analogous to `mpirun -n NUMBER_OF_MPI_TASKS testWithMPI`.

6.1.7 Integrated Tests

About

integratedTests is a submodule of *GEOSX* residing at the top level of the directory structure. It will run *GEOSX* with various *.xml* files and compare the output against a baseline.

Note: you may need to install `h5py` and `mpi4py`.

Structure

The *integratedTests* directory is composed of two main directories, *integratedTests/geosxats* and *integratedTests/update/run*. The *integratedTests/geosxats* directory contains all the machinery involved in running the tests including the main executable *integratedTests/geosxats/geosxats*. The *integratedTests/update/run* directory contains all the actual tests themselves, including the *.xml* files and the baselines.

```
- integratedTests/
  - geosxats/
    - geosxats
  - update/
    - run/
      - sedov/
      - beamBending/
        - baselines/
          - beamBending/
            - <baseline-files>
        - beamBending.ats
      - beamBending.xml
```

Arguments

The program takes a number of arguments, the most important ones are

- The path to the *GEOSX* binary directory (*<build-dir>/bin*)
- *--workingDir WORKINGDIR* which sets the working (test) directory.
- *-N NUMNODES* specifies the number of nodes to use the default is the minimum number that the tests require.
- *-a {veryclean, rebaseline, ...}* specify a specific action, *veryclean* deletes all output files, and *rebaseline* lets you rebaseline all or some of the tests.
- *-h* prints out the *geosxats* help.

How to Run the Tests

To run all the tests from the top level *GEOSX* directory you would do

```
integratedTests/geosxats/geosxats <build-path>/bin --workingDir integratedTests/
↪update/run
```

To run only the *sedov* tests you would do

```
integratedTests/geosxats/geosxats <build-path>/bin --workingDir integratedTests/
↪update/run/sedov
```

However if you want to run all the tests there's an easier way. In the build directory there's a symbolic link to *integratedTests/update/run* called *integratedTests* and a bash script *geosxats.sh* that wraps *integratedTests/geosxats/geosxats* and passes it the path to the binary directory and the *integratedTests/update/run* directory along with any other command line arguments. To run this script do the following

```
cd <build-path>
./geosxats.sh
```

If the script or symbolic link is not present you will need to reconfigure by running the *config-build.py* script.

When the program has finished running you will see something like this

```

FAIL RUN : 0

UNEXPECTEDPASS : 0

FAIL RUN (OPTIONAL STEP) : 0

FAIL CHECK : 1
( beamBending )

FAIL CHECK (MINOR) : 0

TIMEOUT : 0

NOT RUN : 0

INPROGRESS : 0

FILTERED : 0

RUNNING : 0

PASSED : 3
( sedov 2D_100x100_incompr_linear 2D_100x100_incompr_linear_faceBC )

EXPECTEDFAIL : 0

SKIPPED : 0

BATCHED : 0

NOT BUILT : 0

TOTAL TIME : 0:01:13
TOTAL PROCESSOR-TIME : 0:02:13
AVAIL PROCESSOR-TIME : 1:28:26
RESOURCE UTILIZATION : 2.51%

LONGEST RUNNING TESTS:
0:01:13 sedov
0:00:23 2D_100x100_incompr_linear_faceBC
0:00:22 2D_100x100_incompr_linear
0:00:14 beamBending

Status      : TestCase      : Directory      : Elapsed : Resources :
↪TestStep
----- : ----- : ----- : ----- : ----- : -
↪-----
FAIL CHECK : beamBending : beamBending/beamBending : 0:00:14 : 0:00:14 :
↪restartcheck
----- : ----- : ----- : ----- : ----- : -
↪-----

Generating HTML documentation files (running 'atddoc')...
Failed to create HTML documentation in /g/g14/corbett5/geosx/mirror/integratedTests/
↪update/doc

Undocumented test problems:
beamBending 2D_100x100_incompr_linear 2D_100x100_incompr_linear_faceBC sedov

```


Ignore the error regarding the failure to create a HTML documentation file and the warning about the undocumented test problems. The only important thing is if any of the tests aren't in the *PASSED* category. For a nice summary of the results open the *test_results.html* file in the *geosxats* working directory.

When running the tests multiple times in a row, only tests that failed to pass will run. If you would like to run all the tests again call

```
./geosxats.sh -a veryclean
```

which will delete all the generated files. Furthermore these generated files are not ignored by *git*, so until you run *veryclean* the *integratedTests* repo will register changes.

Note: On some development machines *geosxats* won't run parallel tests by default (e.g. on an linux laptop or workstation), and as a result many tests will be skipped. We highly recommend running tests on an MPI-aware platform.

Output Created By a Test

Since the *beamBending* test failed let's look at it's output. The output for the *beamBending* test is stored in *integratedTests/update/run/beamBending/beamBending* directory. In addition to any files *GEOSX* itself creates you will find

- *beamBending.data* which holds all of the standard output of the various steps.
- *beamBending.err* which holds all of the standard error output of the various steps.
- *beamBending.geosx.out* which holds all of the standard output for only the *geosx* step.
- *beamBending_restart_000000003.restartcheck* which holds all of the standard output for only the *restartcheck* step.
- *beamBending_restart_000000003_diff.hdf5* which mimicks the hierarchy of the restart file and has links to the differing data datasets.

The RestartCheck File

Currently the only manner of check that we support is a restart check, this check compares a restart file output at the end of a run against a baseline. The program that does the diff is *integratedTests/geosxats/helpers/restartcheck.py*. The program compares the two restart files and writes out a *.restart_check* file with the results, as well as exiting with an error code if the files compare differently.

This program takes a number of arguments and they are as follows

- Regex specifying the restart file. If the regex matches multiple files the one with the greater string is selected. For example *restart_100.hdf5* wins out over *restart_088.hdf5*.
- Regex specifying the baseline file.
- *-r* The relative tolerance for floating point comparison, the default is 0.0.
- *-a* The absolute tolerance for floating point comparison, the default is 0.0.
- *-e* A list of regex expressions that match paths in the restart file tree to exclude from comparison. The default is *[/commandLine]*.
- *-w* Force warnings to be treated as errors, default is false.
- *-s* Suppress output to stdout, default is False.

The *.restart_check* file itself starts off with a summary of the arguments. The program then compares the *.root* files and if they are similar proceeds to compare all the *.hdf5* data files.

If the program encounters any differences it will spit out an error message. An error message for scalar values looks as follows

```
Error: /datagroup_0000000/sidre/external/ProblemManager/domain/ConstitutiveManager/
↳shale/YoungsModulus
  Scalar values of types float64 and float64 differ: 22500000000.0, 10000022399.9.
```

Where the first value is the value in the test's restart file and the second is the value in the baseline.

An example of an error message for arrays is

```
Error: /datagroup_0000000/sidre/external/ProblemManager/domain/MeshBodies/mesh1/
↳Level0/nodeManager/TotalDisplacement
  Arrays of types float64 and float64 have 1836 values of which 1200 have differing_
↳values.
  Statistics of the differences greater than 0:
    max_index = (1834,), max = 2.47390764755, mean = 0.514503482629, std = 0.
↳7021288881
```

This means that the max absolute difference is 2.47 which occurs at value 1834. Of the values that are not equal the mean absolute difference is 0.514 and the standard deviation of the absolute difference is 0.702.

When the tolerances are non zero the comparison is a bit more complicated. From the *FileComparison.compareFloatArrays* method documentation

```
Entries x1 and x2 are considered equal iff
  |x1 - x2| <= ATOL or |x1 - x2| <= RTOL * |x2|.
To measure the degree of difference a scaling factor q is introduced. The goal is now_
↳to minimize q such that
  |x1 - x2| <= ATOL * q or |x1 - x2| <= RTOL * |x2| * q.
If RTOL * |x2| > ATOL
  q = |x1 - x2| / (RTOL * |x2|)
else
  q = |x1 - x2| / ATOL.
If the maximum value of q over all the entries is greater than 1.0 then the arrays_
↳are considered different and an error message is produced.
```

An sample error message is

```
Error: /datagroup_0000000/sidre/external/ProblemManager/domain/MeshBodies/mesh1/
↳Level0/nodeManager/TotalDisplacement
  Arrays of types float64 and float64 have 1836 values of which 1200 fail both the_
↳relative and absolute tests.
  Max absolute difference is at index (1834,): value = 2.07474948094, base_value =_
↳4.54865712848
  Max relative difference is at index (67,): value = 0.00215842135281, base_value =_
↳0.00591771127792
  Statistics of the q values greater than 1.0 defined by the absolute tolerance: N =_
↳1200
    max = 16492717650.3, mean = 3430023217.52, std = 4680859258.74
  Statistics of the q values greater than 1.0 defined by the relative tolerance: N = 0
```

The restart check step can be run in parallel using mpi via

```
mpirun -n NUM_PROCESSES python -m mpi4py restartcheck.py ...
```

In this case rank zero reads in the restart root file and then each rank parses a subset of the data files creating a *.\$RANK.restartcheck* file. Rank zero then merges the output from each of these files into the main *.restartcheck* file and prints it to standard output.

The `.diff.hdf5` File

Each error generated in the *restartcheck* step creates a group with three children in the `_diff.hdf5` file. For example the error given above will generate a hdf5 group

```
/FILENAME/datagroup_0000000/sidre/external/ProblemManager/domain/MeshBodies/mesh1/  
↪Level0/nodeManager/TotalDisplacement
```

with datasets *baseline*, *run* and *message* where *FILENAME* is the name of the restart data file being compared. The *message* dataset contains a copy of the error message while *baseline* is a symbolic link to the baseline dataset and *run* is a symbolic link to the dataset generated by the run. This allows for easy access to the raw data underlying the diff without data duplication. For example if you want to extract the datasets into python you could do this:

```
import h5py
file_path = "beamBending_restart_000000003_diff.hdf5"
path_to_data = "/beamBending_restart_000000011_00000000.hdf5/datagroup_0000000/sidre/  
↪external/ProblemManager/domain/MeshBodies/mesh1/Level0/nodeManager/TotalDisplacement  
↪"
f = h5py.File("file_path", "r")
error_message = f["path_to_data/message"]
run_data = f["path_to_data/run"][:]
baseline_data = f["path_to_data/baseline"][:]

# Now run_data and baseline_data are numpy arrays that you may use as you see fit.
rtol = 1e-10
atol = 1e-15
absolute_diff = np.abs(run_data - baseline_data) < atol
hybrid_diff = np.close(run_data, baseline_data, rtol, atol)
```

When run in parallel each rank creates a `.$RANK.diff.hdf5` file which contains the diff of each data file processed by that rank.

The `.ats` File

The `.ats` file is a python script that describes the *TestCases* to run and steps for each *TestCase*. Each `.ats` file needs to have at least one *TestCase* and each *TestCase* needs to have at least one step.

A simple example is the *beamBending.ats* file

```
TestCase(  
    name = "beamBending",  
    desc = "Tests beam bending.",  
    label = "auto",  
    owner = "Ben Corbett",  
    independent = True,  
    steps = (geosx(deck="beamBending.xml"),)
```

This creates a *TestCase* called *beamBending* with a single step that runs *GEOSX* with the *beamBending.xml* input file, a *restartcheck* step automatically follows each *geosx* step. So this file describes a test that runs the *beamBending* problem and compares the restart file against the baseline.

A slightly more complicated example is the *singlePhaseFlow.ats* file.

```
decks = ("2D_100x100_incompr_linear",  
         "2D_100x100_incompr_linear_faceBC")  
descriptions = ("Testing the single phase incompressible flow solver.",
```

(continues on next page)

(continued from previous page)

```

        "Testing the single phase incompressible flow solver with face_
↪boundary conditions.")

for i in range(len(decks)):
    deck = decks[i]
    description = descriptions[i]
    TestCase(
        name = deck,
        desc = description,
        label = "auto",
        owner = "Ben Corbett",
        independent = True,
        steps = (geosx(deck=deck + ".xml"),)
    )

```

This creates two *TestCases* each of which runs a different problem. The *independent* parameter means that the two *TestCases* can be executed independently of each other. When a *TestCase* executes it uses it's name to create a directory where all the output files are stored so if you have multiple *TestCases* in an *.ats* file it's imperative that they have unique names.

Finally there's the *sedov.ats* file which tests that starting from a restart file has no impact on the final solution.

```

import os

TestCase(
    name = "sedov",
    desc = "Test the basic sedov problem and restart capabilities.",
    label = "auto",
    owner = "Ben Corbett",
    independent = True,
    steps = (geosx(deck="sedov.xml",
                  name="0to100"),
            geosx(deck="sedov.xml",
                  name="50to100",
                  restart_file=os.path.join(testcase_name, "0to100_restart_000000050.
↪root"),
                  baseline_pattern="0to100_restart_[0-9]+\root",
                  allow_rebaseline=False)
    )
)

```

This creates a single *TestCase* That executes *GEOSX* twice. The first step does 100 time steps followed by a *restartcheck* step. The second *geosx* step executes the original 100 time step *xml* file but restarts using the restart file output half way through the first run. Each *geosx* step gets its name from the *xml* file, but this can be overridden by the *name* parameter Furthermore the default behavior is to look for a baseline in the *baselines/<TestCaseName>* directory named *TestStepName_restart_[0-9]+\root*, however the second step overrides this to instead compare against the "0to100" baseline. Because of this it does not allow rebaselining.

You can pass parameters to the *restartcheck* step in a dictionary passed as an argument to the *geosx* step. For example to set the tolerance you would do

```

restartcheck_params={}
restartcheck_params["atol"] = 1.5E-10
restartcheck_params["rtol"] = 1E-12

TestCase(

```

(continues on next page)

(continued from previous page)

```

name = "sedov",
desc = "Test the basic sedov problem and restart capabilities.",
label = "auto",
owner = "Ben Corbett",
independent = True,
steps = (geosx(deck="sedov.xml",
               name="0to100",
               restartcheck_params=restartcheck_params))
)

```

For more info see *integratedTests/geosxats/GeosxAtsTestSteps.py* and *integratedTests/geosxats/GeosxAtsTestCase.py*

Adding a Test

To add a new test create a new folder in the `integratedTests/update/run*` directory. At a minimum this new folder needs to include an `.ats` file. Using the `beamBending` example, after creating `beamBending.ats` the directory should look like

```

- integratedTests/update/run/beamBending/
  - beamBending.ats
  - beamBending.xml

```

At this point you should run the test. Assuming the `geosx` step is successful the `restartcheck` step will fail because there are no baselines. At this point the directory should look like

```

- integratedTests/update/run/beamBending/
  - beamBending/
    - <geosx files>...
    - <ats files>...
  - beamBending.ats
  - beamBending.xml
  - <ats files>...

```

Now run

```
./geosxats.sh -a rebaseline
```

and rebaseline your test. Finally run the test a second time and confirm that it passes. Note that unless you disable the `restartcheck` step you will need to output a restart file. Although not strictly necessary it is best to put the `xml` file in the main *GEOSX* repo and create an relative symbolic link to it in the test directory.

Rebaselining Tests

Occasionally it is necessary to rebaseline one or more tests due to feature changes in the code. We suggest the following workflow:

In the *GEOSX* repository, create a branch with your modifications:

```

cd <GEOSX-path>
git checkout -b user/feature/newFeature

```

Add your changes, confirm it passes all the continuous integration tests, and get approval for a pull request.

Now, confirm that your `integratedTests` submodule is up to date:

```
git submodule
```

This will list the commit hash for all submodules. Check that the integrated tests submodule is on develop and that the commit hash is the same one as the latest GEOSX develop branch points to. If you have somehow fallen behind, go into integratedTests, checkout develop, and pull.

Now go to the integratedTests submodule and check out a branch for your new baselines. It is a good idea to name branch something similar to your feature branch so it is obvious the two branches are related.

```
cd <integratedTests-path>
git checkout -b user/rebase/newFeature
```

Go back to your GEOSX build directory and run the integrated tests

```
cd <build-path>
./geosxats.sh
```

Confirm that any tests that fail need to be **legitimately** rebaselined. Arbitrarily changing baselines defeats the purpose of the integrated tests. In your PR discussion, please identify which tests will change and any unusual behavior.

We can now actually rebaseline the tests

```
./geosxats -a rebaseline
```

You'll be prompted to confirm whether rebaselining is required for every integrated test, one at a time, via a [y/n] prompt. Make sure to only answer y to the tests that you actually want to rebaseline, otherwise correct baselines for already passing tests will still be updated and bloat your pull request and repository size.

Confirm that the rebaselines are working as expected, by cleaning the test dir and re-running the checks:

```
./geosxats -a veryclean
./geosxats
```

At this point you should pass all the integratedTests. Clean the branch and commit your changes to the baseline branch.

```
./geosxats -a veryclean
cd <integratedTests-path>
git status
git add *
git commit -m "Updating baselines"
git push
```

If you haven't already set up your local branch to point to a remote branch, you will be prompted to do so when attempting to push. You will then want to create a pull request in the integratedTests repository. Once you have merge approval for your PR, you can merge your rebase branch into integratedTests/develop.

At this point, you need to get your GEOSX user/feature/newFeature branch pointing to the head commit on integratedTests/develop. We will check out the latest version of the test repo and add it to our feature branch:

```
cd <integratedTests-path>
git checkout develop
git pull

cd <GEOSX-path>
git add integratedTests
git commit -m "Updating integratedTests hash"
git push
```

You may also want to run `git submodule` to confirm the submodule hash is what we expect, the last commit in `integratedTests/develop`. Once your feature branch passes all Continuous Integration tests, it can be successfully merged into `GEOSX/develop`.

Tips

Parallel Tests: On some development machines `geosxats` won't run parallel tests by default (e.g. on an linux laptop or workstation), and as a result many baselines will be skipped. We highly recommend running tests and rebaselining on an MPI-aware platform.

Filtering Checks: A common reason for rebaselining is that you have changed the name of an XML node in the input files. While the baselines may be numerically identical, the restarts will fail because they contain different node names. In this situation, it can be useful to add a filter to the restart check script. If you open `integratedTests/geosxats/helpers/restartcheck.py`, at line 12 you will find a line:

```
EXCLUDE_DEFAULT = [".*/commandLine", ".*schema$", ".*globalToLocalMap"]
```

This variable contains paths to be excluded from the restart checks. For example, we recently renamed the XML block `<SystemSolverParameters/>` to `<LinearSolverParameters/>`. In doing so, we had to rebaseline every test even though we expected no numerical differences. Temporarily adding the following filter helped us rapidly check this was indeed the case:

```
EXCLUDE_DEFAULT = [".*/SystemSolverParameters", ".*LinearSolverParameters", ".*/  
↪commandLine", ".*schema$", ".*globalToLocalMap"]
```

You may find this approach useful for quickly filtering tests to distinguish between expected and unexpected failures.

6.1.8 Benchmarks

In addition to the integrated tests which track code correctness we have a suite of benchmarks that track performance.

Running the benchmarks

Because performance is system specific we currently only support running the benchmarks on the LLNL machines Quartz and Lassen. If you are on either of these machines the script `benchmarks/runBenchmarks.py` can be used to run the benchmarks.

```
> python ../benchmarks/runBenchmarks.py --help
usage: runBenchmarks.py [-h] [-t TIMELIMIT] [-o TIMINGCOLLECTIONDIR]
                        [-e ERRORCOLLECTIONDIR]
                        geosxPath outputDirectory

positional arguments:
  geosxPath              The path to the GEOSX executable to benchmark.
  outputDirectory        The parent directory to run the benchmarks in.

optional arguments:
  -h, --help              show this help message and exit
  -t TIMELIMIT, --timeLimit TIMELIMIT
                        Time limit for the entire script in minutes, the
                        default is 60.
  -o TIMINGCOLLECTIONDIR, --timingCollectionDir TIMINGCOLLECTIONDIR
                        Directory to copy the timing files to.
```

(continues on next page)

(continued from previous page)

```
-e ERRORCOLLECTIONDIR, --errorCollectionDir ERRORCOLLECTIONDIR
    Directory to copy the output from any failed runs to.
```

At a minimum you need to pass the script the path to the GEOSX executable and a directory to run the benchmarks in. This directory will be created if it doesn't exist. The script will collect a list of benchmarks to be run and submit a job to the system's scheduler for each benchmark. This means that you don't need to be in an allocation to run the benchmarks. Note that this is different from the integrated tests where you need to already be in an allocation and an internal scheduler is used to run the individual tests. Since a benchmark is a measure of performance to get consistent results it is important that each time a benchmark is run it has access to the same resources. Using the system scheduler guarantees this.

In addition to whatever outputs the input would normally produce (plot files, restart files, ...) each benchmark will produce an output file `output.txt` containing the standard output and standard error of the run and a `.cali` file containing the Caliper timing data in a format that [Spot](#) can read.

Note: A future version of the script will be able to run only a subset of the benchmarks.

Specifying a benchmark

A group of benchmarks is specified with a standard GEOSX input XML file with an extra `Benchmarks` block added at the top level. This block is ignored by GEOSX itself and only used by the `runBenchmarks.py` script.

```
<Benchmarks>
  <quartz>
    <Run
      name="OMP "
      nodes="1 "
      tasksPerNode="1"
      autoPartition="On"
      timeLimit="10"/>
    <Run
      name="MPI_OMP "
      nodes="1 "
      tasksPerNode="2"
      autoPartition="On"
      timeLimit="10"
      strongScaling="{ 1, 2, 4, 8 }"/>
    <Run
      name="MPI "
      nodes="1 "
      tasksPerNode="36"
      autoPartition="On"
      timeLimit="10"
      strongScaling="{ 1, 2, 4, 8 }"/>
  </quartz>

  <lassen>
    <Run
      name="OMP_CUDA"
      nodes="1 "
      tasksPerNode="1"
      autoPartition="On"
      timeLimit="10"/>
```

(continues on next page)

(continued from previous page)

```

<Run
  name="MPI_OMP_CUDA"
  nodes="1"
  tasksPerNode="4"
  autoPartition="On"
  timeLimit="10"
  strongScaling="{ 1, 2, 4, 8 }"/>
</lassen>
</Benchmarks>

```

[Source: *benchmarks/SSLE-small.xml*]

The Benchmarks block consists of a block for each machine the benchmarks are to run on. Currently the only options are quartz and lassen.

The Run block

Each machine block contains a number of Run blocks each of which specify a family of benchmarks to run. Each Run block must have the following required attributes

- **name:** The name of the family of benchmarks, must be unique among all the other Run blocks on that system.
- **nodes:** An integer which specifies the base number of nodes to run the benchmark with.
- **tasksPerNode:** An integer that specifies the number of tasks to launch per node.

Each Run block may contain the following optional attributes

- **threadsPerTask:** An integer specifying the number of threads to allocate each task.
- **timeLimit:** An integer specifying the time limit in minutes to pass to the system scheduler when submitting the benchmark.
- **args:** containing any extra command line arguments to pass to GEOSX.
- **autoPartition:** Either On or Off, not specifying autoPartition is equivalent to autoPartition="Off". When auto partitioning is enabled the script will compute the number of x, y and z partitions such that the the resulting partition is close to a perfect cube as possible, ie with 27 tasks $x = 3, y = 3, z = 3$ and with 36 tasks $x = 4, y = 3, z = 3$. This is optimal when the domain itself is a cube, but will be suboptimal otherwise.
- **strongScaling:** A list of unique integers specifying the factors to scale the number of nodes by. If N number are provided then N benchmarks are run and benchmark i uses $\text{nodes} * \text{strongScaling}[i]$ nodes. Not specifying strongScaling is equivalent to strongScaling="{ 1 }".

Looking at the example Benchmarks block above on Lassen one benchmark from the OMP_CUDA family will be run with one node and one task. Four benchmarks from the MPI_OMP_CUDA family will be run with one, two, four and eight nodes and four tasks per node.

Note that specifying a time limit for each benchmark family can greatly decrease the time spent waiting in the scheduler's queue. A good rule of thumb is that the time limit should be twice as long as it takes to run the longest benchmark in the family.

Adding a benchmark problem

To add a new group of benchmarks you need to create an XML input file describing the problem to be run. Then you need to add the Benchmarks block described above which specifies the specific benchmarks. Finally add a symbolic link to the input file in benchmarks and run the benchmarks to make sure everything works as expected.

Viewing the results

Each night the [NightlyTests](#) repository runs the benchmarks on both Quartz and Lassen, the `timingFiles` directory contains all of the resulting caliper output files. If you're on LC then these files are duplicated at `/usr/gapps/GEOSX/timingFiles/` and if you have LC access you can view them in [Spot](#). You can also open these files in Python and analyse them (See [Opening Spot caliper files in Python](#)).

If you want to run the benchmarks on your local branch and compare the results with develop you can use the `benchmarks/compareBenchmarks.py` python script. This requires that you run the benchmarks on your branch and on develop. It will print out a table with the initialization time speed up and run time speed up, so a run speed up of 2x means your branch runs twice as fast as develop where as a initialization speed up of 0.5x means the set up takes twice as long.

Note: A future version of the script will be able to pull timing results straight from the `.cali` files so that if you have access to the [NightlyTests](#) timing files you won't need to run the benchmarks on develop. Furthermore it will be able to provide more detailed information than just initialization and run times.

6.1.9 Basic profiling with CALIPER

GEOSX is equipped with [Caliper](#) timers. We integrate Caliper into GEOSX by marking source-code sections of interest such as computational kernels or initialization steps. Caliper is included in the GEOSX TPL library and is built by adding the following cmake configuration to a host-config file.

```
option( ENABLE_CALIPER "Enables CALIPER" On )
```

GEOSX/Caliper Annotation Macros

The following macros may be used to annotate GEOSX:

- `GEOSX_MARK_SCOPE (name)` - Marks a scope with the given name.
- `GEOSX_MARK_FUNCTION` - Marks a function with the name of the function. The name includes the namespace the function is in but not any of the template arguments or parameters. Therefore overloaded function all show up as one entry. If you would like to mark up a specific overload use `GEOSX_MARK_SCOPE` with a unique name.
- `GEOSX_MARK_BEGIN (name)` - Marks the beginning of a user defined code region.
- `GEOSX_MARK_END (name)` - Marks the end of user defined code region.

Configuring Caliper

Caliper configuration is done by specifying a string to initialize Caliper with via the `-t` option. A few options are listed below but we refer the reader to [Caliper Config](#) for the full Caliper tutorial.

- `-t runtime-report,max_column_width=200` Will make Caliper print aggregated timing information to standard out, with a column width large enough that it doesn't truncate most function names.
- `-t runtime-report,max_column_width=200,profile.cuda` Does the same as the above, but also instruments CUDA API calls. This is only an option when building with CUDA.
- `-t runtime-report,aggregate_across_ranks=false` Will make Caliper write per rank timing information to standard out. This isn't useful when using more than one rank but it does provide more information for single rank runs.

- `-t spot()` Will make Caliper output a *.cali* timing file that can be viewed in the Spot web server.

Using Adiak

Adiak is a library that allows the addition of meta-data to the Caliper Spot output, it is enabled with Caliper. This meta-data allows you to easily slice and dice the timings available in the Spot web server. To export meta-data use the *adiak::value* function.

See [Adiak API](#) for the full Adiak documentation.

Using Spot

To use Spot you will need an LC account and a directory full of *.cali* files you would like to analyse. Point your browser to [Spot](#) and open up the directory containing the timing files.

Opening Spot caliper files in Python

An example Python program for analyzing Spot Caliper files in Python is provided below. Note that it requires *pandas* and *hatchet* both of which can be installed with a package manager. In addition it requires that *cali-query* is in the *PATH* variable, this is built with Caliper so we can just point it into the TPLs.

```
import sys
import subprocess
import json
import os

import pandas as pd
from IPython.display import display, HTML

# Import hatchet, on LC this can be done by adding hatchet to PYTHONPATH
sys.path.append('/usr/gapps/spot/live/hatchet')
import hatchet as ht

# Add cali-query to PATH
cali_query_path = "/usr/gapps/GEOSX/thirdPartyLibs/2020-06-12/install-quartz-gcc@8.1.
↳0-release/caliper/bin"
os.environ["PATH"] += os.pathsep + cali_query_path

CALI_FILES = [
{ "cali_file": "/usr/gapps/GEOSX/timingFiles/200612-04342891243.cali", "metric_name":
↳"avg#inclusive#sum#time.duration"},
{ "cali_file": "/usr/gapps/GEOSX/timingFiles/200611-044740108300.cali", "metric_name
↳": "avg#inclusive#sum#time.duration"},
]

grouping_attribute = "prop:nested"
default_metric = "avg#inclusive#sum#time.duration"
query = "select %s,sum(%s) group by %s format json-split" % (grouping_attribute,
↳default_metric, grouping_attribute)

gf1 = ht.GraphFrame.from_caliper(CALI_FILES[0]['cali_file'], query)
gf2 = ht.GraphFrame.from_caliper(CALI_FILES[1]['cali_file'], query)

# Print the tree representation using the default metric
# Also print the resulting dataframe with metadata
```

(continues on next page)

(continued from previous page)

```

print(gf1.tree(color=True, metric="sum#"+default_metric))
display(HTML(gf1.dataframe.to_html()))

# Print the tree representation using the default metric
# Also print the resulting dataframe with metadata
print(gf2.tree(color=True, metric="sum#"+default_metric))
display(HTML(gf2.dataframe.to_html()))

# Compute the speedup between the first two cali files (exclusive and inclusive,
↪metrics only)
gf3 = (gf1 - gf2) / gf2
print(gf3.tree(color=True, metric="sum#"+default_metric))

# Compute the difference between the first two cali files (exclusive and inclusive,
↪metrics only)
# Print the resulting tree
gf4 = gf1 - gf2
print(gf4.tree(color=True, metric="sum#"+default_metric))

# Compute the sum of the first two cali files (exclusive and inclusive metrics only)
# Print the resulting tree
gf5 = gf1 + gf2
print(gf5.tree(color=True, metric="sum#"+default_metric))

```

6.2 Code Components

The main code components are described here.

6.2.1 Data Repository

The GEOSX “Data Repository” is intended to provide the building blocks for the code structure within GEOSX. The “Data Repository” provides a general capability to store arbitrary data and objects in a hierarchical structure, similar to a standard file system.

The components/classes of the data structure that a developer will require some knowledge of are:

MappedVector

A superposition of a contiguous and an associative container.

Description

The container stores pointers to objects (which are themselves heap-allocated). Each element may be optionally *owned* by the container, in which case it will be deleted upon removal or container destruction. The pointers are stored in a contiguous memory allocation, and thus are accessible through an integral index lookup. In addition, there is a map that provides a key lookup capability to the container if that is the preferred interface.

The container template has four type parameters:

- T is the object type pointed to by container entries
- T_PTR is a pointer-to-T type which must be either T * (default) or std::unique_ptr<T>

- `KEY_TYPE` is the type of key used in associative lookup
- `INDEX_TYPE` is the type used in index lookup

Element access

`MappedVector` provides three main types of data access using `[]` operator:

- **Index lookup** is the fastest way of element random access if the ordinal index is known.
- **Key lookup** is similar to key lookup of any associative container and incurs similar cost.
- **KeyIndex lookup** uses a special type, `KeyIndex`, that contains both a key and an index. Initially the index is unknown and the key is used for the lookup. The `KeyIndex` is modified during lookup, storing the index located. If the user persists the `KeyIndex` object, they may reuse it in subsequent accesses and get the benefit of direct index access.

In addition to these, an STL-conformant iterator interface is available via `begin()` and `end()` methods. The type iterated over is a key-pointer pair (provided as `value_type` alias).

API documentation

`MappedVector`

Group

`dataRepository::Group` serves as a base class for most objects in GEOSX. In GEOSX, the `Group` may be thought of as an analogy to the file folder in a hierarchical filesystem-like structure. As such, a `Group` is used as a container class that holds a collection of other `Groups`, or sub-`Groups`, a pointer to the parent of the `Group`, and a collection of `Wrappers`. The `Group` also defines a general capability to create and traverse/access the objects in the hierarchy. The `Wrappers` contained in a `Group` may be of arbitrary type, but in the case of an `LvArray` object, a `Group` size and capacity may be translated down to array, thereby keeping a collection of wrapped `Array` objects that will resize in unison with the `Group`. Each group has a string “name” that defines its key in the parent `Group`. This key string must be unique in the scope of the parent `Group`.

Implementation Details

Some noteworthy implementation details inside the declaration of `dataRepository::Group` are:

```
/// The default key type for entries in the hierarchy.
using keyType = string;

/// The default index type for entries the hierarchy.
using indexType = localIndex;
```

- In the GEOSX repository, the `keyType` is specified to be a `string` for all collection objects, while the `indexType` is specified to be a `localIndex`. The types are set in the `common/DataTypes.hpp` file, but are typically a `string` and a `std::ptrdiff_t` respectively.

```
/// The template specialization of MappedVector to use for the collection of sub-
↪Group objects.
using subGroupMap = MappedVector< Group, Group *, keyType, indexType >;
```

(continues on next page)

(continued from previous page)

```

/// The template specialization of MappedVector to use for the collection wrappers_
↳objects.
using wrapperMap = MappedVector< WrapperBase, WrapperBase *, keyType, indexType >;

```

- The subGroupMap and wrapperMap aliases represent the type of container that the collection of sub-Groups and Wrappers are stored in for each Group. These container types are template specializations of the MappedVector class, which store a pointer to a type, and provides functionality for a key or index based lookup. More details may be found in the documentation for MappedVector.

```

/// The parent Group that contains "this" Group in its "sub-Group" collection.
Group * m_parent = nullptr;

/// Specification that this group will have the same m_size as m_parent.
integer m_sizedFromParent;

/// The container for the collection of all wrappers contained in "this" Group.
wrapperMap m_wrappers;

/// The container for the collection of all sub-groups contained in "this" Group.
subGroupMap m_subGroups;

/// The size/length of this Group...and all Wrapper<> that are are specified to_
↳have the same size as their
/// owning group.
indexType m_size;

/// The capacity for wrappers in this group...and all Wrapper<> that are specified_
↳to have the same size as their
/// owning group.
indexType m_capacity;

/// The name/key of this Group in its parent collection of sub-Groups.
string m_name;

/// Verbosity flag for group logs
integer m_logLevel;

```

- The m_parent member is a pointer to the Group that contains the current Group as part of its collection of sub-Groups.

Warning: The existence of the non-const m_parent gives the current Group access to alter the parent Group. Special care should be taken to avoid using this access whenever possible. Remember... with great power comes great responsibility.

- The m_wrappers member is the collection of Wrappers contained in the current Group.
- The m_subGroups member is the collection of Groups contained in the current Group.
- The m_size and m_capacity members are used to set the size and capacity of any objects contained in the m_wrappers collection that have been specified to be set by their owning Group. This is typically only useful for Array types and is implemented within the WrapperBase object.
- The m_name member is the key of this Group in the collection of m_parent->m_subGroups. This key is unique in the scope of m_parent, so some is required when constructing the hierarchy.

Interface Functions

The public interface for `dataRepository::Group` provides functionality for constructing a hierarchy, and traversing that hierarchy, as well as accessing the contents of objects stored in the `Wrapper` containers stored within a `Group`.

Adding New Groups

To add new sub-Groups there are several `registerGroup` functions that add a new `Group` under the calling `Group` scope. A listing of these functions is provided:

```
/**
 * @name Sub-group registration interface
 */
///<

/**
 * @brief Register a new Group as a sub-group of current Group.
 *
 * @tparam T The type of the Group to add/register. This should be a type that
↳ derives from Group.
 * @param[in] name The name of the group to use as a string key.
 * @param[in] newObject A unique_ptr to the object that is being registered.
 * @return A pointer to the newly registered Group.
 *
 * Registers a Group or class derived from Group as a subgroup of this Group and
↳ takes ownership.
 */
template< typename T = Group >
T & registerGroup( string const & name, std::unique_ptr< T > newObject )
{
    newObject->m_parent = this;
    return dynamicCast< T & >( *m_subGroups.insert( name, newObject.release(), true )_
↳ );
}

/**
 * @brief @copybrief registerGroup(string const &,std::unique_ptr<T>)
 *
 * @tparam T The type of the Group to add/register. This should be a type that
↳ derives from Group.
 * @param[in] name The name of the group to use as a string key.
 * @param[in] newObject A unique_ptr to the object that is being registered.
 * @return A pointer to the newly registered Group.
 *
 * Registers a Group or class derived from Group as a subgroup of this Group but
↳ does not take ownership.
 */
template< typename T = Group >
T & registerGroup( string const & name, T * newObject )
{ return dynamicCast< T & >( *m_subGroups.insert( name, newObject, false ) ); }

/**
 * @brief @copybrief registerGroup(string const &,std::unique_ptr<T>)
 */
```

(continues on next page)

(continued from previous page)

```

    * @tparam T The type of the Group to add/register. This should be a type that
↳ derives from Group.
    * @param[in] name The name of the group to use as a string key.
    * @return      A pointer to the newly registered Group.
    *
    * Creates and registers a Group or class derived from Group as a subgroup of this
↳ Group.
    */
    template< typename T = Group >
    T & registerGroup( string const & name )
    { return registerGroup< T >( name, std::move( std::make_unique< T >( name, this ) ) );
↳ }; }

/**
    * @brief @copybrief registerGroup(string const &,std::unique_ptr<T>)
    *
    * @tparam T The type of the Group to add/register. This should be a type that
↳ derives from Group.
    * @param keyIndex A KeyIndexT object that will be used to specify the name of
    * the new group. The index of the KeyIndex will also be set.
    * @return A pointer to the newly registered Group, or @c nullptr if no group was
↳ registered.
    *
    * Creates and registers a Group or class derived from Group as a subgroup of this
↳ Group.
    */
    template< typename T = Group >
    T & registerGroup( subGroupMap::KeyIndex const & keyIndex )
    {
        T & rval = registerGroup< T >( keyIndex.key(), std::move( std::make_unique< T >(
↳ keyIndex.key(), this ) ) );
        keyIndex.setIndex( m_subGroups.getIndex( keyIndex.key() ) );
        return rval;
    }

/**
    * @brief @copybrief registerGroup(string const &,std::unique_ptr<T>)
    *
    * @tparam T The type of the Group to add/register. This should be a type that
↳ derives from Group.
    * @tparam TBASE The type whose type catalog will be used to look up the new sub-
↳ group type
    * @param[in] name      The name of the group to use as a string key.
    * @param[in] catalogName The catalog name of the new type.
    * @return      A pointer to the newly registered Group.
    *
    * Creates and registers a Group or class derived from Group as a subgroup of this
↳ Group.
    */
    template< typename T = Group, typename TBASE = Group >
    T & registerGroup( string const & name, string const & catalogName )
    {
        std::unique_ptr< TBASE > newGroup = TBASE::CatalogInterface::Factory( catalogName,
↳ name, this );
        return registerGroup< T >( name, std::move( newGroup ) );
    }

```

(continues on next page)

(continued from previous page)

```

/**
 * @brief Removes a child group from this group.
 * @param name the name of the child group to remove from this group.
 */
void deregisterGroup( string const & name );

/**
 * @brief Creates a new sub-Group using the ObjectCatalog functionality.
 * @param[in] childKey The name of the new object type's key in the
 *                  ObjectCatalog.
 * @param[in] childName The name of the new object in the collection of
 *                  sub-Groups.
 * @return A pointer to the new Group created by this function.
 */
virtual Group * createChild( string const & childKey, string const & childName );

///<@}

```

These functions all take in a name for the new Group, which will be used as the key when trying to access the Group in the future. Some variants create a new Group, while some variants take in an existing Group. The template argument is to specify the actual type of the Group as it is most likely a type that derives from Group that is we would like to create in the repository. Please see the doxygen documentation for a detailed description of each option.

Getting Groups

The collection of functions to retrieve a Group and their descriptions are taken from source and shown here:

```

/**
 * @name Sub-group retrieval methods.
 *
 * This collection of functions are used to get a sub-Group from the current group.
↳ Various methods
 * for performing the lookup are provided (localIndex, string, KeyIndex), and each
↳ have their
 * advantages and costs. The lowest cost lookup is the "localIndex" lookup. The
↳ KeyIndex lookup
 * will add a cost for checking to make sure the index stored in KeyIndex is valid
↳ (a string
 * compare, and a hash if it is incorrect). The string lookup is the full cost hash
↳ lookup every
 * time that it is called.
 *
 * The template parameter specifies the "type" that the caller expects to lookup,
↳ and thus attempts
 * to cast the pointer that is stored in m_subGroups to a pointer of the desired
↳ type. If this
 * cast fails, then a @p nullptr is returned. If no template parameter is specified,
↳ then a default
 * type of Group is assumed.
 */
///<@{

/**
 * @brief Return a pointer to a sub-group of the current Group.

```

(continues on next page)

(continued from previous page)

```

    * @tparam T The type of subgroup.
    * @tparam KEY The type of the lookup.
    * @param key The key used to perform the lookup.
    * @return A pointer to @p T that refers to the sub-group, if the Group does not_
→ exist or it
    * has an incompatible type a @c nullptr is returned.
    */
    template< typename T = Group, typename KEY = void >
    T * getGroupPointer( KEY const & key )
    { return dynamicCast< T * >( m_subGroups[ key ] ); }

    /**
    * @copydoc getGroupPointer(KEY const &)
    */
    template< typename T = Group, typename KEY = void >
    T const * getGroupPointer( KEY const & key ) const
    { return dynamicCast< T const * >( m_subGroups[ key ] ); }

    /**
    * @brief Return a reference to a sub-group of the current Group.
    * @tparam T The type of subgroup.
    * @tparam KEY The type of the lookup.
    * @param key The key used to perform the lookup.
    * @return A reference to @p T that refers to the sub-group.
    * @throw std::domain_error If the Group does not exist is thrown.
    */
    template< typename T = Group, typename KEY = void >
    T & getGroup( KEY const & key )
    {
        Group * const child = m_subGroups[ key ];
        GEOSX_THROW_IF( child == nullptr, "Group " << getPath() << " doesn't have a child
→ " << key, std::domain_error );
        return dynamicCast< T & >( *child );
    }

    /**
    * @copydoc getGroup( KEY const & )
    */
    template< typename T = Group, typename KEY = void >
    T const & getGroup( KEY const & key ) const
    {
        Group const * const child = m_subGroups[ key ];
        GEOSX_THROW_IF( child == nullptr, "Group " << getPath() << " doesn't have a child
→ " << key, std::domain_error );
        return dynamicCast< T const & >( *child );
    }

    /**
    * @brief Retrieve a group from the hierarchy using a path.
    * @tparam T type of subgroup
    * @param[in] path a unix-style string (absolute, relative paths valid)
    * to lookup the Group to return. Absolute paths search
    * from the tree root, while relative - from current group.
    * @return A reference to @p T that refers to the sub-group.
    * @throw std::domain_error If the Group doesn't exist.
    */
    template< typename T = Group >

```

(continues on next page)

(continued from previous page)

```

T & getGroupByPath( string const & path )
{ return dynamicCast< T & >( const_cast< Group & >( getBaseGroupByPath( path ) ) );  

↪ }

/**
 * @copydoc getGroupByPath(string const &)
 */
template< typename T = Group >
T const & getGroupByPath( string const & path ) const
{ return dynamicCast< T const & >( getBaseGroupByPath( path ) ); }

```

Register Wrappers

```

/**
 * @name Wrapper registration interface
 */
///  

/**
 * @brief Create and register a Wrapper around a new object.
 * @tparam T The type of the object allocated.
 * @tparam TBASE The type of the object that the Wrapper holds.
 * @param[in] name the name of the wrapper to use as a string key
 * @param[out] rkey a pointer to a index type that will be filled with the new
 *   Wrapper index in this Group
 * @return A reference to the newly registered/created Wrapper
 */
template< typename T, typename TBASE=T >
Wrapper< TBASE > & registerWrapper( string const & name,
                                   wrapperMap::KeyIndex::index_type * const rkey =  

↪ nullptr );

/**
 * @copybrief registerWrapper(string const &,wrapperMap::KeyIndex::index_type *  

↪ const)
 * @tparam T the type of the wrapped object
 * @tparam TBASE the base type to cast the returned wrapper to
 * @param[in] viewKey The KeyIndex that contains the name of the new Wrapper.
 * @return A reference to the newly registered/created Wrapper
 */
template< typename T, typename TBASE=T >
Wrapper< TBASE > & registerWrapper( Group::wrapperMap::KeyIndex const & viewKey );

/**
 * @brief Register a Wrapper around a given object and take ownership.
 * @tparam T the type of the wrapped object
 * @param[in] name the name of the wrapper to use as a string key
 * @param[in] newObject an owning pointer to the object that is being registered
 * @return A reference to the newly registered/created Wrapper
 */
template< typename T >
Wrapper< T > & registerWrapper( string const & name, std::unique_ptr< T > newObject,  

↪ );

```

(continues on next page)

(continued from previous page)

```

/**
 * @brief Register a Wrapper around an existing object, does not take ownership of
↳the object.
 * @tparam T the type of the wrapped object
 * @param[in] name the name of the wrapper to use as a string key
 * @param[in] newObject a pointer to the object that is being registered
 * @return A reference to the newly registered/created Wrapper
 */
template< typename T >
Wrapper< T > & registerWrapper( string const & name,
                               T * newObject );

/**
 * @brief Register and take ownership of an existing Wrapper.
 * @param name The name of the wrapper to use as a string key
 * @param wrapper A pointer to the an existing wrapper.
 * @return An un-typed pointer to the newly registered/created wrapper
 */
WrapperBase & registerWrapper( string const & name,
                              std::unique_ptr< WrapperBase > wrapper );

/**
 * @brief Removes a Wrapper from this group.
 * @param name the name of the Wrapper to remove from this group.
 */
void deregisterWrapper( string const & name );

//@}

```

Getting Wrappers/Wrapped Objects

```

/**
 * @name Untyped wrapper retrieval methods
 *
 * These functions query the collection of Wrapper objects for the given
 * index/name/KeyIndex and returns a WrapperBase pointer to the object if
 * it exists. If it is not found, nullptr is returned.
 */
//@{

/**
 * @brief Return a reference to a WrapperBase stored in this group.
 * @tparam KEY The lookup type.
 * @param key The value used to lookup the wrapper.
 * @return A reference to the WrapperBase that resulted from the lookup.
 * @throw std::domain_error if the wrapper doesn't exist.
 */
template< typename KEY >
WrapperBase const & getWrapperBase( KEY const & key ) const
{
    WrapperBase const * const wrapper = m_wrappers[ key ];
    GEOSX_THROW_IF( wrapper == nullptr, "Group " << getPath() << " doesn't have a
↳child " << key, std::domain_error );
    return *wrapper;
}

```

(continues on next page)

(continued from previous page)

```

}

/**
 * @copydoc getWrapperBase(KEY const &) const
 */
template< typename KEY >
WrapperBase & getWrapperBase( KEY const & key )
{
    WrapperBase * const wrapper = m_wrappers[ key ];
    GEOSX_THROW_IF( wrapper == nullptr, "Group " << getPath() << " doesn't have a_
→child " << key, std::domain_error );
    return *wrapper;
}

/**
 * @brief
 * @param name
 * @return
 */
indexType getWrapperIndex( string const & name ) const
{ return m_wrappers.getIndex( name ); }

/**
 * @brief Get access to the internal wrapper storage.
 * @return a reference to wrapper map
 */
wrapperMap const & wrappers() const
{ return m_wrappers; }

/**
 * @copydoc wrappers() const
 */
wrapperMap & wrappers()
{ return m_wrappers; }

/**
 * @brief Return the number of wrappers.
 * @return The number of wrappers.
 */
indexType numWrappers() const
{ return m_wrappers.size(); }

///<}

/**
 * @name Typed wrapper retrieval methods
 *
 * These functions query the collection of Wrapper objects for the given
 * index/key and returns a Wrapper<T> pointer to the object if
 * it exists. The template parameter @p T is used to perform a cast
 * on the WrapperBase pointer that is returned by the lookup, into
 * a Wrapper<T> pointer. If the wrapper is not found, or the
 * WrapperBase pointer cannot be cast to a Wrapper<T> pointer, then nullptr
 * is returned.
 */
///<{

```

(continues on next page)

(continued from previous page)

```

/**
 * @brief Check if a wrapper exists
 * @tparam LOOKUP_TYPE the type of key used to perform the lookup.
 * @param[in] lookup a lookup value used to search the collection of wrappers
 * @return @p true if wrapper exists (regardless of type), @p false otherwise
 */
template< typename LOOKUP_TYPE >
bool hasWrapper( LOOKUP_TYPE const & lookup ) const
{ return m_wrappers[ lookup ] != nullptr; }

/**
 * @brief Retrieve a Wrapper stored in this group.
 * @tparam T the object type contained in the Wrapper
 * @tparam LOOKUP_TYPE the type of key used to perform the lookup
 * @param[in] index a lookup value used to search the collection of wrappers
 * @return A reference to the Wrapper<T> that resulted from the lookup.
 * @throw std::domain_error if the Wrapper doesn't exist.
 */
template< typename T, typename LOOKUP_TYPE >
Wrapper< T > const & getWrapper( LOOKUP_TYPE const & index ) const
{
    WrapperBase const & wrapper = getWrapperBase( index );
    return dynamicCast< Wrapper< T > const &>( wrapper );
}

/**
 * @copydoc getWrapper(LOOKUP_TYPE const &) const
 */
template< typename T, typename LOOKUP_TYPE >
Wrapper< T > & getWrapper( LOOKUP_TYPE const & index )
{
    WrapperBase & wrapper = getWrapperBase( index );
    return dynamicCast< Wrapper< T > &>( wrapper );
}

/**
 * @brief Retrieve a Wrapper stored in this group.
 * @tparam T the object type contained in the Wrapper
 * @tparam LOOKUP_TYPE the type of key used to perform the lookup
 * @param[in] index a lookup value used to search the collection of wrappers
 * @return A pointer to the Wrapper<T> that resulted from the lookup, if the Wrapper
 *         doesn't exist or has a different type a @c nullptr is returned.
 */
template< typename T, typename LOOKUP_TYPE >
Wrapper< T > const * getWrapperPointer( LOOKUP_TYPE const & index ) const
{ return dynamicCast< Wrapper< T > const * >( m_wrappers[ index ] ); }

/**
 * @copydoc getWrapperPointer(LOOKUP_TYPE const &) const
 */
template< typename T, typename LOOKUP_TYPE >
Wrapper< T > * getWrapperPointer( LOOKUP_TYPE const & index )
{ return dynamicCast< Wrapper< T > * >( m_wrappers[ index ] ); }

///  


```

(continues on next page)

(continued from previous page)

```

* @name Wrapper data access methods.
*
* These functions can be used to get referece/pointer access to the data
* stored by wrappers in this group. They are essentially just shortcuts for
* @p Group::getWrapper() and @p Wrapper<T>::getReference().
* An additional template parameter can be provided to cast the return pointer
* or reference to a base class pointer or reference (e.g. Array to ArrayView).
*/
///@{

/**
* @brief Look up a wrapper and get reference to wrapped object.
* @tparam T return value type
* @tparam WRAPPEDTYPE wrapped value type (by default, same as return)
* @tparam LOOKUP_TYPE type of value used for wrapper lookup
* @param lookup value for wrapper lookup
* @return reference to @p T
* @throw A std::domain_error if the Wrapper does not exist.
*/
template< typename T, typename LOOKUP_TYPE >
GEOSX_DECLTYPE_AUTO_RETURN
getReference( LOOKUP_TYPE const & lookup ) const
{ return getWrapper< T >( lookup ).reference(); }

/**
* @copydoc getReference(LOOKUP_TYPE const &) const
*/
template< typename T, typename LOOKUP_TYPE >
T & getReference( LOOKUP_TYPE const & lookup )
{ return getWrapper< T >( lookup ).reference(); }

```

Looping Interface

```

/**
* @name Functor-based subgroup iteration
*
* These functions loop over sub-groups and executes a functor that uses the sub-
↳group as an
* argument. The functor is only executed if the group can be cast to a certain_
↳type specified
* by the @p ROUPTYPE/S pack. The variadic list consisting of @p GROUPTYPE/S will_
↳be used recursively
* to check if the group is able to be cast to the one of these types. The first_
↳type in the
* @p GROUPTYPE/S list will be used to execute the functor, and the next sub-group_
↳will be processed.
*/
///@{

/**
* @brief Apply the given functor to subgroups that can be casted to one of_
↳specified types.
* @tparam GROUPTYPE the first type that will be used in the attempted casting of_
↳group.

```

(continues on next page)

(continued from previous page)

```

    * @tparam GROUPTYPES a variadic list of types that will be used in the attempted_
↳casting of group.
    * @tparam LAMBDA      the type of functor to call
    * @param[in] lambda    the functor to call on subgroups
    */
    template< typename GROUPTYPE = Group, typename ... GROUPTYPES, typename LAMBDA >
    void forSubGroups( LAMBDA && lambda )
    {
        for( auto & subGroupIter : m_subGroups )
        {
            applyLambdaToContainer< GROUPTYPE, GROUPTYPES... >( *subGroupIter.second, [&](
↳auto & castedSubGroup )
            {
                lambda( castedSubGroup );
            } );
        }
    }

    /**
    * @copydoc forSubGroups(LAMBDA &&)
    */
    template< typename GROUPTYPE = Group, typename ... GROUPTYPES, typename LAMBDA >
    void forSubGroups( LAMBDA && lambda ) const
    {
        for( auto const & subGroupIter : m_subGroups )
        {
            applyLambdaToContainer< GROUPTYPE, GROUPTYPES... >( *subGroupIter.second, [&](
↳auto const & castedSubGroup )
            {
                lambda( castedSubGroup );
            } );
        }
    }

    /**
    * @copybrief forSubGroups(LAMBDA &&)
    * @tparam GROUPTYPE      the first type that will be used in the attempted_
↳casting of group.
    * @tparam GROUPTYPES    a variadic list of types that will be used in the_
↳attempted casting of group.
    * @tparam LOOKUP_CONTAINER type of container of subgroup lookup keys (names or_
↳indices), must support range-based for
    * loop
    * @tparam LAMBDA        type of functor callable with an index in lookup_
↳container and a reference to casted
    * subgroup
    * @param[in] subGroupKeys container with subgroup lookup keys (e.g. names or_
↳indices) to apply the functor to
    * @param[in] lambda      the functor to call
    */
    template< typename GROUPTYPE = Group, typename ... GROUPTYPES, typename LOOKUP_
↳CONTAINER, typename LAMBDA >
    void forSubGroups( LOOKUP_CONTAINER const & subGroupKeys, LAMBDA && lambda )
    {
        localIndex counter = 0;
        for( auto const & subgroup : subGroupKeys )
        {

```

(continues on next page)

(continued from previous page)

```

    applyLambdaToContainer< GROUPTYPE, GROUPTYPES... >( getGroup( subgroup ), [&](
↳auto & castedSubGroup )
    {
        lambda( counter, castedSubGroup );
    } );
    ++counter;
}

/**
 * @copybrief forSubGroups(LAMBDA &&)
 * @tparam GROUPTYPE      the first type that will be used in the attempted_
↳casting of group.
 * @tparam GROUPTYPES     a variadic list of types that will be used in the_
↳attempted casting of group.
 * @tparam LOOKUP_CONTAINER type of container of subgroup lookup keys (names or_
↳indices), must support range-based for
 * loop
 * @tparam LAMBDA         type of functor callable with an index in lookup_
↳container and a reference to casted
 * subgroup
 * @param[in] subGroupKeys container with subgroup lookup keys (e.g. names or_
↳indices) to apply the functor to
 * @param[in] lambda       the functor to call
 */
template< typename GROUPTYPE = Group, typename ... GROUPTYPES, typename LOOKUP_
↳CONTAINER, typename LAMBDA >
void forSubGroups( LOOKUP_CONTAINER const & subGroupKeys, LAMBDA && lambda ) const
{
    localIndex counter = 0;
    for( auto const & subgroup : subGroupKeys )
    {
        applyLambdaToContainer< GROUPTYPE, GROUPTYPES... >( getGroup( subgroup ), [&](
↳auto const & castedSubGroup )
        {
            lambda( counter, castedSubGroup );
        } );
        ++counter;
    }
}
///  

/**
 * @name Functor-based wrapper iteration
 *
 * These functions loop over the wrappers contained in this group, and executes a_
↳functor that
 * uses the Wrapper as an argument. The functor is only executed if the Wrapper can_
↳be casted to
 * a certain type specified by the @p TYPE/S pack. The variadic list consisting of
 * @p TYPE/S will be used recursively to check if the Wrapper is able to be casted_
↳to the
 * one of these types. The first type in the @p WRAPPERTYPE/S list will be used to_
↳execute the
 * functor, and the next Wrapper will be processed.
 */
///  


```

(continues on next page)

(continued from previous page)

```

/**
 * @brief Apply the given functor to wrappers.
 * @tparam LAMBDA the type of functor to call
 * @param[in] lambda the functor to call
 */
template< typename LAMBDA >
void forWrappers( LAMBDA && lambda )
{
    for( auto & wrapperIter : m_wrappers )
    {
        lambda( *wrapperIter.second );
    }
}

/**
 * @copydoc forWrappers(LAMBDA &&)
 */
template< typename LAMBDA >
void forWrappers( LAMBDA && lambda ) const
{
    for( auto const & wrapperIter : m_wrappers )
    {
        lambda( *wrapperIter.second );
    }
}

/**
 * @brief Apply the given functor to wrappers that can be cast to one of specified_
↳types.
 * @tparam TYPE the first type that will be used in the attempted casting of_
↳Wrapper
 * @tparam TYPES a variadic list of types that will be used in the attempted_
↳casting of Wrapper
 * @tparam LAMBDA the type of functor to call
 * @param[in] lambda the functor to call
 */
template< typename TYPE, typename ... TYPES, typename LAMBDA >
void forWrappers( LAMBDA && lambda )
{
    for( auto & wrapperIter : m_wrappers )
    {
        applyLambdaToContainer< Wrapper< TYPE >, Wrapper< TYPES >... >( *wrapperIter.
↳second,
                                                                    std::forward<_
↳LAMBDA >( lambda ));
    }
}

/**
 * @brief Apply the given functor to wrappers that can be cast to one of specified_
↳types.
 * @tparam TYPE the first type that will be used in the attempted casting of_
↳Wrapper
 * @tparam TYPES a variadic list of types that will be used in the attempted_
↳casting of Wrapper
 * @tparam LAMBDA the type of functor to call

```

(continues on next page)

(continued from previous page)

```

    * @param[in] lambda the functor to call
    */
template< typename TYPE, typename ... TYPES, typename LAMBDA >
void forWrappers( LAMBDA && lambda ) const
{
    for( auto const & wrapperIter : m_wrappers )
    {
        applyLambdaToContainer< Wrapper< TYPE >, Wrapper< TYPES >... >( *wrapperIter.
↪second,
                                                                    std::forward<
↪LAMBDA >( lambda ));
    }
}

//@}

```

Wrapper

This class encapsulates an object for storage in a `Group` and provides an interface for performing some common operations on that object.

Description

In the filesystem analogy, a `Wrapper` may be thought of as a file that stores actual data. Each `Wrapper` belong to a single `Group` much like a file belongs to a filesystem directory. In general, more than one wrapper in the tree may refer to the same wrapped object, just like symlinks in the file system may refer to the same file. However, only one wrapper should be *owning* the data (see below).

In the XML input file, `Wrapper` correspond to attribute of an XML element representing the containing `Group`. See [XML Input](#) for the relationship between XML input files and Data Repository.

`Wrapper<T>` is templated on the type of object it encapsulates, thus providing strong type safety when retrieving the objects. As each `Wrapper` class instantiation will be a distinct type, `Wrapper` derives from a non-templated `WrapperBase` class that defines a common interface. `WrapperBase` is the type of pointer that is stored in the `MappedVector` container within a `Group`.

`WrapperBase` provides several interface functions that delegate the work to the wrapped object if it supports the corresponding method signature. This allows a collection of heterogeneous wrappers (i.e. over different types) to be treated uniformly. Examples include:

- `size()`
- `resize(newSize)`
- `reserve(newCapacity)`
- `capacity()`
- `move(LvArray::MemorySpace)`

A `Wrapper` may be *owning* or *non-owning*, depending on how it's constructed. An *owning* `Wrapper` will typically either take a previously allocated object via `std::unique_ptr<T>` or no pointer at all and itself allocate the object. It will delete the wrapped object when destroyed. A *non-owning* `Wrapper` may be used to register with the data repository objects that are not directly heap-allocated, for example data members of other objects. It will take a raw pointer as input and not delete the wrapped object when destroyed.

Attributes

Each instance of `Wrapper` has a set of attributes that control its function in the data repository. These attributes are:

- **InputFlags**

A strongly typed enum that defines the relationship between the `Wrapper` and the XML input. Possible values are:

Value	Explanation
FALSE	Data is not read from XML input (default).
OPTIONAL	Data is read from XML if an attribute matching <code>Wrapper</code> 's name is found.
REQUIRED	Data is read from XML and an error is raised if the attribute is not found.

Other values of `InputFlags` enumeration are reserved for `Group` objects.

Note: A runtime error will occur when attempting to read from XML a wrapped type `T` that does not have `operator>>` defined.

- **RestartFlags**

Enumeration that describes how the `Wrapper` interacts with restart files.

Value	Explanation
NO_WRITE	Data is not written into restart files.
WRITE	Data is written into restart files but not read upon restart.
WRITE_AND_READ	Data is both written and read upon restart (default).

Note: A runtime error will occur when attempting to write a wrapped type `T` that does not support buffer packing. Therefore, when registering custom types (i.e. not a basic C++ type or an `LvArray` container) we recommend setting the flag to `NO_WRITE`. A future documentation topic will explain how to extend buffer packing capabilities to custom user-defined types.

- **PlotLevel**

Enumeration that describes how the `Wrapper` interacts with plot (visualization) files.

Value	Explanation
LEVEL_0	Data always written to plot files.
LEVEL_1	Data written to plot when <code>plotLevel</code> ≥ 1 is specified in input.
LEVEL_2	Data written to plot when <code>plotLevel</code> ≥ 2 is specified in input.
LEVEL_3	Data written to plot when <code>plotLevel</code> ≥ 3 is specified in input.
NO_PLOT	Data never written to plot files.

Note: Only data stored in `LvArray`'s `Array<T>` containers is currently written into plot files.

Default Values

Wrapper supports setting a default value for its wrapped object. The default value is used if a wrapper with `InputFlags::OPTIONAL` attribute does not match an attribute in the input file. For *LvArray* containers it is also used as a default value for new elements upon resizing the container.

Default value can be set via one of the following two methods:

- `setDefaultValue` sets the default value but does not affect the actual value stored in the wrapper.
- `setApplyDefaultValue` sets the default value *and* applies it to the stored value.

Note: A runtime error is raised if a default value is not set for a wrapper with `InputFlags::OPTIONAL` attribute.

The type `DefaultValue<T>` is used to store the default value for the wrapper.

Todo: `DefaultValue` is actually not a type but an alias for another internal struct. As such, it cannot currently be specialized for a user's custom type.

API documentation

Wrapper

ObjectCatalog

The “ObjectCatalog” is a collection of classes that acts as a statically initialized factory. It functions in a similar manner to a classic [factory method](#), except that there is no maintained list of derived objects that is required to create new objects. In other words, there is no case-switch/if-elseif block to maintain. Instead, the `ObjectCatalog` creates a “catalog” of derived objects using a `std::unordered_map`. The “catalog” is filled when new types are declared through the declaration of a helper class named `CatalogEntryConstructor`.

The key functional features of the “ObjectCatalog” concept may be summarized as:

- Anonymous addition of new objects to the catalog. Because we use a statically initialized singleton map object to store the catalog, no knowledge of the contents of the catalog is required in the main code. Therefore if a proprietary/sensitive catalog entry is desired, it is only required that the object definition be outside of the main repository and tied into the build system through some non-specific mechanism (i.e. a link in the `src/externalComponents` directory) and the catalog entry will be registered in the catalog without sharing any knowledge of its existence. Then a proprietary input file may refer to the object to call for its creation.
- Zero maintenance catalog. Again, because we use a singleton map to store the catalog, there is no updating of code required to add new entries into the catalog. The only modifications required are the actual source files of the catalog entry, as described in the [Usage](#) section below.

Implementation Details

There are three key objects that are used to provide the `ObjectCatalog` functionality.

CatalogInterface

The `CatalogInterface` class provides the base definitions and interface for the `ObjectCatalog` concept. It is templated on the common base class of all derived objects that are creatable by the “ObjectCatalog”. In addition, `CatalogInterface` is templated on a variadic parameter pack that allows for an arbitrary constructor argument list as shown in the declaration shown below:

```
template< typename BASETYPE, typename ... ARGS >
class CatalogInterface
```

The `CatalogInterface` also defines the actual catalog type using the template arguments:

```
typedef std::unordered_map< std::string,
                           std::unique_ptr< CatalogInterface< BASETYPE, ARGS... > >
    > CatalogType;
```

The `CatalogInterface::CatalogType` is a `std::unordered_map` with a string “key” and a value type that is a pointer to the `CatalogInterface` that represents a specific combination of `BASETYPE` and constructor arguments.

After from setting up and populating the catalog, which will be described in the “Usage” section, the only interface with the catalog will typically be when the `Factory()` method is called. The definition of the method is given as:

```
static std::unique_ptr< BASETYPE > factory( std::string const & objectTypeName,
                                           ARGS... args )
{
    // We stop the simulation if the product is not found
    if( !hasKeyName( objectTypeName ) )
    {
        std::list< typename CatalogType::key_type > keys = getKeys();
        string const tmp = stringutilities::join( keys.cbegin(), keys.cend(), ",\n" );

        string errorMsg = "Could not find keyword \"" + objectTypeName + "\" in this_
        context. ";
        errorMsg += "Please be sure that all your keywords are properly spelled or that_
        input file parameters have not changed.\n";
        errorMsg += "All available keys are: [\n" + tmp + "\n]";
        GEOSX_ERROR( errorMsg );
    }

    // We also stop the simulation if the builder is not here.
    CatalogInterface< BASETYPE, ARGS... > const * builder = getCatalog().at(
        objectTypeName ).get();
    if( builder == nullptr )
    {
        const string errorMsg = "\"" + objectTypeName + "\" could be found. But the_
        builder is invalid.\n";
        GEOSX_ERROR( errorMsg );
    }

    return builder->allocate( args ... );
}
```

It can be seen that the static `Factory` method is simply a wrapper that calls the virtual `Allocate` method on a the catalog which is returned by `getCatalog()`. The usage of the `Factory` method will be further discussed in the *Usage* section.

Note: The method for organizing constructing new objects relies on a common constructor list between the derived type and the `BASETYPE`. This means that there is a single catalog for each combination of `BASETYPE` and the variadic parameter pack representing the constructor arguments. In the future, we can investigate removing this restriction and allowing for construction of a hierarchy of objects with an arbitrary constructor parameter list.

CatalogEntry

The `CatalogEntry` class derives from `CatalogInterface` and adds the a `TYPE` template argument to the arguments of the `CatalogInterface`.

```
template< typename BASETYPE, typename TYPE, typename ... ARGS >
class CatalogEntry final : public CatalogInterface< BASETYPE, ARGS... >
```

The `TYPE` template argument is the type of the object that you would like to be able to create with the “ObjectCatalog”. `TYPE` must be derived from `BASETYPE` and have a constructor that matches the variadic parameter pack specified in the template parameter list. The main purpose of the `CatalogEntry` is to override the `CatalogInterface::Allocate()` virtual function s.t. when key is retrieved from the catalog, then it is possible to create a new `TYPE`. The `CatalogEntry::Allocate()` function is a simple creation of the underlying `TYPE` as shown by its definition:

```
virtual std::unique_ptr< BASETYPE > allocate( ARGS... args ) const override
{
    #if OBJECTCATALOGVERBOSE > 0
        GEOSX_LOG( "Creating type " << LvArray::system::demangle( typeid(TYPE).name() )
                    << " from catalog of " << LvArray::system::demangle(
->typeid(BASETYPE).name() ) );
    #endif
    #if ( __cplusplus >= 201402L )
        return std::make_unique< TYPE >( args ... );
    #else
        return std::unique_ptr< BASETYPE >( new TYPE( args ... ) );
    #endif
}
```

CatalogEntryConstructor

The `CatalogEntryConstructor` is a helper class that has a sole purpose of creating a new `CatalogEntry` and adding it to the catalog. When a new `CatalogEntryConstructor` is created, a new `CatalogEntry` entry is created and inserted into the catalog automatically.

Usage

Creating A New Catalog

When creating a new “ObjectCatalog”, it typically is done within the context of a specific `BASETYPE`. A simple example of a class hierarchy in which we would like to use the “ObjectCatalog” to use to generate new objects is given in the unit test located in `testObjectCatalog.cpp`.

The base class for this example is defined as:

```

class Base
{
public:
    Base( int & junk, double const & junk2 )
    {
        GEOSX_LOG( "calling Base constructor with arguments ("<<junk<<" "<<junk2<<")" );
    }

    virtual ~Base()
    {
        GEOSX_LOG( "calling Base destructor" );
    }

    using CatalogInterface = dataRepository::CatalogInterface< Base, int &, double_
↪const & >;
    static CatalogInterface::CatalogType & getCatalog()
    {
        static CatalogInterface::CatalogType catalog;
        return catalog;
    }

    virtual string getCatalogName() = 0;
};

```

There a couple of things to note in the definition of Base:

- Base has a convenience alias to use in place of the fully templated CatalogInterface name.
- Base defines a getCatalog() function that returns a static instantiation of a CatalogInterface::CatalogType. The CatalogInterface::getCatalog() function actually calls this function within the base class. This means that the base class actually owns the catalog, and the CatalogInterface is only operating on that Base::getCatalog(), and that the definition of this function is required.

Adding A New Type To The Catalog

Once a Base class is defined with the required features, the next step is to add a new derived type to the catalog defined in Base. There are three requirements for the new type to be registered in the catalog:

- The derived type must have a constructor with the arguments specified by the variadic parameter pack specified in the catalog.
- There must be a static function static string catalogName() that returns the name of the type that will be used to as keyname when it is registered Base's catalog.
- The new type must be registered with the catalog held in Base. To accomplish this, a convenience macro REGISTER_CATALOG_ENTRY() is provided. The arguments to this macro are the name type of Base, the type of the derived class, and then the variadic pack of constructor arguments.

A pair of of simple derived class that have the required methods are used in the unit test.

```

class Derived1 : public Base
{
public:
    Derived1( int & junk, double const & junk2 ):
        Base( junk, junk2 )
    {

```

(continues on next page)

(continued from previous page)

```

    GEOSX_LOG( "calling Derived1 constructor with arguments ("<<junk<<" "<<junk2<<") " );
    ↪);
}

~Derived1()
{
    GEOSX_LOG( "calling Derived1 destructor" );
}
static string catalogName() { return "derived1"; }
string getCatalogName() { return catalogName(); }

};
REGISTER_CATALOG_ENTRY( Base, Derived1, int &, double const & )

```

```

class Derived2 : public Base
{
public:
    Derived2( int & junk, double const & junk2 ):
        Base( junk, junk2 )
    {
        GEOSX_LOG( "calling Derived2 constructor with arguments ("<<junk<<" "<<junk2<<") " );
        ↪);
    }

    ~Derived2()
    {
        GEOSX_LOG( "calling Derived2 destructor" );
    }
    static string catalogName() { return "derived2"; }
    string getCatalogName() { return catalogName(); }

};
REGISTER_CATALOG_ENTRY( Base, Derived2, int &, double const & )

```

Allocating A New Object From The Catalog

The test function in the unit test shows how to allocate a new object of one of the derived types from Factory method. Note the call to Factory is scoped by Base::CatalogInterface, which is an alias to the full templated instantiation of CatalogInterface. The arguments for Factory

```

TEST( testObjectCatalog, testRegistration )
{
    GEOSX_LOG( "EXECUTING MAIN" );
    int junk = 1;
    double junk2 = 3.14;

    // allocate a new Derived1 object
    std::unique_ptr< Base >
    derived1 = Base::CatalogInterface::factory( "derived1", junk, junk2 );

    // allocate a new Derived2 object
    std::unique_ptr< Base >
    derived2 = Base::CatalogInterface::factory( "derived2", junk, junk2 );
}

```

(continues on next page)

(continued from previous page)

```

EXPECT_STREQ( derived1->getCatalogName().c_str(),
              Derived1::catalogName().c_str() );

EXPECT_STREQ( derived2->getCatalogName().c_str(),
              Derived2::catalogName().c_str() );
GEOSX_LOG( "EXITING MAIN" );
}

```

The unit test creates two new objects of type `Derived1` and `Derived2` using the catalogs `Factory` method. Then the test checks to see that the objects that were created are of the correct type. This unit test has some extra output to screen to help with understanding of the sequence of events. The result of running this test is:

```

$ tests/testObjectCatalog
Calling constructor for CatalogEntryConstructor< Derived1 , Base , ... >
Calling constructor for CatalogInterface< Base , ... >
Calling constructor for CatalogEntry< Derived1 , Base , ... >
Registered Base catalog component of derived type Derived1 where
↳Derived1::catalogName() = derived1
Calling constructor for CatalogEntryConstructor< Derived2 , Base , ... >
Calling constructor for CatalogInterface< Base , ... >
Calling constructor for CatalogEntry< Derived2 , Base , ... >
Registered Base catalog component of derived type Derived2 where
↳Derived2::catalogName() = derived2
Running main() from gtest_main.cc
[=====] Running 1 test from 1 test case.
[-----] Global test environment set-up.
[-----] 1 test from testObjectCatalog
[ RUN      ] testObjectCatalog.testRegistration
EXECUTING MAIN
Creating type Derived1 from catalog of Base
calling Base constructor with arguments (1 3.14)
calling Derived1 constructor with arguments (1 3.14)
Creating type Derived2 from catalog of Base
calling Base constructor with arguments (1 3.14)
calling Derived2 constructor with arguments (1 3.14)
EXITING MAIN
calling Derived2 destructor
calling Base destructor
calling Derived1 destructor
calling Base destructor
[ OK ] testObjectCatalog.testRegistration (0 ms)
[-----] 1 test from testObjectCatalog (0 ms total)

[-----] Global test environment tear-down
[=====] 1 test from 1 test case ran. (0 ms total)
[ PASSED ] 1 test.
Calling destructor for CatalogEntryConstructor< Derived2 , Base , ... >
Calling destructor for CatalogEntryConstructor< Derived1 , Base , ... >
Calling destructor for CatalogEntry< Derived2 , Base , ... >
Calling destructor for CatalogInterface< Base , ... >
Calling destructor for CatalogEntry< Derived1 , Base , ... >
Calling destructor for CatalogInterface< Base , ... >

```

In the preceding output, it is clear that the static catalog in `Base::getCatalog()` is initialized prior the execution of `main`, and destroyed after the completion of `main`. In practice, there have been no indicators of problems due to the use of a statically initialized/deinitialized catalog.

6.2.2 XML Input

In this document, you will learn how GEOSX classes interact with external information parsed from XML files, and how to add a new XML block that can be interpreted by GEOSX. Flow solvers and relative permeability are used as examples.

GEOSX data structure overview

Group : the base class of GEOSX

All GEOSX classes derive from a base class called `dataRepository::Group`. The `Group` class provides a way to organize all GEOSX objects in a filesystem-like structure. One could think of `Group`s as *file folders* that can bear data (stored in `Wrapper`s), have a parent folder (another `Group`), and have possibly multiple subfolders (referred to as the subgroups). Below, we briefly review the data members of the `Group` class that are essential to understand the correspondence between the GEOSX data structure and the XML input. For more details, we refer the reader to the extensive documentation of the *Data Repository*, including the *Group* class documentation.

In the code listing below, we see that each `Group` object is at minimum equipped with the following member properties:

- A pointer to the parent `Group` called `m_parent` (member classes are prefixed by `m_`),
- The `Group`'s own data, stored for flexibility in an array of generic data `Wrapper`s called `m_wrappers`,
- A map of one or many children (also of type `Group`) called `m_subGroups`.
- The `m_size` and `m_capacity` members, that are used to set the size and capacity of any objects contained.
- The name of the `Group`, stored as a string in `m_name`. This name can be seen as the object unique ID.

```

/// The parent Group that contains "this" Group in its "sub-Group" collection.
Group * m_parent = nullptr;

/// Specification that this group will have the same m_size as m_parent.
integer m_sizedFromParent;

/// The container for the collection of all wrappers contained in "this" Group.
wrapperMap m_wrappers;

/// The container for the collection of all sub-groups contained in "this" Group.
subGroupMap m_subGroups;

/// The size/length of this Group...and all Wrapper<> that are are specified to_
↳ have the same size as their
/// owning group.
indexType m_size;

/// The capacity for wrappers in this group...and all Wrapper<> that are specified_
↳ to have the same size as their
/// owning group.
indexType m_capacity;

/// The name/key of this Group in its parent collection of sub-Groups.
string m_name;

/// Verbosity flag for group logs
integer m_logLevel;
//END_SPHINX_INCLUDE_02

```

(continues on next page)

```
/// Restart flag for this group... and subsequently all wrappers in this group.
```

[Source: `src/coreComponents/dataRepository/Group.hpp`]

A few words about the ObjectCatalog

What is an ObjectCatalog and why do we need it?

Some classes need external information (physical and/or algorithmic parameters for instance) provided by the user to be instantiated. This is the case when the `m_input_flags` data member of one of the `Group's Wrapper's` has an entry set to `REQUIRED` (we will illustrate this below). In this situation, the required information must be supplied in the XML input file, and if it is absent, an error is raised by GEOSX.

To connect the external (XML) and internal (C++) data structures, GEOSX uses an **ObjectCatalog** that maps keys (of type `string`) to the corresponding classes (one unique key per mapped class). These string keys, referred to as `catalogName's`, are essential to transfer the information from the XML file to the factory functions in charge of object instantiation (see below).

What is a CatalogName?

The `catalogName` of an object is a *key* (of type `string`) associated with this object's class. On the one hand, in the XML file, the key is employed by the user as an XML tag to specify the type of object (e.g., the type of solver, constitutive model, etc) to create and use during the simulation. On the other hand, internally, the key provides a way to access the appropriate factory function to instantiate an object of the desired class.

Most of the time, the `catalogName` and the C++ class name are identical. This helps make the code easier to debug and allows the XML/C++ correspondence to be evident. But strictly speaking, the `catalogName` can be anything, as long as it refers uniquely to a specific class. The `catalogName` must not be confused with the object's *name* (`m_name` is a data member of the class that stores the object's unique ID, not its class key). You can have several objects of the same class and hence the same `catalogName`, but with different names (i.e. unique ID): several fluid models, several solvers, etc.

How can I add my new externally-accessible class to the ObjectCatalog?

Let us consider a flow solver class derived from `FlowSolverBase`, that itself is derived from `SolverBase`. To instantiate and use this solver, the developer needs to make the derived flow solver class reachable from the XML file, via an XML tag. Internally, this requires adding the derived class information to `ObjectCatalog`, which is achieved with two main ingredients: 1) a `CatalogName()` method in the class that lets GEOSX know *what* to search for in the internal `ObjectCatalog` to instantiate an object of this class, 2) a macro that specifies *where* to search in the `ObjectCatalog`.

1. To let GEOSX know what to search for in the catalog to instantiate an object of the derived class, the developer must equip the class with a `CatalogName()` method that returns a `string`. In this document, we have referred to this returned `string` as the object's `catalogName`, but in fact, the method `CatalogName()` is what matters since the `ObjectCatalog` contains all the `CatalogName()` return values. Below, we illustrate this with the `CompositionalMultiphaseFlow` solver. The first code listing defines the class name, which in this case is the same as the `catalogName` shown in the second listing.

```
/**
 * @class CompositionalMultiphaseBase
 *
 * A compositional multiphase solver
 */
class CompositionalMultiphaseBase : public FlowSolverBase
{
```

[Source: *src/coreComponents/physicsSolvers/fluidFlow/CompositionalMultiphaseBase.hpp*]

2. To let GEOSX know where to search in the `ObjectCatalog`, a macro needs to be added at the end of the .cpp file implementing the class. This macro (illustrated below) must contain the type of the base class (in this case, `SolverBase`), and the name of the derived class (continuing with the example used above, this is `CompositionalMultiphaseFlow`). As a result of this construct, the `ObjectCatalog` is not a flat list of `strings` mapping the C++ classes. Instead, the `ObjectCatalog` forms a tree that reproduces locally the structure of the class diagram, from the base class to the derived classes.

```
REGISTER_CATALOG_ENTRY( SolverBase, CompositionalMultiphaseFVM, string const &, Group_
↳ * const )
```

[Source: *src/coreComponents/physicsSolvers/fluidFlow/CompositionalMultiphaseFVM.cpp*]

Summary: All GEOSX objects form a filesystem-like structure. If an object needs to be accessible externally, it must be registered in the `ObjectCatalog`. This is done by adding `CatalogName()` method that returns a `string` key to the object's class, and by adding the appropriate macro. The catalog has the same tree structure as the class diagram.

Registration: parsing XML input files to instantiate GEOSX objects

In this section, we describe with more details the connection between **internal GEOSX objects** and **external XML tags** parsed from parameter files. We call this process *Registration*. The registration process works in three steps:

1. The XML document is parsed. Each time a new XML tag is found, the current local scope of the `ObjectCatalog` is inspected. The goal is to find a `catalogName` string that matches the XML tag.
2. If it is the case (the current local scope of the `ObjectCatalog` contains a `catalogName` identical to the XML tag), then the code creates a new instance of the class that the `catalogName` refers to. This new object is inserted in the `Group` tree structure at the appropriate location, as a subgroup.
3. By parsing the XML attributes of the tag, the new object properties are populated. Some checks are performed to ensure that the data supplied is conform, and that all the required information is present.

Let's look at this process in more details.

Creating a new object and giving it a Catalog name

Consider again that we are registering a flow solver deriving from `FlowSolverBase`, and assume that this solver is called `CppNameOfMySolver`. This choice of name is not recommended (we want names that reflect what the solver does!), but for this particular example, we just need to know that this name is the class name inside the C++ code.

To specify parameters of this new solver from an XML file, we need to be sure that the XML tag and the `catalogName` of the class are identical. Therefore, we equip the `CppNameOfMySolver` class with a `CatalogName()` method that returns the solver `catalogName` (=XML name). Here, this method returns the string "`XmlNameOfMySolver`".

We have deliberately distinguished the class name from the catalog/XML name for the sake of clarity in this example. It is nevertheless a best practice to use the same name for the class and for the `catalogName`. This is the case below for the existing `CompositionalMultiphaseFVM` class.

```
class CompositionalMultiphaseFVM : public CompositionalMultiphaseBase
{
```

```

/**
 * @brief name of the solver in the object catalog
 * @return string that contains the catalog name to generate a new object through_
↳ the object catalog.
 */
static string catalogName() { return "CompositionalMultiphaseFVM"; }

```

[Source: src/coreComponents/physicsSolvers/fluidFlow/CompositionalMultiphaseFVM.hpp]

Parsing XML and searching the ObjectCatalog in scope

Now that we have implemented a `CatalogName()` method returning a specific key (of type `string`), we can have a block in our XML input file with a tag that corresponds to the `catalogName` “`XmlNameOfMySolver`”. This is how the XML block would look like.

```

<Problem>
  <Solvers
    gravityVector="{ 0.0, 0.0, -9.81 }">
      <XmlNameOfMySolver name="nameOfThisSolverInstance"
        verboseLevel="1"
        gravityFlag="1"
        temperature="297.15" />
        <LinearSolverParameters newtonTol="1.0e-6"
          maxIterNewton="15"
          useDirectSolver="1"/>
      </XmlNameOfMySolver>
    </Solvers>
  </Problem>

```

Here, we see that the XML structure defines a parent node “Problem”, that has (among many others) a child node “Solvers”. In the “Solvers” block, we have placed the new solver block as a child node of the “Solvers” block with the XML tag corresponding to the `catalogName` of the new class. We will see in details next how the GEOSX internal structure constructed from this block mirrors the XML file structure.

Instantiating the new solver

Above, we have specified an XML block with the tag “`XmlNameOfMySolver`”. Now, when reading the XML file and encountering an “`XmlNameOfMySolver`” solver block, we add a new instance of the class `CppNameOfMySolver` in the filesystem structure as explained below.

We saw that in the XML file, the new solver block appeared as child node of the XML block “Solvers”. The internal construction mirrors this XML structure. Specifically, the new object of class `CppNameOfMySolver` is registered as a subgroup (to continue the analogy used so far, as a subfolder) of its parent Group, the class `PhysicsSolverManager` (that has a `catalogName` “Solvers”). To do this, the method `CreateChild` of the `PhysicsSolverManager` class is used.

```

// Variable values in this example:
// -----
// childKey = "XmlNameOfMySolver" (string)
// childName = "nameOfThisSolverInstance" (string)
// SolverBase::CatalogInterface = the Catalog attached to the base Solver class
// hasKeyName = bool method to test if the childKey string is present in the Catalog
// registerGroup = method to create a new instance of the solver and add it to the_
↳ group tree

```

```

Group * PhysicsSolverManager::createChild( string const & childKey, string const &
↳childName )
{
    Group * rval = nullptr;
    if( SolverBase::CatalogInterface::hasKeyName( childKey ) )
    {
        GEOSX_LOG_RANK_0( "Adding Solver of type " << childKey << ", named " << childName_
↳);
        rval = &registerGroup( childName,
                               SolverBase::CatalogInterface::factory( childKey, childName,
↳this ) );
    }
    return rval;
}

```

[Source: *src/coreComponents/physicsSolvers/PhysicsSolverManager.cpp*]

In the code listing above, we see that in the `PhysicsSolverManager` class, the `ObjectCatalog` is searched to find the catalogName “`CompositionalMultiphaseFlow`” in the scope of the `SolverBase` class. Then, the factory function of the base class `SolverBase` is called. The catalogName (stored in `childKey`) is passed as an argument of the factory function to ensure that it instantiates an object of the desired derived class.

As explained above, this is working because 1) the XML tag matches the catalogName of the `CompositionalMultiphaseFlow` class and 2) a macro is placed at the end of the .cpp file implementing the `CompositionalMultiphaseFlow` class to let the `ObjectCatalog` know that `CompositionalMultiphaseFlow` is a derived class of `SolverBase`.

Note that several instances of the same type of solver can be created, as long as they each have a different name.

Filling the objects with data (wrappers)

After finding and placing the new solver `Group` in the filesystem hierarchy, properties are read and stored. This is done by registering *data wrappers*. We refer to the documentation of the [Data Repository](#) for additional details about the `Wrapper`s. The method used to do that is called `registerWrapper` and is placed in the class constructor when the data is required in the XML file. Note that some properties are registered at the current (derived) class level, and other properties can also be registered at a base class level.

Here, the only data (=wrapper) that is defined at the level of our `CppNameOfMySolver` class is temperature, and everything else is registered at the base class level. We register a property of temperature, corresponding to the member class `m_temperature` of `CppNameOfMySolver`. The registration also checks if a property is required or optional (here, it is required), and provides a brief description that will be used in the auto-generated code documentation.

```

this->registerWrapper( viewKeyStruct::temperatureString(), &m_temperature ).
    setInputFlag( InputFlags::REQUIRED ).
    setDescription( "Temperature" );
//END_SPHINX_INCLUDE_00

```

[Source: *src/coreComponents/physicsSolvers/fluidFlow/CompositionalMultiphaseBase.cpp*]

This operation is done recursively if XML tags are nested.

To summarize:

- Every class in GEOSX derive from a `Group` in a filesystem-like structure. A `Group` must have a parent `Group`, can have data (in `Wrapper`s), and can have one or many children (the subgroups). There is an `ObjectCatalog` in which the classes derived from `Group` are identified by a key called the `catalogName`.

- When parsing XML input files, GEOSX inspects each object's scope in the `ObjectCatalog` to find classes with the same `catalogName` as the XML tag. Once it finds an XML tag in the `ObjectCatalog`, it registers it inside the filesystem structure.
- In the registration process, properties from the XML file are parsed and used to allocate member data `Wrapper s` and fully instantiate the `Group` class.
- If XML tags are nested, subgroups are allocated and processed in a nested manner.

The correspondence between XML and class hierarchy is thus respected, and the internal object hierarchy mirrors the XML structure.

Example: adding a new relative permeability model

This example is taken from the class `BrooksCoreyRelativePermeability`, derived from `RelativePermeabilityBase`.

Implement a `CatalogName` function (.hpp):

As explained above we add the class to the `ObjectCatalog` in two steps. First we implement the `CatalogName` function:

```
static string catalogName() { return "BrooksCoreyRelativePermeability"; }
```

[source: `src/coreComponents/constitutive/relativePermeability/BrooksCoreyRelativePermeability.hpp`]

Then in the .cpp file we add the macro to register the catalog entry:

```
REGISTER_CATALOG_ENTRY( ConstitutiveBase, BrooksCoreyRelativePermeability, string_  
↳const &, Group * const )
```

[source: `src/coreComponents/constitutive/relativePermeability/BrooksCoreyRelativePermeability.cpp`]

Now every time a “`BrooksCoreyRelativePermeability`” string is encountered inside a `RelativePermeability` catalog, we will instantiate a class `BrooksCoreyRelativePermeability`.

Declare the `Wrapper s` keys (.hpp):

When attaching properties (i.e. data `Wrapper s`) to a class, a similar registration process must be done. Every property is accessed through its `ViewKey` namespace. In this namespace, we define strings that correspond to the tags of XML attributes of the “`BrooksCoreyRelativePermeability`” block.

```
struct viewKeyStruct : RelativePermeabilityBase::viewKeyStruct  
{  
    static constexpr char const * phaseMinVolumeFractionString() { return  
↳"phaseMinVolumeFraction"; }  
    static constexpr char const * phaseRelPermExponentString() { return  
↳"phaseRelPermExponent"; }  
    static constexpr char const * phaseRelPermMaxValueString() { return  
↳"phaseRelPermMaxValue"; }  
    static constexpr char const * volFracScaleString() { return "volFracScale"; }  
  
    dataRepository::ViewKey phaseMinVolumeFraction = { phaseMinVolumeFractionString()_  
↳};  
    dataRepository::ViewKey phaseRelPermExponent    = { phaseRelPermExponentString() };
```

(continues on next page)

(continued from previous page)

```
dataRepository::ViewKey phaseRelPermMaxValue    = { phaseRelPermMaxValueString() };
} vieKeysBrooksCoreyRelativePermeability;
```

[source: *src/coreComponents/constitutive/relativePermeability/BrooksCoreyRelativePermeability.hpp*]

Declare data members (.hpp):

The data members are defined in the class. They will ultimately contain the data read from the XML file (other data members not read from the XML file can also exist).

```
array1d< real64 > m_phaseMinVolumeFraction;
array1d< real64 > m_phaseRelPermExponent;
array1d< real64 > m_phaseRelPermMaxValue;

real64 m_volFracScale;
```

[source: *src/coreComponents/constitutive/relativePermeability/BrooksCoreyRelativePermeability.hpp*]

Implement the data registration process (registerWrapper):

The registration process done in the class constructor puts everything together. It connects the attributes values in the XML file to class member data. For instance, in the listing below, the first `registerWrapper` call means that we want to read in the XML file the attribute value corresponding to the attribute tag “`phaseMinVolumeFraction`” defined in the .hpp file, and that we want to store the read values into the `m_phaseMinVolumeFraction` data members. We see that this input is not required. If it is absent from the XML file, the default value is used instead. The short description that completes the registration will be added to the auto-generated documentation.

```
BrooksCoreyRelativePermeability::BrooksCoreyRelativePermeability( string const & name,
                                                                    Group * const _
↳parent )
    : RelativePermeabilityBase( name, parent )
{
    registerWrapper( viewKeyStruct::phaseMinVolumeFractionString(), &m_
↳phaseMinVolumeFraction ).
        setApplyDefaultValue( 0.0 ).
        setInputFlag( InputFlags::OPTIONAL ).
        setDescription( "Minimum volume fraction value for each phase" );

    registerWrapper( viewKeyStruct::phaseRelPermExponentString(), &m_
↳phaseRelPermExponent ).
        setApplyDefaultValue( 1.0 ).
        setInputFlag( InputFlags::OPTIONAL ).
        setDescription( "MinimumRel perm power law exponent for each phase" );

    registerWrapper( viewKeyStruct::phaseRelPermMaxValueString(), &m_
↳phaseRelPermMaxValue ).
        setApplyDefaultValue( 0.0 ).
        setInputFlag( InputFlags::OPTIONAL ).
        setDescription( "Maximum rel perm value for each phase" );

    registerWrapper( viewKeyStruct::volFracScaleString(), &m_volFracScale ).
```

(continues on next page)

(continued from previous page)

```

    setApplyDefaultValue( 1.0 ).
    setDescription( "Factor used to scale the phase capillary pressure, defined as:
    ↪one minus the sum of the phase minimum volume fractions." );

```

[source: *src/coreComponents/constitutive/relativePermeability/BrooksCoreyRelativePermeability.cpp*]

The XML block

We are ready to use the relative permeability model in GEOSX. The corresponding XML block (child node of the “Constitutive” block) reads:

```

<Constitutive>
  <BrooksCoreyBakerRelativePermeability name="relperm"
    phaseNames="{oil, gas, water}"
    phaseMinVolumeFraction="{0.05, 0.05, 0.05}"
    waterOilRelPermExponent="{2.5, 1.5}"
    waterOilRelPermMaxValue="{0.8, 0.9}"
    gasOilRelPermExponent="{3, 3}"
    gasOilRelPermMaxValue="{0.4, 0.9}"/>
</Constitutive>

```

With this construct, we instruct the `ConstitutiveManager` class (whose `catalogName` is “Constitutive”) to instantiate a subgroup of type `BrooksCoreyRelativePermeability`. We also fill the data members of the values that we want to use for the simulation. For a simulation with multiple regions, we could define multiple relative permeability models in the “Constitutive” XML block (yielding multiple `relperm` subgroups in GEOSX), with a unique name attribute for each model.

For more examples on how to contribute to GEOSX, please read [Adding a new Physics Solver](#)

Input Schema Generation

A schema file is a useful tool for validating input .xml files and constructing user-interfaces. Rather than manually maintaining the schema during development, GEOSX is designed to automatically generate one by traversing the documentation structure.

To generate the schema, run GEOSX with the input, schema, and the (optional) `schema_level` arguments, i.e.: `geosx -i input.xml -s schema.xsd`. There are two ways to limit the scope of the schema:

1. Setting the verbosity flag for an object in the documentation structure. If the schema-level argument is used, then only objects (and their children) and attributes with `(verbosity < schema-level)` will be output.
2. By supplying a limited input xml file. When GEOSX builds its data structure, it will only include objects that are listed within the xml (or those that are explicitly appended when those objects are initialized). The code will add all available *attributes* for these objects to the schema.

To take advantage of this design it is necessary to use the automatic xml parsing approach that relies upon the documentation node. If values are read in manually, then the schema can not be used to validate xml those inputs.

Note: the lightweight xml parser that is used in GEOSX cannot be used to validate inputs with the schema directly. As such, it is necessary to use an external tool for validation, such as the `geosx_tools` python module.

6.2.3 Working with data in GEOSX

In GEOSX, data is typically registered in the [Data Repository](#). This allows for the writing/reading of data to/from restart and plot files. Any object that derives from [Group](#) may have data registered on it through the methods de-

scribed in *Group*. Similarly, accessing data from outside the scope of an object is possible through one of the various `Group::get()`. Of course, for temporary data that does not need to persist between cycles, or across physics packages, you may simply define member or local variables which will not be registered with the *Data Repository*.

Working with data on the Mesh objects

the mesh objects in GEOSX such as the `FaceManager` or `NodeManager`, are derived from `ObjectManagerBase`, which in turn derives from *Group*. The important distinction is that `ObjectManagerBase` contains various members that are useful when defining mesh object managers. When considering data that is attached to a mesh object, we group the data into two categories:

- **Intrinsic** data is data that is required to describe the object. For instance, to define a `Node`, the `NodeManager` contains an array of positions corresponding to each `Node` it contains. Thus the `ReferencePosition` is Intrinsic data. Intrinsic data is almost always a member of the mesh object, and is registered on the mesh object in the constructor of mesh object itself.
- **Extrinsic** data is data that is not required to define the object. For instance, a physics package may request that a `Velocity` value be stored on the nodes. Appropriately the data will be registered on the `NodeManager`. However, this data is not required to define a `Node`, and is viewed as Extrinsic. Extrinsic data is never a member of the mesh object, and is typically registered on the mesh object outside of the definition of the mesh object (i.e. from a physics solver).

Registering Intrinsic data on a Mesh Object

As mentioned above, Intrinsic data is typically a member of the mesh object, and is registered in the constructor of the mesh Object. Taking the `NodeManager` and the `referencePosition` as an example, we point out that the reference position is actually a member in the `NodeManager`.

```
/**
 * @brief Get the mutable reference position array. This table will contain all the
 *↪node coordinates.
 * @return reference position array
 */
array2d< real64, nodes::REFERENCE_POSITION_PERM > & referencePosition() { return m_
↪referencePosition; }

/**
 * @brief Provide an immutable arrayView of the reference position. This table will
 *↪contain all the node coordinates.
 * @return an immutable arrayView of the reference position.
 */

arrayView2d< real64 const, nodes::REFERENCE_POSITION_USD > referencePosition() const
{ return m_referencePosition; }
```

This member is registered in the constructor for the `NodeManager`.

```
NodeManager::NodeManager( string const & name,
                          Group * const parent ):
    ObjectManagerBase( name, parent ),
    m_referencePosition( 0, 3 )
{
    registerWrapper( viewKeyStruct::referencePositionString(), &m_referencePosition );
}
```

Finally in order to access this data, the `NodeManager` provides explicit accessors.

```

/**
 * @brief Get the mutable reference position array. This table will contain all the
 * node coordinates.
 * @return reference position array
 */
array2d< real64, nodes::REFERENCE_POSITION_PERM > & referencePosition() { return m_
referencePosition; }

/**
 * @brief Provide an immutable arrayView of the reference position. This table will
 * contain all the node coordinates.
 * @return an immutable arrayView of the reference position.
 */

arrayView2d< real64 const, nodes::REFERENCE_POSITION_USD > referencePosition() const
{ return m_referencePosition; }

```

Thus the interface for Intrinsic data is set by the object that it is a part of, and the developer may only access the data through the accessors from outside of the mesh object class scope.

Registering Extrinsic data on a Mesh Object

To register Extrinsic data, there are many ways a developer may proceed. We will use the example of registering a TotalDisplacement on the NodeManager from the SolidMechanics solver. The most general approach is to define a string key and call one of the `Group::registerWrapper()` functions from `SolverBase::RegisterMeshData()`. Then when you want to use the data, you can call `Group::getReference()`. For example this would look something like:

```

void SolidMechanicsLagrangianFEM::RegisterDataOnMesh( Group * const MeshBodies )
{
    for( auto & mesh : MeshBodies->GetSubGroups() )
    {
        NodeManager & nodes = mesh.second->groupCast< MeshBody * >()->getMeshLevel( 0 ).
        getNodeManager();

        nodes.registerWrapper< array2d< real64, nodes::TOTAL_DISPLACEMENT_PERM > >(
        keys::TotalDisplacement ).
        setPlotLevel( PlotLevel::LEVEL_0 ).
        setRegisteringObjects( this->getName() ).
        setDescription( "An array that holds the total displacements on the nodes." ).
        reference().resizeDimension< 1 >( 3 );
    }
}

```

and

```

arrayView2d< real64, nodes::TOTAL_DISPLACEMENT_USD > const & u = nodes.getReference
<array2d<real64, nodes::TOTAL_DISPLACEMENT_PERM >(keys::TotalDisplacement);
... do something with u

```

This approach is flexible and extendible, but is potentially error prone due to its verbosity. Therefore we also provide a more controlled/uniform method by which to register and extract commonly used data on the mesh. The trait approach requires the definition of a traits struct for each data object that will be supported. To apply the trait approach to the example use case shown above, there should be the following definition somewhere in a header file:

```

namespace extrinsicMeshData
{
struct TotalDisplacement
{
    static constexpr auto key = "totalDisplacement";
    using DataType = real64;
    using Type = array2d< DataType, nodes::TOTAL_DISPLACEMENT_PERM >;
    static constexpr DataType defaultValue = 0;
    static constexpr auto plotLevel = dataRepository::PlotLevel::LEVEL_0;

    /// Description of the data associated with this trait.
    static constexpr auto description = "An array that holds the total displacements on_
↪the nodes.";
};
}

```

Then the registration is simplified as follows:

```

void SolidMechanicsLagrangianFEM::RegisterDataOnMesh( Group * const MeshBodies )
{
    for( auto & mesh : MeshBodies->GetSubGroups() )
    {
        NodeManager & nodes = mesh.second->groupCast< MeshBody * >()->getMeshLevel( 0 ).
↪getNodeManager();
        nodes.registerExtrinsicData< extrinsicMeshData::TotalDisplacement >( this->
↪getName() ).resizeDimension< 1 >( 3 );
    }
}

```

And to extract the data, the call would be:

```

arrayView2d< real64, nodes::TOTAL_DISPLACEMENT_USD > const & u = nodes.
↪getExtrinsicData< extrinsicMeshData::TotalDisplacement >();
... do something with u

```

The end result of the trait approach to this example is that the developer has defined a standard specification for TotalDisplacement, which may be used uniformly across the code.

6.2.4 Mesh Hierarchy

In GEOSX, the mesh structure consists of a hierarchy of classes intended to encapsulate data and functionality for each topological type. Each class in the mesh hierarchy represents a distinct topological object, such as a nodes, edges, faces, elements, etc. The mesh data structure is illustrated in an object instantiation hierarchy. The object instantiation hierarchy differs from a “class hierarchy” in that it shows how instantiations of each class relate to each other in the data hierarchy rather than how each class type relates to each other in an inheritance diagram.

To illustrate the mesh hierarchy, we propose to present it along with a model with two regions (Top and Bottom) (Fig. 6.2).

DomainPartition

In Fig. 6.1 the top level object DomainPartition represents a partition of the decomposed physical domain. At this time there is a unique DomainPartition for every MPI rank.

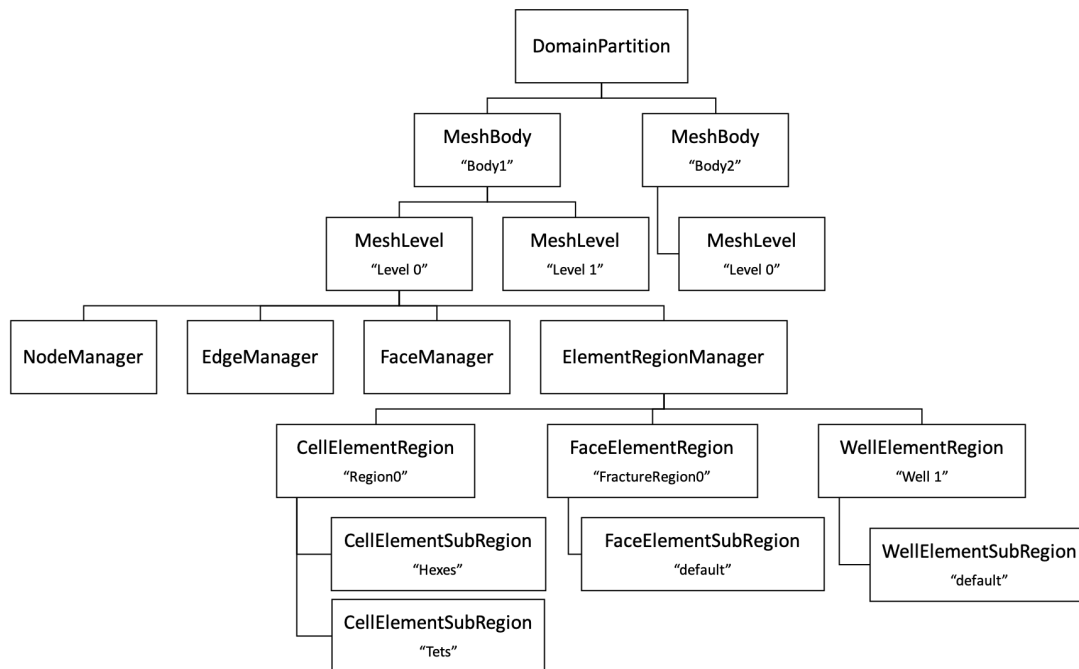


Fig. 6.1: Object Instantiation Hierarchy for the Mesh Objects.




Fig. 6.2: Example of a model with two regions

Note: Hypothetically, there may be more than one `DomainPartition` in cases where the ranks are overloaded. Currently GEOSX does not support overloading multiple `DomainPartition`'s onto a rank, although this may be a future option if its use is properly motivated.

For instance, the model presented as example can be split into two different domains (Fig. 6.3).




Fig. 6.3: Mesh partitioned in two `DomainPartition`

MeshBody

The `MeshBody` represents a topologically distinct mesh body. For instance if a simulation of two separate spheres was required, then one option would be to have both spheres as part of a single mesh body, while another option would be to have each sphere be a individual body.

Note: While not currently utilized in GEOSX, the intent is to have the ability to handle the bodies in a multi-body mesh on an individual basis. For instance, when conducting high resolution crush simulations of granular materials (i.e. sand), it may be advantageous to represent each particle as a `MeshBody`.

MeshLevel

The `MeshLevel` is intended to facilitate the representation of a multi-level discretization of a `MeshBody`.

Note: In current practice, the code utilizes a single `MeshLevel` until such time as we implement a proper multi-level mesh capability. The `MeshLevel` contains the main components that compose a discretized mesh in GEOSX.

Topological Mesh Objects

Each of the “Manager” objects are responsible for holding child objects, data, and providing functionality specific to a single topological object. Each topological object that is used to define a discretized mesh has a “Manager” to allow for simple traversal over the hierarchy, and to provide modular access to data. As such, the `NodeManager` manages data for the “nodes”, the `EdgeManager` manages data for the edges, the `FaceManager` holds data for the faces and the `ElementRegionManager` manages the physical groups within the `MeshLevel` (regions, fractures, wells etc. ...). Additionally each manager contains index maps to the other types objects that are connected to the objects in that manager. For instance, the `FaceManager` contains a downward pointing map that gives the nodes that comprise each face in the mesh. Similarly the `FaceManager` contains an upward pointing map that gives the elements that are connected to a face.

ElementRegionManager

The element data structure is significantly more complicated than the other Managers. While the other managers are “flat” across the `MeshLevel`, the element data structure seeks to provide a hierarchy in order to define groupings of the physical problem, as well as collecting discretization of similar topology. At the top of the element branch of the hierarchy is the `ElementRegionManager`. The `ElementRegionManager` holds a collection of instantiations of `ElementRegionBase` derived classes.

ElementRegion

Conceptually the `ElementRegion` are used to defined regions of the problem domain where a `PhysicsSolver` will be applied.

- The `CellElementRegion` is related to all the polyhedra
- The `FaceElementRegion` is related to all the faces that have physical meaning in the domain, such as fractures and faults. This object should not be mistaken with the `FaceManager`. The `FaceManager` handles all the faces of the mesh, not only the faces of interest.
- The `WellElementRegion` is related to the well geometry.

An `ElementRegion` also has a list of materials allocated at each quadrature point across the entire region. One example of the utility of the `ElementRegion` is the case of the simulation of the mechanics and flow within sub-surface reservoir with an overburden. We could choose to have two `ElementRegion`, one being the reservoir, and one for the overburden. The mechanics solver would be applied to the entire problem, while the flow problem would be applied only to the reservoir region.

Each `ElementRegion` holds some number of `ElementSubRegion`. The `ElementSubRegion` is meant to hold all the element topologies present in an `ElementSubRegion` in their own groups. For instance, for a `CellElementRegion`, there may be one `CellElementSubRegion` for all tetrahedra, one for all hexahedra, one for all wedges and one for all the pyramids (Fig. 6.4).

Now that all the classes of the mesh hierarchy has been described, we propose to adapt the diagram presented in Fig. 6.1 to match with the example presented in Fig. 6.2.

6.2.5 DoF Manager

This will contains a description of the DoF manager in GEOSX.




Fig. 6.4: Model meshed with different cell types

Brief description

The main aim of the Degrees-of-Freedom (DoF) Manager class is to handle all degrees of freedom associated with fields that exist on mesh elements, faces, edges and nodes. It creates a map between local mesh objects and global DoF indices. Additionally, `DofManager` simplifies construction of system matrix sparsity patterns.

Key concepts are locations and connectors. Locations, that can be elements, faces, edges or nodes, represent where the DoF is assigned. For example, a DoF for pressure in a two-point flux approximation will be on a cell (i.e. element), while a displacement DoF for structural equations will be on a node. The counterparts of locations are connectors, that are the geometrical entities that link together different DoFs to create the sparsity pattern. Connectors can be elements, faces, edges, nodes or none. Using the same example as before, connectors will be faces and cells, respectively. The case of a mass matrix, where every element is linked only to itself, is an example when there are no connectors, i.e. these have to be set to none.

DoFs located on a mesh object are owned by the same rank that owns the object in parallel mesh partitioning. Two types of DoF numbering are supported, with the difference only showing in parallel runs of multi-field problems.

- Initially, each field is assigned an independent DoF numbering that starts at 0 and is contiguous across all MPI ranks. Within each rank, locally owned DoFs are numbered sequentially across mesh locations, and within each mesh location (e.g. node) - sequentially according to component number. With this numbering, sparsity patterns can be constructed for individual sub-matrices that represent diagonal/off-diagonal blocks of the global coupled system matrix.
- After all fields have been declared, the user can call `DofManager::reorderByRank()`, which constructs a globally contiguous DoF numbering across all fields. Specifically, all DoFs owned by rank 0 are numbered field-by-field starting from 0, then those on rank 1, etc. This makes global system sparsity pattern compatible with linear algebra packages that only support contiguous matrix rows on each rank. At this point, coupled system matrix sparsity pattern can be constructed.

Thus, each instance of `DofManager` only supports one type of numbering. If both types are required, the user is advised to maintain two separate instances of `DofManager`.

`DofManager` allocates a separate “DOF index” array for each field on the mesh. It is an array of global indices, where

each value represents the first DoF index for that field and location (or equivalently, the row and column offset of that location's equations and variables for the field in the matrix). For example, if index array for a field with 3 components contains the value N, global DoF numbers for that location will be N, N+1, N+2. DoF on ghosted locations have the same indices as on the owning rank. The array is stored under a generated key, which can be queried from the DoF manager, and is typically used in system assembly.

Methods

The main methods of DoF Manager are:

- **setMesh**: sets which portion of the mesh the DoF manager instance is referring to. `domain` identifies the global mesh, while `meshLevelIndex` and `meshBodyIndex` determine the specific level and body, respectively.

```
void setMesh( DomainPartition * const domain,
              localIndex const meshLevelIndex = 0,
              localIndex const meshBodyIndex = 0 );
```

- **addField**: creates a new set of DoF, labeled `field`, with specific location. Default number of components is 1, like for pressure in flux. Default regions is the empty string, meaning all domain.

```
void addField( string const & fieldName,
              Location const location,
              localIndex const components,
              arrayViewId< string const > const & regions );
```

- **addCoupling**: creates a coupling between two fields (`rowField` and `colField`) according to a given connectivity in the regions defined by `regions`. Both fields (row and column) must have already been defined on the regions where is required the coupling among them. Default value for `regions` is the whole intersection between the regions where the first and the second fields are defined. This method also creates the coupling between `colField` and `rowField`, i.e. the transpose of the rectangular sparsity pattern. This default behaviour can be disabled by passing `symmetric = false`.

```
void addCoupling( string const & rowField,
                  string const & colField,
                  Connectivity const connectivity,
                  arrayViewId< string const > const & regions,
                  bool const symmetric );
```

- **reorderByRank**: finish populating field and coupling information and apply DoF re-numbering

```
void reorderByRank();
```

- **getKey**: returns the “key” associated with the field, that can be used to access the index array on the mesh object manager corresponding to field's location.

```
string const & getKey( string const & fieldName );
```

- **clear**: removes all fields, releases memory and re-opens the DofManager

```
void clear();
```

- **setSparsityPattern**: populates the sparsity for the given `rowField` and `colField` into matrix. Closes the matrix if `closePattern` is true.

```
void setSparsityPattern( MATRIX & matrix,
                        string const & rowField,
                        string const & colField,
                        bool closePattern = true) const;
```

- `setSparsityPattern`: populates the sparsity for the full system matrix into `matrix`. Closes the matrix if `closePattern` is `true`.

```
void setSparsityPattern( MATRIX & matrix,
                        bool closePattern = true ) const;
```

- `numGlobalDofs`: returns the total number of DoFs across all processors for the specified name `field` (if given) or all fields (if empty).

```
globalIndex numGlobalDofs( string const & field = "" ) const;
```

- `numLocalDofs`: returns the number of DoFs on this process for the specified name `field` (if given) or all fields (if empty).

```
localIndex numLocalDofs( string const & field = "" ) const;
```

- `printFieldInfo`: prints a short summary of declared fields and coupling to the output stream `os`.

```
void printFieldInfo( std::ostream & os = std::cout ) const;
```

Example

Here we show how the sparsity pattern is computed for a simple 2D quadrilateral mesh with 6 elements. Unknowns are pressure, located on the element center, and displacements (x and y components), located on the nodes. For fluxes, a two-point flux approximation (TPFA) is used. The representation of the sparsity pattern of the C_L matrix (connectors/locations) for the simple mesh, shown in Fig. 6.5, is reported in Fig. 6.6. It can be noticed that the two unknowns for the displacements x and y are grouped together. Elements are the connectivity for DoF on nodes (Finite Element Method for displacements) and on elements (pressures). Faces are the connectivity for DoF on elements (Finite Volume Method for pressure), being the flux computation based on the pressure on the two adjacent elements.

Fig. 6.5: Small 2D quadrilateral mesh used for this examples. Nodes are labeled with black numbers, elements with light gray numbers and faces with italic dark gray numbers.

Fig. 6.6: Sparsity pattern of the binary matrix connections/locations.

The global sparsity pattern, shown in Fig. 6.7, is obtained through the symbolic multiplication of the transpose of the matrix C_L and the matrix itself, i.e. $P = C_L^T C_L$.

Fig. 6.7: Sparsity pattern of the global matrix, where red and green entries are related to the displacement field and to the pressure field, respectively. Blue entries represent coupling blocks.

Real mesh and patterns

Now we build the pattern of the Jacobian matrix for a simple 3D mesh, shown in Fig. 6.8. Fields are:

- displacement (location: node, connectivity: element) defined on the blue, orange and red regions;
- pressure (location: element, connectivity: face) defined on the green, orange and red regions;
- mass matrix (location: element, connectivity: element) defined on the green region only.

Moreover, following coupling are imposed:

- displacement-pressure (connectivity: element) on the orange region only;
- pressure-mass matrix and transpose (connectivity: element) everywhere it is possible.

Fig. 6.8: Real mesh used to compute the Jacobian pattern.

Fig. 6.9 shows the global pattern with the field-based ordering of unknowns. Different colors mean different fields. Red unknowns are associated with displacement, yellow ones with pressure and blue ones with mass matrix. Orange means the coupling among displacement and pressure, while green is the symmetric coupling among pressure and mass matrix.

Fig. 6.9: Global pattern with field-based ordering. Red is associated with displacement unknowns, yellow with pressure ones and blue with those of mass matrix field. Orange means the coupling among displacement and pressure, while green is the symmetric coupling among pressure and mass matrix.

Fig. 6.10 shows the global pattern with the MPI rank-based ordering of unknowns. In this case, just two processes are used. Again, different colors indicate different ranks.

Fig. 6.10: Global pattern with MPI rank-based ordering. Red unknowns are owned by rank 0 and green ones by rank 1. Blue indicates the coupling among the two processes.

6.2.6 LvArray

Use in GEOSX

LvArray containers are used in GEOSX as primary storage mechanism for mesh topology, field data and any other type of “large” data sets (i.e. ones that scale with the size of the problem). When allocating a new field, using one of LvArray containers is mandatory if the data is meant to be used in any computational kernel. The file `common/DataTypes.hpp` provides shorthand aliases for commonly used containers:

```
/**
 * @name Aliases for LvArray::Array class family.
 */
///<@{

/// Multidimensional array type. See LvArray:Array for details.
template< typename T,
          int NDIM,
          typename PERMUTATION=camp::make_idx_seq_t< NDIM > >
using Array = LvArray::Array< T, NDIM, PERMUTATION, localIndex, LvArray::ChaiBuffer >;

/// Multidimensional array view type. See LvArray:ArrayView for details.
template< typename T,
          int NDIM,
          int USD = NDIM - 1 >
```

(continues on next page)

(continued from previous page)

```

using ArrayView = LvArray::ArrayView< T, NDIM, USD, localIndex, LvArray::ChaiBuffer >;

/// Multidimensional array slice type. See LvArray:ArraySlice for details.
template< typename T, int NDIM, int USD = NDIM - 1 >
using ArraySlice = LvArray::ArraySlice< T, NDIM, USD, localIndex >;

/// Multidimensional stack-based array type. See LvArray:StackArray for details.
template< typename T, int NDIM, int MAXSIZE, typename PERMUTATION=camp::make_idx_seq_t
↳< NDIM > >
using StackArray = LvArray::StackArray< T, NDIM, PERMUTATION, localIndex, MAXSIZE >;

///@}

/**
 * @name Short-hand aliases for commonly used array types.
 */
///@{

/// Alias for a local (stack-based) rank-1 tensor type
using R1Tensor = Tensor< real64, 3 >;

/// Alias for a local (stack-based) rank-2 Voigt tensor type
using R2SymTensor = Tensor< real64, 6 >;

/// Alias for 1D array.
template< typename T >
using array1d = Array< T, 1 >;

/// Alias for 1D array view.
template< typename T >
using arrayView1d = ArrayView< T, 1 >;

/// Alias for 1D array slice.
template< typename T, int USD = 0 >
using arraySlice1d = ArraySlice< T, 1, USD >;

/// Alias for 1D stack array.
template< typename T, int MAXSIZE >
using stackArray1d = StackArray< T, 1, MAXSIZE >;

/// Alias for 2D array.
template< typename T, typename PERMUTATION=camp::make_idx_seq_t< 2 > >
using array2d = Array< T, 2, PERMUTATION >;

/// Alias for 2D array view.
template< typename T, int USD = 1 >
using arrayView2d = ArrayView< T, 2, USD >;

/// Alias for 2D array slice.
template< typename T, int USD = 1 >
using arraySlice2d = ArraySlice< T, 2, USD >;

/// Alias for 2D stack array.
template< typename T, int MAXSIZE >
using stackArray2d = StackArray< T, 2, MAXSIZE >;

```

(continues on next page)

(continued from previous page)

```

/// Alias for 3D array.
template< typename T, typename PERMUTATION=camp::make_idx_seq_t< 3 > >
using array3d = Array< T, 3, PERMUTATION >;

/// Alias for 3D array view.
template< typename T, int USD=2 >
using arrayView3d = ArrayView< T, 3, USD >;

/// Alias for 3D array slice.
template< typename T, int USD=2 >
using arraySlice3d = ArraySlice< T, 3, USD >;

/// Alias for 3D stack array.
template< typename T, int MAXSIZE >
using stackArray3d = StackArray< T, 3, MAXSIZE >;

/// Alias for 4D array.
template< typename T, typename PERMUTATION=camp::make_idx_seq_t< 4 > >
using array4d = Array< T, 4, PERMUTATION >;

/// Alias for 4D array view.
template< typename T, int USD=3 >
using arrayView4d = ArrayView< T, 4, USD >;

/// Alias for 4D array slice.
template< typename T, int USD=3 >
using arraySlice4d = ArraySlice< T, 4, USD >;

/// Alias for 4D stack array.
template< typename T, int MAXSIZE >
using stackArray4d = StackArray< T, 4, MAXSIZE >;

/// Alias for 5D array.
template< typename T, typename PERMUTATION=camp::make_idx_seq_t< 5 > >
using array5d = Array< T, 5, PERMUTATION >;

/// Alias for 5D array view.
template< typename T, int USD=4 >
using arrayView5d = ArrayView< T, 5, USD >;

/// Alias for 5D array slice.
template< typename T, int USD=4 >
using arraySlice5d = ArraySlice< T, 5, 4 >;

/// Alias for 5D stack array.
template< typename T, int MAXSIZE >
using stackArray5d = StackArray< T, 5, MAXSIZE >;

///  

/**
 * @name Aliases for sorted arrays and set types.
 */
///  

/// A set of local indices.
template< typename T >

```

(continues on next page)

(continued from previous page)

```

using set = std::set< T >;

/// A sorted array of local indices.
template< typename T >
using SortedArray = LvArray::SortedArray< T, localIndex, LvArray::ChaiBuffer >;

/// A sorted array view of local indices.
template< typename T >
using SortedArrayView = LvArray::SortedArrayView< T, localIndex, LvArray::ChaiBuffer >
↪;

///  

/**
 * @name Aliases for LvArray::ArrayOfArrays class family.
 */
///  

/// Array of variable-sized arrays. See LvArray::ArrayOfArrays for details.
template< typename T >
using ArrayOfArrays = LvArray::ArrayOfArrays< T, localIndex, LvArray::ChaiBuffer >;

/// View of array of variable-sized arrays. See LvArray::ArrayOfArraysView for ↪
↪details.
template< typename T, bool CONST_SIZES=std::is_const< T >::value >
using ArrayOfArraysView = LvArray::ArrayOfArraysView< T, localIndex const, CONST_
↪SIZES, LvArray::ChaiBuffer >;

/// Array of variable-sized sets. See LvArray::ArrayOfSets for details.
template< typename T >
using ArrayOfSets = LvArray::ArrayOfSets< T, localIndex, LvArray::ChaiBuffer >;

/// View of array of variable-sized sets. See LvArray::ArrayOfSetsView for details.
template< typename T >
using ArrayOfSetsView = LvArray::ArrayOfSetsView< T, localIndex const, ↪
↪LvArray::ChaiBuffer >;

/// Alias for Sparsity pattern class.
template< typename COL_INDEX, typename INDEX_TYPE=localIndex >
using SparsityPattern = LvArray::SparsityPattern< COL_INDEX, INDEX_TYPE, ↪
↪LvArray::ChaiBuffer >;

/// Alias for Sparsity pattern View.
template< typename COL_INDEX, typename INDEX_TYPE=localIndex >
using SparsityPatternView = LvArray::SparsityPatternView< COL_INDEX, INDEX_TYPE const,
↪ LvArray::ChaiBuffer >;

/// Alias for CRS Matrix class.
template< typename T, typename COL_INDEX=globalIndex >
using CRSMatrix = LvArray::CRSMatrix< T, COL_INDEX, localIndex, LvArray::ChaiBuffer >;

/// Alias for CRS Matrix View.
template< typename T, typename COL_INDEX=globalIndex >
using CRSMatrixView = LvArray::CRSMatrixView< T, COL_INDEX, localIndex const, ↪
↪LvArray::ChaiBuffer >;

///  


```

LvArray documentation

Please refer to the full [LvArray](#) documentation for details on each of the classes.

6.2.7 Kernel interface

Finite Element Method Kernel Interface

The finite element method kernel interface (FEMKI) specifies an API for the launching of computational kernels for solving physics discretized using the finite element method. Using this approach, a set of generic element looping patterns and kernel launching functions may be implemented, and reused by various physics solvers that contain kernels conforming to the FEMKI.

There are several main components of the FEMKI:

1. A collection of element looping functions that provide various looping patterns, and call the `launch` function.
2. The kernel interface, which is specified by the `finiteElement::KernelBase` class. Each physics solver will define a class that contains its kernels functions, most likely deriving, or conforming to the API specified by the `KernelBase` class. Also part of this class will typically be a nested `StackVariables` class that defines a collection of stack variables for use in the various kernel interface functions.
3. A `launch` function, which launches the kernel, and calls the kernel interface functions conforming to the interface defined by `KernelBase`. This function is actually a member function of the `Kernel` class, so it may be overridden by a specific physics kernel, allowing complete customization of the interface, while maintaining the usage of the looping patterns.

A Generic Element Looping Pattern

One example of a looping pattern is the `regionBasedKernelApplication` function.

The contents of the looping function are displayed here:

```
/**
 * @brief Performs a loop over specific regions (by type and name) and calls a kernel_
 * launch on the subregions
 * with compile time knowledge of sub-loop bounds such as number of nodes and_
 * quadrature points per element.
 * @tparam POLICY The RAJA launch policy to pass to the kernel launch.
 * @tparam CONSTITUTIVE_BASE The common base class for constitutive pass-thru/
 * dispatch which gives the kernel
 * launch compile time knowledge of the constitutive model. This is achieved_
 * through a call to the
 * ConstitutivePassThru function which should have a specialization for_
 * CONSTITUTIVE_BASE implemented in
 * order to perform the compile time dispatch.
 * @tparam SUBREGION_TYPE The type of subregion to loop over. TODO make this a_
 * parameter pack?
 * @tparam KERNEL_FACTORY The type of @p kernelFactory, typically an instantiation of_
 * @c KernelFactory, and
 * must adhere to that interface.
 * @param mesh The MeshLevel object.
 * @param targetRegions The names of the target regions(of type @p SUBREGION_TYPE) to_
 * apply the @p KERNEL_TEMPLATE.
 * @param finiteElementName The name of the finite element.
 * @param constitutiveNames The names of the constitutive models present in the_
 * region.
```

(continues on next page)

(continued from previous page)

```

* @param kernelFactory The object used to construct the kernel.
* @return The maximum contribution to the residual, which may be used to scale the
↳residual.
*
* @details Loops over all regions Applies/Launches a kernel specified by the @p
↳KERNEL_TEMPLATE through
* #::geosx::finiteElement::KernelBase::kernelLaunch().
*/
template< typename POLICY,
          typename CONSTITUTIVE_BASE,
          typename SUBREGION_TYPE,
          typename KERNEL_FACTORY >
static
real64 regionBasedKernelApplication( MeshLevel & mesh,
                                     arrayViewld< string const > const &
↳targetRegions,
                                     string const & finiteElementName,
                                     arrayViewld< string const > const &
↳constitutiveNames,
                                     KERNEL_FACTORY & kernelFactory )
{
    GEOSX_MARK_FUNCTION;
    // save the maximum residual contribution for scaling residuals for convergence
↳criteria.
    real64 maxResidualContribution = 0;

    NodeManager & nodeManager = mesh.getNodeManager();
    EdgeManager & edgeManager = mesh.getEdgeManager();
    FaceManager & faceManager = mesh.getFaceManager();
    ElementRegionManager & elementRegionManager = mesh.getElemManager();

    // Loop over all sub-regions in regions of type SUBREGION_TYPE, that are listed in
↳the targetRegions array.
    elementRegionManager.forElementSubRegions< SUBREGION_TYPE >( targetRegions,
                                                                [&constitutiveNames,
                                                                 &
↳maxResidualContribution,
                                                                &nodeManager,
                                                                &edgeManager,
                                                                &faceManager,
                                                                &kernelFactory,
                                                                &finiteElementName]
                                                                ( localIndex const
↳targetRegionIndex, auto & elementSubRegion )
    {
        localIndex const numElems = elementSubRegion.size();

        // Get the constitutive model...and allocate a null constitutive model if
↳required.
        constitutive::ConstitutiveBase * constitutiveRelation = nullptr;
        constitutive::NullModel * nullConstitutiveModel = nullptr;
        if( targetRegionIndex <= constitutiveNames.size()-1 )
        {
            constitutiveRelation = &elementSubRegion.template getConstitutiveModel(
↳constitutiveNames[targetRegionIndex] );
        }
        else

```

(continues on next page)

(continued from previous page)

```

{
    nullConstitutiveModel = &elementSubRegion.template registerGroup<
↳constitutive::NullModel >( "nullModelGroup" );
    constitutiveRelation = nullConstitutiveModel;
}

    // Call the constitutive dispatch which converts the type of constitutive model
↳into a compile time constant.
    constitutive::ConstitutivePassThru< CONSTITUTIVE_BASE >::execute(
↳*constitutiveRelation,
                                                                    [&
↳maxResidualContribution,
                                                                    &nodeManager,
                                                                    &edgeManager,
                                                                    &faceManager,
                                                                    ]
↳targetRegionIndex,
                                                                    &kernelFactory,
                                                                    &
↳elementSubRegion,
                                                                    &
↳finiteElementName,
                                                                    numElems]
↳castedConstitutiveRelation )
    {
        FiniteElementBase &
        subRegionFE = elementSubRegion.template getReference< FiniteElementBase >(
↳finiteElementName );

        finiteElement::dispatch3D( subRegionFE,
                                    [&maxResidualContribution,
                                    &nodeManager,
                                    &edgeManager,
                                    &faceManager,
                                    targetRegionIndex,
                                    &kernelFactory,
                                    &elementSubRegion,
                                    numElems,
                                    &castedConstitutiveRelation] ( auto const
↳finiteElement )
        {
            auto kernel = kernelFactory.createKernel( nodeManager,
                                                        edgeManager,
                                                        faceManager,
                                                        targetRegionIndex,
                                                        elementSubRegion,
                                                        finiteElement,
                                                        castedConstitutiveRelation );

            using KERNEL_TYPE = decltype( kernel );

            // Call the kernelLaunch function, and store the maximum contribution to the
↳residual.
            maxResidualContribution =
                std::max( maxResidualContribution,
                          KERNEL_TYPE:template kernelLaunch< POLICY, KERNEL_TYPE >(
↳numElems, kernel ) );

```

(continues on next page)

(continued from previous page)

```

    } );
} );

// Remove the null constitutive model (not required, but cleaner)
if( nullConstitutiveModel )
{
    elementSubRegion.deregisterGroup( "nullModelGroup" );
}

} );

return maxResidualContribution;
}

```

This pattern may be used with any kernel class that either:

1. Conforms to the `KernelBase` interface by defining each of the kernel functions in `KernelBase`.
2. Defines its own `kernelLaunch` function that conforms the the signature of `KernelBase::kernelLaunch`. This option essentially allows for a custom kernel that does not conform to the interface defined by `KernelBase` and `KernelBase::kernelLaunch`.

The `KernelBase::kernelLaunch` Interface

The `kernelLaunch` function is a member of the kernel class itself. As mentioned above, a physics implementation may use the existing `KernelBase` interface, or define its own. The `KernelBase::kernelLaunch` function defines a launching policy, and an internal looping pattern over the quadrature points, and calls the functions defined by the `KernelBase` as shown here:

```

template< typename POLICY,
          typename KERNEL_TYPE >
static
real64
kernelLaunch( localIndex const numElems,
              KERNEL_TYPE const & kernelComponent )
{
    GEOSX_MARK_FUNCTION;

    // Define a RAJA reduction variable to get the maximum residual contribution.
    RAJA::ReduceMax< ReducePolicy< POLICY >, real64 > maxResidual( 0 );

    forAll< POLICY >( numElems,
                    [=] GEOSX_HOST_DEVICE ( localIndex const k )
                    {
                        typename KERNEL_TYPE::StackVariables stack;

                        kernelComponent.setup( k, stack );
                        for( integer q=0; q<numQuadraturePointsPerElem; ++q )
                        {
                            kernelComponent.quadraturePointKernel( k, q, stack );
                        }
                        maxResidual.max( kernelComponent.complete( k, stack ) );
                    } );
    return maxResidual.get();
}

```

Each of the `KernelBase` functions called in the `KernelBase::kernelLaunch` function are intended to provide a certain amount of modularity and flexibility for the physics implementations. The general purpose of each function is described by the function name, but may be further described by the function documentation found [here](#).

6.2.8 Adding a new Physics Solver

In this tutorial, you will learn how to construct a new GEOSX Physics Solver class. We will use *LaplaceFEM* solver, computing the solution of the Laplace problem in a specified material, as a starting point.

$$\begin{aligned} D^* \Delta X &= f \quad \text{in } \Omega \\ X &= X^g \quad \text{on } \Gamma_g \end{aligned}$$

It is advised to read *XML Input* preliminary to this tutorial. The goal of this document is to explain how to develop a new solver that solves Laplace's equation with a constant diffusion coefficient that is specified by users in the XML input.

For readability, member functions in the text will be referenced by their names but their arguments will be omitted.

LaplaceFEM overview

The *LaplaceFEM* solver can be found in `./src/coreComponents/physicsSolvers/simplePDE/`. Let us inspect declarations in `LaplaceFEM.hpp` and implementations in `LaplaceFEM.cpp` before diving into specifying a new solver class that meets our needs.

Declaration file (reference)

The included header is `physicsSolvers/simplePDE/LaplaceBaseH1.hpp` which declares the base class `LaplaceBaseH1`, shared by all Laplace solvers. Moreover, `physicsSolver/simplePDE/LaplaceBaseH1.hpp` includes the following headers:

- `common/EnumStrings.hpp` which includes facilities for enum-string conversion (useful for reading enum values from input);
- `physicsSolver/SolverBase.hpp` which declares the abstraction class shared by all physics solvers.
- `managers/FieldSpecification/FieldSpecificationManager.hpp` which declares a manager used to access and to set field on the discretized domain.

Let us jump forward to the class enum and variable as they contain the data used specifically in the implementation of *LaplaceFEM*.

class enums and variables (reference)

The class exhibits two member variables:

- `m_fieldName` which stores the name of the diffused variable (*e.g.* the temperature) as a *string*;
- `m_timeIntegrationOption` an *enum* value allowing to dispatch with respect to the transient treatment.

`TimeIntegrationOption` is an *enum* specifying the transient treatment which can be chosen respectively between *SteadyState* and *ImplicitTransient* depending on whether we are interested in the transient state.

```
enum class TimeIntegrationOption : integer
{
    SteadyState,
    ImplicitTransient
};
```

In order to register an enumeration type with the Data Repository and have its value read from input, we must define stream insertion/extraction operators. This is a common task, so GEOSX provides a facility for automating it. Upon including `common/EnumStrings.hpp`, we can call the following macro at the namespace scope (in this case, right after the `LaplaceBaseH1` class definition is complete):

```
ENUM_STRINGS( LaplaceBaseH1::TimeIntegrationOption,
              "SteadyState",
              "ImplicitTransient" );
```

Once explained the main variables and enum, let us start reading through the different member functions:

```
class LaplaceFEM : public LaplaceBaseH1
{
public:
    /// The default nullary constructor is disabled to avoid compiler auto-generation:
    LaplaceFEM() = delete;

    /// The constructor needs a user-defined "name" and a parent Group (to place this_
    ↪instance in the
    /// tree structure of classes)
    LaplaceFEM( const string & name,
                Group * const parent );

    /// Destructor
    virtual ~LaplaceFEM() override;

    /// "CatalogName()" return the string used as XML tag in the input file. It ties_
    ↪the XML tag with
    /// this C++ classes. This is important.
    static string catalogName() { return "LaplaceFEM"; }
```

Start looking at the class *LaplaceFEM* constructor and destructor declarations shows the usual *string* name and *Group** pointer to parent that are required to build the global file-system like structure of GEOSX (see *Group : the base class of GEOSX* for details). It can also be noted that the nullary constructor is deleted on purpose to avoid compiler automatic generation and user misuse.

The next method `catalogName()` is static and returns the key to be added to the *Catalog* for this type of solver (see *A few words about the ObjectCatalog* for details). It has to be paired with the following macro in the implementation file.

```
REGISTER_CATALOG_ENTRY( SolverBase, LaplaceFEM, string const &, Group * const )
```

Finally, the member function `registerDataOnMesh()` is declared in the `LaplaceBaseH1` class as

```
/// This method ties properties with their supporting mesh
virtual void registerDataOnMesh( Group & meshBodies ) override final;
```

It is used to assign fields onto the discretized mesh object and will be further discussed in the *Implementation File (reference)* section.

The next block consists in solver interface functions. These member functions set up and specialize every time step from the system matrix assembly to the solver stage.

```

virtual void
setupSystem( DomainPartition & domain,
             DofManager & dofManager,
             CRSMatrix< real64, globalIndex > & localMatrix,
             array1d< real64 > & localRhs,
             array1d< real64 > & localSolution,
             bool const setSparsity = false ) override;

virtual void
assembleSystem( real64 const time,
               real64 const dt,
               DomainPartition & domain,
               DofManager const & dofManager,
               CRSMatrixView< real64, globalIndex const > const & localMatrix,
               arrayView1d< real64 > const & localRhs ) override;

```

Furthermore, the following functions are inherited from the base class.

```

virtual real64 solverStep( real64 const & time_n,
                        real64 const & dt,
                        integer const cycleNumber,
                        DomainPartition & domain ) override;

virtual void
implicitStepSetup( real64 const & time_n,
                  real64 const & dt,
                  DomainPartition & domain ) override;

virtual void
setupDofs( DomainPartition const & domain,
           DofManager & dofManager ) const override;

virtual void
applyBoundaryConditions( real64 const time,
                       real64 const dt,
                       DomainPartition & domain,
                       DofManager const & dofManager,
                       CRSMatrixView< real64, globalIndex const > const &
↳localMatrix,
                       arrayView1d< real64 > const & localRhs ) override;

virtual void
solveSystem( DofManager const & dofManager,
             ParallelMatrix & matrix,
             ParallelVector & rhs,
             ParallelVector & solution ) override;

virtual void
applySystemSolution( DofManager const & dofManager,
                    arrayView1d< real64 const > const & localSolution,
                    real64 const scalingFactor,
                    DomainPartition & domain ) override;

virtual void updateState( DomainPartition & domain ) override final;

virtual void
resetStateToBeginningOfStep( DomainPartition & GEOSX_UNUSED_PARAM( domain ) )
↳override;

```

(continues on next page)

(continued from previous page)

```

virtual void
implicitStepComplete( real64 const & time,
                      real64 const & dt,
                      DomainPartition & domain ) override;

/// This method is specific to this Laplace solver.
/// It is used to apply Dirichlet boundary condition
/// and called when the base class applyBoundaryConditions() is called.
virtual void applyDirichletBCImplicit( real64 const time,
                                      DofManager const & dofManager,
                                      DomainPartition & domain,
                                      CRSMatrixView< real64, globalIndex const >
↳ const & localMatrix,
                                      arrayView1d< real64 > const & localRhs );

```

Eventually, `applyDirichletBCImplicit()` is the working specialized member functions called when `applyBoundaryConditions()` is called in this particular class override.

Browsing the base class `SolverBase`, it can be noted that most of the solver interface functions are called during either `SolverBase::linearImplicitStep()` or `SolverBase::nonlinearImplicitStep()` depending on the solver strategy chosen.

Switching to protected members, `postProcessInput()` is a central member function and will be called by `Group` object after input is read from XML entry file. It will set and dispatch solver variables from the base class `SolverBase` to the most derived class. For `LaplaceFEM`, it will allow us to set the right time integration scheme based on the XML value as will be further explored in the next *Implementation File (reference)* section.

Let us focus on a struct that plays an important role: the `viewKeyStruct` structure.

viewKeyStruct structure (reference)

This embedded instantiated structure is a common pattern shared by all solvers. It stores `dataRepository::ViewKey` type objects that are used as binding data between the input XML file and the source code.

```

struct viewKeyStruct : public SolverBase::viewKeyStruct
{
    static constexpr char const * timeIntegrationOption() { return
↳ "timeIntegrationOption"; }
    static constexpr char const * fieldVarName() { return "fieldName"; }
};

```

We can check that in the *LaplaceFEM* companion `integratedTest`

```

<LaplaceFEM
  name="laplace"
  discretization="FE1"
  timeIntegrationOption="SteadyState"
  fieldName="Temperature"
  targetRegions="{ Region1 }">
  <LinearSolverParameters
    directParallel="0"/>
</LaplaceFEM>

```

In the following section, we will see where this binding takes place.

Implementation File (reference)

Switching to implementation, we will focus on few implementations, leaving details to other tutorials. The LaplaceFEM constructor is implemented as follows.

```
LaplaceFEM::LaplaceFEM( const string & name,
                        Group * const parent ):
    LaplaceBaseH1( name, parent )
{ }
```

As we see, it calls the LaplaceBaseH1 constructor, that is implemented as follows.

```
LaplaceBaseH1::LaplaceBaseH1( const string & name,
                               Group * const parent ):
    SolverBase( name, parent ),
    m_fieldName( "primaryField" ),
    m_timeIntegrationOption( TimeIntegrationOption::ImplicitTransient )
{
    this->registerWrapper( viewKeyStruct::timeIntegrationOption(), &m_
    ↪timeIntegrationOption ).
        setInputFlag( InputFlags::REQUIRED ).
        setDescription( "Time integration method. Options are:\n* " + EnumStrings<
    ↪TimeIntegrationOption >::concat( "\n* " ) );

    this->registerWrapper( viewKeyStruct::fieldVarName(), &m_fieldName ).
        setInputFlag( InputFlags::REQUIRED ).
        setDescription( "Name of field variable" );
}
```

Checking out the constructor, we can see that the use of a registerWrapper<T>(...) allows us to register the key value from the *enum* viewKeyStruct defining them as:

- InputFlags::OPTIONAL if they are optional and can be provided;
- InputFlags::REQUIRED if they are required and will throw error if not;

and their associated descriptions for auto-generated docs.

```
void LaplaceBaseH1::registerDataOnMesh( Group & meshBodies )
{
    meshBodies.forSubGroups< MeshBody >( [&] ( MeshBody & meshBody )
    {
        NodeManager & nodes = meshBody.getMeshLevel( 0 ).getNodeManager();

        nodes.registerWrapper< real64_array >( m_fieldName ).
            setApplyDefaultValue( 0.0 ).
            setPlotLevel( PlotLevel::LEVEL_0 ).
            setDescription( "Primary field variable" );
    } );
}
```

registerDataOnMesh() is browsing all subgroups in the mesh Group object and for all nodes in the sub group:

- register the observed field under the chosen m_fieldName key;
- apply a default value;
- set the output verbosity level (here PlotLevel::LEVEL_0);
- set the field associated description for auto generated docs.


```

void LaplaceFEM::assembleSystem( real64 const GEOSX_UNUSED_PARAM( time_n ),
                                real64 const GEOSX_UNUSED_PARAM( dt ),
                                DomainPartition & domain,
                                DofManager const & dofManager,
                                CRSMatrixView< real64, globalIndex const > const &
    ↪ localMatrix,
                                arrayViewld< real64 > const & localRhs )
{
    MeshLevel & mesh = domain.getMeshBody( 0 ).getMeshLevel( 0 );

    NodeManager & nodeManager = mesh.getNodeManager();
    string const dofKey = dofManager.getKey( m_fieldName );
    arrayViewld< globalIndex const > const &
    dofIndex = nodeManager.getReference< arrayld< globalIndex > >( dofKey );

    LaplaceFEMKernelFactory kernelFactory( dofIndex, dofManager.rankOffset(),
    ↪ localMatrix, localRhs, m_fieldName );

    finiteElement::
        regionBasedKernelApplication< parallelDevicePolicy< 32 >,
                                    constitutive::NullModel,
                                    CellElementSubRegion >( mesh,
                                                            targetRegionNames(),
                                                            this->
    ↪ getDiscretizationName(),
                                                            arrayViewld< string const >
    ↪ (),
                                                            kernelFactory );
}

```

`assembleSystem()` will be our core focus as we want to change the diffusion coefficient from its hard coded value to a XML read user-defined value. One can see that this method is in charge of constructing in a parallel fashion the FEM system matrix. Bringing `nodeManager` and `ElementRegionManager` from domain local `MeshLevel` object together with `FiniteElementDiscretizationManager` from the `NumericalMethodManager`, it uses nodes embedded loops on degrees of freedom in a local index embedded loops to fill a matrix and a rhs container.

As we spotted the place to change in a code to get a user-defined diffusion coefficient into the game, let us jump to writing our new *LaplaceDiffFEM* solver.

Note: We might want to remove final keyword from `postProcessInput()` as it will prevent you from overriding it.

Start doing your own Physic solver

As we will extend *LaplaceFEM* capabilities, we will derive publicly from it.

Declaration File

As there is only few places where we have to change, the whole declaration file is reported below and commented afterwards.

```
#include "physicsSolvers/simplePDE/LaplaceFEM.hpp"

namespace geosx
{
class LaplaceDiffFEM : public LaplaceFEM
{
public:

    LaplaceDiffFEM() = delete;

    LaplaceDiffFEM( const string& name,
                    Group * const parent );

    virtual ~LaplaceDiffFEM() override;

    static string catalogName() { return "LaplaceDiffFEM"; }

    virtual void
    assembleSystem( real64 const time,
                    real64 const dt,
                    DomainPartition * const domain,
                    DofManager const & dofManager,
                    ParallelMatrix & matrix,
                    ParallelVector & rhs ) override;

    struct viewKeyStruct : public LaplaceFEM::viewKeyStruct
    {
        dataRepository::ViewKey diffusionCoeff = { "diffusionCoeff" };
    } laplaceDiffFEMViewKeys;

    protected:
    virtual void postProcessInput() override final;

private:
    real64 m_diffusion;
};
```

We intend to have a user-defined diffusion coefficient, we then need a *real64* class variable `m_diffusion` to store it.

Consistently with *LaplaceFEM*, we will also delete the nullary constructor and declare a constructor with the same arguments for forwarding to *Group* master class. Another mandatory step is to override the static `CatalogName()` method to properly register any data from the new solver class.

Then as mentioned in *Implementation File (reference)*, the diffusion coefficient is used when assembling the matrix coefficient. Hence we will have to override the `assembleSystem()` function as detailed below.

Moreover, if we want to introduce a new binding between the input XML and the code we will have to work on the three struct `viewKeyStruct`, `postProcessInput()` and the constructor.

Our new solver `viewKeyStruct` will have its own structure inheriting from the *LaplaceFEM* one to have the `timeIntegrationOption` and `fieldName` field. It will also create a `diffusionCoeff` field to be bound to the user defined homogeneous coefficient on one hand and to our `m_diffusion` class variable on the other.

Implementation File

As we have seen in *Implementation File (reference)*, the first place where to implement a new register from XML input is in the constructor. The `diffusionCoeff` entry we have defined in the `laplaceDiffFEMViewKeys` will then be asked as a required input. If not provided, the error thrown will ask for it described as an “input uniform diffusion coefficient for the Laplace equation”.

```
LaplaceDiffFEM::LaplaceDiffFEM( const string& name,
                                Group * const parent ):
LaplaceFEM( name, parent ), m_diffusion(0.0)
{
    registerWrapper<string>(laplaceDiffFEMViewKeys.diffusionCoeff.Key()).
        setInputFlag(InputFlags::REQUIRED).
        setDescription("input uniform diffusion coeff for the laplace equation");
}
```

Another important spot for binding the value of the XML read parameter to our `m_diffusion` is in `postProcessInput()`.

```
void LaplaceDiffFEM::postProcessInput()
{
    LaplaceFEM::postProcessInput();

    string sDiffCoeff = this->getReference<string>(laplaceDiffFEMViewKeys.
    ↪diffusionCoeff);
    this->m_diffusion = std::stof(sDiffCoeff);
}
```

Now that we have required, read and bind the user-defined diffusion value to a variable, we can use it in the construction of our matrix into the overridden `assembleSystem()`.

```
// begin element loop, skipping ghost elements
for( localIndex k=0 ; k<elementSubRegion->size() ; ++k )
{
    if(elemGhostRank[k] < 0)
    {
        element_rhs = 0.0;
        element_matrix = 0.0;
        for( localIndex q=0 ; q<n_q_points ; ++q)
        {
            for( localIndex a=0 ; a<numNodesPerElement ; ++a)
            {
                elemDofIndex[a] = dofIndex[ elemNodes( k, a ) ];

                for( localIndex b=0 ; b<numNodesPerElement ; ++b)
                {
                    element_matrix(a,b) += detJ[k][q] *
                                            m_diffusion *
                                            + Dot( dNdX[k][q][a], dNdX[k][q][b] );
                }
            }
        }
        matrix.add( elemDofIndex, elemDofIndex, element_matrix );
        rhs.add( elemDofIndex, element_rhs );
    }
}
```

This completes the implementation of our new solver *LaplaceDiffFEM*.

Nonetheless, the compiler should complain that `m_fieldName` is privately as inherited from *LaplaceFEM*. One should then either promote `m_fieldName` to protected or add a getter in *LaplaceFEM* class to correct the error. The getter option has been chosen and the fix in our solver is then:

```
arrayld<globalIndex> const & dofIndex =  
    nodeManager->getReference< arrayld<globalIndex> >( dofManager.getKey(     
    ↪getFieldName() ) );
```

Note: For consistency do not forget to change *LaplaceFEM* to *LaplaceDiffFEM* in the guards comments

Last steps

After assembling both declarations and implementations for our new solver, the final steps go as:

- add declarations to parent CMakeLists.txt (here add to `physicsSolvers_headers`);
- add implementations to parent CMakeLists.txt (here add to `physicsSolvers_sources`);
- check that Doxygen comments are properly set in our solver class;
- uncrustify it to match the code style;
- write unit tests for each new features in the solver class;
- write an `integratedTests` for the solver class.

CHAPTER 7

Doxygen

The c++ source in GEOSX is annotated using doxygen. Our doxygen pages are linked below.

Developers may find it helpful to review the key code components described in the Developer Guide before diving into the doxygen.

Welcome to the GEOSX build guide.

8.1 System prerequisites

To configure and build GEOSX you will need the following tools available on your system.

8.1.1 List of prerequisites

Minimal requirements:

- [CMake](#) build system generator (3.13+).
- build tools ([GNU make](#) or [ninja](#) on Linux, XCode on MacOS).
- a C++ compiler with full c++14 standard support ([gcc](#) 8.3+ or [clang](#) 8.0+ are recommended).
- [python](#) (2.7+ or 3.6+).
- [zlib](#), [blas](#) and [lapack](#) libraries
- any compatible MPI runtime and compilers (if building with MPI)

If you want to build from a repository check out (instead of a release tarball):

- [git](#) (2.20+ is tested, but most versions should work fine)

If you plan on building bundled third-party library (TPLs) dependencies yourself:

- Compatible C and Fortran compilers

If you will be checking out and running integrated tests (a submodule of GEOSX, currently not publicly available):

- [git-lfs](#) (Git Large File Storage extension)
- [h5py](#) and [mpi4py](#) python modules

If you are interested in building Doxygen documentation:

- GNU bison
- LaTeX
- ghostscript
- Graphviz

In order for XML validation to work (executed as an optional build step):

- xmlint

8.1.2 Installing prerequisites

On a local development machine with sudo/root privileges, most of these dependencies can be installed with a system package manager. For example, on a Debian-based system (check your package manager for specific package names):

```
sudo apt install build-essential git git-lfs gcc g++ gfortran cmake libopenmpi-dev_  
↳ libblas-dev liblapack-dev zlib1g-dev python3 python3-h5py python3-mpi4py libxml2-  
↳ utils
```

On HPC systems it is typical for these tools to be installed by system administrators and provided via [modules](#). To list available modules, type:

```
module avail
```

Then load the appropriate modules using `module load` command. Please contact your system administrator if you need help choosing or installing appropriate modules.

8.2 Third-party dependencies

GEOSX makes use of multiple third-party libraries (TPLs) and tools, some of which are mandatory and some optional. We only test against specific versions, and sometimes even require development snapshots (specific git commits). Not all of these guarantee backwards compatibility, so we strongly recommend building with these specific versions.

8.2.1 List of third-party libraries and tools

The two tables below lists the dependencies with their specific versions and relevant CMake variables. Some of these libraries may have their own system prerequisites.

Libraries

The following libraries are linked to by GEOSX:

Name	Version	Enable option	Path variable	Description
Adiak	0.2.0	ENABLE_CALIPER	ADIAK_DIR	Library for collecting metadata from HPC application runs, and distributing that metadata to subscriber tools.
Caliper	2.4.0	ENABLE_CALIPER	CALIPER_DIR	Instrumentation and performance profiling library.
conduit	0.5.0	<i>mandatory</i>	CONDUIT_DIR	Simplified Data Exchange for HPC Simulations.
CHAI	2.2.2	<i>mandatory</i>	CHAI_DIR	Copy-hiding array abstraction to automatically migrate data between memory spaces.
RAJA	0.12.1	<i>mandatory</i>	RAJA_DIR	Collection of C++ software abstractions that enable architecture portability for HPC applications.
hdf5	1.10.5	<i>mandatory</i>	HDF5_DIR	High-performance data management and storage suite.
math-presso	2015-12-15	ENABLE_MATHPRESSO	MATHPRESSO_DIR	Mathematical Expression Parser and JIT Compiler.
pugixml	1.8.0	<i>mandatory</i>	PUGIXML_DIR	Light-weight, simple and fast XML parser for C++ with XPath support.
parmetis	4.0.3	<i>mandatory</i> (with MPI)	PARMETIS_DIR	Parallel Graph Partitioning library. Should be built with 64-bit <code>idx_t</code> type.
suites-parse	5.8.1	ENABLE_SUITESPARSE	SUITESPARSE_DIR	A suite of sparse matrix software.
superlu-dist	0f6efc3	ENABLE_SUPERLU	SUPERLU_DIST_DIR	General purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations.
hypre	2186a8f	ENABLE_HYPRE	HYPRE_DIR	Library of high performance preconditioners and solvers for large sparse linear systems on massively parallel computers.
PETSc	3.13.0	ENABLE_PETSC	PETSC_DIR	Suite of data structures and routines for the scalable (parallel) solution of scientific applications.
Trilinos	12.18.1	ENABLE_TRILINOS	TRILINOS_DIR	Collection of reusable scientific software libraries, known in particular for linear solvers.
siilo	4.10.3	<i>mandatory</i>	SILO_DIR	A Mesh and Field I/O Library and Scientific Database.
VTK	9.0.0-rc3	ENABLE_VTK	VTK_DIR	Open source software for manipulating and displaying scientific data.

Tools

The following tools are used as part of the build process to support GEOSX development:

Name	Version	Enable option	Path variable	Description
doxygen	1.8.20	ENABLE_DOXYGEN	DOXYGEN_EXECUTABLE	De facto standard tool for generating documentation from annotated C++ sources.
sphinx	1.8.5	ENABLE_SPHINX	SPHINX_EXECUTABLE	A tool that makes it easy to create intelligent and beautiful documentation.
uncrustify	401a409	ENABLE_UNCRUSTIFY	UNCrustify_EXECUTABLE	A source code beautifier for C, C++, C#, ObjectiveC, D, Java, Pawn and VALA.

Some other dependencies ([GoogleTest](#), [GoogleBenchmark](#)) are provided through [BLT](#) build system which is embedded in GEOSX source. No actions are needed to build them.

If you would like to create a Docker image with all dependencies, take a look at [Dockerfiles](#) that are used in our CI process.

8.2.2 Building bundled dependencies

To simplify the process of building TPLs, we provide a git repository [thirdPartyLibs](#). It contains source copies of exact TPL versions required and is updated periodically. It also contains a CMake script for building all TPLs in a single command.

The recommended steps to build TPLs are:

- Create a host-config file that sets all system-specific CMake variables (compiler and library paths, configuration flags, etc.) Take a look at [host-config examples](#).
- Configure via `config-build.py` script:

```
cd thirdPartyLibs
python scripts/config-build.py --hostconfig=/path/to/host-config.cmake --
    ↪ buildtype=Release --installpath=/path/to/install/dir -DNUM_PROC=8
```

where

- `--buildpath` or `-bp` is the build directory (by default, created under current).
- `--installpath` or `-ip` is the installation directory (wraps `CMAKE_INSTALL_PREFIX`).
- `--buildtype` or `-bt` is a wrapper to the `CMAKE_BUILD_TYPE` option.
- `--hostconfig` or `-hc` is a path to host-config file.
- all other command-line options are passed to CMake.

- Run the build:

```
cd <buildpath>
make
```

Warning: Do not provide `-j` argument to `make` here, since the top-level `make` only launches sub-project builds. Instead use `-DNUM_PROC` option above, which is passed to each sub-project's `make` command.

You may also run the CMake configure step manually instead of relying on `config-build.py`. The full TPL build may take anywhere between 15 minutes and 2 hours, depending on your machine, number of threads and libraries enabled.

Note: An exception from the above pattern, `sphinx` is currently not a part of the TPL bundle and must be installed with your Python or package manager.

Note: PETSc build currently downloads `pt-scotch` from the internet. If you do not have access to internet, modify the `./configure` step of `petsc` in `CMakeLists.txt` and change the `--download-ptscotch` option accordingly. `pt-scotch` also relies on `bison` and `flex`.

8.2.3 Installing dependencies individually

You may also install each individual TPL separately, either manually or through a package manager. This is a more difficult route, since you are responsible for configuring dependencies in a compatible manner. Again, we strongly recommend using the exact versions listed above, to avoid possible build problems.

You may look at [our TPL CMake script](#) to see how we configure TPL builds.

8.3 Building GEOSX

8.3.1 Build steps

- Create a host-config file that sets all system-specific CMake variables. Take a look at [host-config examples](#). We recommend the same host-config is used for both TPL and GEOSX builds. In particular, certain options (such as `ENABLE_MPI` or `ENABLE_CUDA`) need to match between the two.
- Provide paths to all enabled TPLs. This can be done in one of two ways:
 - Provide each path via a separate CMake variable (see [Third-party dependencies](#) for path variable names).
 - If you built TPLs from the `tplMirror` repository, you can set `GEOSX_TPL_DIR` variable in your host-config to point to the TPL installation path, and

```
include("/path/to/GEOSX/host-configs/tpls.cmake")
```

which will set all the individual TPL paths for you.

- Configure via `config-build.py` script:

```
cd GEOSX
python scripts/config-build.py --hostconfig=/path/to/host-config.cmake --
➔ buildtype=Release --installpath=/path/to/install/dir
```

where

- `--buildpath` or `-bp` is the build directory (by default, created under current working dir).
- `--installpath` or `-ip` is the installation directory (wraps `CMAKE_INSTALL_PREFIX`).
- `--buildtype` or `-bt` is a wrapper to the `CMAKE_BUILD_TYPE` option.
- `--hostconfig` or `-hc` is a path to host-config file.
- all unrecognized options are passed to CMake.

- Run the build:

```
cd <buildpath>
make -j
```

You may also run the CMake configure step manually instead of relying on `config-build.py`. A full build typically takes between 10 and 30 minutes, depending on chosen compilers, options and number of cores.

8.3.2 Configuration options

Below is a list of CMake configuration options, in addition to TPL options above. Some options, when enabled, require additional settings (e.g. `ENABLE_CUDA`). Please see [host-config examples](#).

Option	De- fault	Explanation
ENABLE_MPI	ON	Build with MPI (also applies to TPLs)
ENABLE_OPENMP	OFF	Build with OpenMP (also applies to TPLs)
ENABLE_CUDA	OFF	Build with CUDA (also applies to TPLs)
ENABLE_DOCS	ON	Build documentation (Sphinx and Doxygen)
ENABLE_WARNINGS_AS_ERRORS	ON	Treat all warnings as errors
ENABLE_PAMELA	ON	Enable PAMELA library (required for external mesh import)
ENABLE_PVTPackage	ON	Enable PVTPackage library (required for compositional flow runs)
ENABLE_TOTALVIEW_OUTPUT	OFF	Enables TotalView debugger custom view of GEOSX data structures
GEOSX_ENABLE_FPE	ON	Enable floating point exception trapping
GEOSX_LA_INTERFACE	Hypre	Choice of Linear Algebra backend (Hypre/Petsc/Trilinos)

8.4 Spack and Uberenv

GEOSX is transitioning to a new [Spack](#) and [Uberenv](#) system for building our dependencies. We refer the reader to the [Spack documentation](#) and [Uberenv documentation](#), in particular the Spack documentation for [specs and dependencies](#), [manual compiler configuration](#) and [external packages](#) are worth reading.

Building the dependencies can be as simple as running

```
./scripts/uberenv/uberenv.py
```

This will create a directory `uberenv_libs` in the current working directory, clone Spack into `uberenv_libs/spack` and install the dependencies into `uberenv_libs/system_dependent_path`. It will then spit out a host-config file in the current directory which you can use to build GEOSX. While the above command **should** work on every system, it should never be used. Invoked as such, Spack will ignore any system libraries you have installed and will go down a rabbit hole building dependencies. Furthermore this does not allow you to choose the compiler to build. Both of these are easily solved by creating a directory with a `packages.yaml` and a `compilers.yaml`.

To prevent this from happening you'll need to create a directory with a `packages.yaml` file and a `compilers.yaml` file. You can find working examples for commonly used systems in [scripts/uberenv/spack_configs](#). It is worth noting that each LC system type has two such directories, for example there is a `toss_3_x85_54_ib` and `toss_3_x85_54_ib_python` directory. This is because when building `pygeosx` Python needs to be built from scratch, and as such cannot be listed in `packages.yaml`. However, when not building `pygeosx` other dependencies depend on python, but an existing system version works just fine, so it can be put in `packages.yaml` to prevent Spack from building it.

Once you have these files setup you can run Uberenv again and instruct it to use them with. If for instance you added Clang 10.0.1 to the `compilers.yaml` file the your command would look something like this:

```
./scripts/uberenv/uberenv.py --spack-config-dir=/path/to/your/config/directory/ --  
↪ spec="%clang@10.0.1"
```

Note: When building `pygeosx`, Spack will build various python packages, however by default they are not installed in python. There are various ways of accomplishing [this](#), but the recommended approach is to use `spack activate`. The command would look something like this `./uberenv_libs/spack/bin/spack activate py-numpy py-scipy py-pip py-mpi4py`

8.4.1 Build Configuration

The GEOSX Spack package has a lot of options for controlling which dependencies you would like to build and how you'd like them built. The GEOSX Spack package file is at ``scripts/uberenv/packages/geosx/package.py`` <<https://github.com/GEOSX/GEOSX/tree/develop/scripts/uberenv/packages/geosx/package.py>>`. The variants for the package are as follows

```
variant('shared', default=True, description='Build Shared Libs.')
variant('caliper', default=True, description='Build Caliper support.')
variant('mkl', default=False, description='Use the Intel MKL library.')
variant('essl', default=False, description='Use the IBM ESSL library.')
variant('suite-sparse', default=True, description='Build SuiteSparse support.')
variant('trilinos', default=True, description='Build Trilinos support.')
variant('hypre', default=True, description='Build HYPRE support.')
variant('hypre-cuda', default=False, description='Build HYPRE with CUDA support.')
variant('petsc', default=True, description='Build PETSc support.')
variant('lai', default='trilinos', description='Linear algebra interface.',
        values=('trilinos', 'hypre', 'petsc'), multi=False)
variant('pygeosx', default=False, description='Build the GEOSX python interface.')
```

For example if you wanted to build with GCC 8.3.1, without Caliper and with PETSC as the Linear Algebra Interface, your spec would be `%gcc@8.3.1 ~caliper lai=petsc`.

The GEOSX Spack package lists out the libraries that GEOSX depends ons. Currently these dependencies are

```
depends_on('cmake@3.8:', type='build')
depends_on('cmake@3.9:', when='+cuda', type='build')

#
# Virtual packages
#
depends_on('mpi')
depends_on('blas')
depends_on('lapack')

#
# Performance portability
#
depends_on('raja@0.12.1 +openmp +shared ~examples ~exercises')
depends_on('raja +cuda', when='+cuda')

depends_on('umpire@4.1.2 ~c +shared +openmp ~examples')
depends_on('umpire +cuda', when='+cuda')

depends_on('chai@2.2.2 +shared +raja ~benchmarks ~examples')
depends_on('chai@2.2.2 +cuda', when='+cuda')

#
# IO
#
depends_on('hdf5@1.10.5: +shared +pic +mpi', when='~vtk')

depends_on('conduit@0.5.0 +shared ~test ~fortran +mpi +hdf5 ~hdf5_compat')

depends_on('silo@4.10: ~fortran +shared ~silex +pic +mpi ~zlib')
```

(continues on next page)

(continued from previous page)

```

depends_on('adiak@0.2: +mpi +shared', when='+caliper')
depends_on('caliper@2.4: +shared +adiak +mpi ~callpath ~libpfm ~gotcha ~sampler',
↳when='+caliper')

depends_on('pugixml@1.8: +shared')

#
# Math
#
depends_on('intel-mkl +shared ~ilp64', when='+mkl')

# depends_on('essl ~ilp64 threads=openmp +lapack +cuda', when='+essl')

depends_on('parmetis@4.0.3: +shared +int64')

depends_on('superlu-dist +int64 +openmp +shared', when='~petsc')
depends_on('superlu-dist@6.3.0 +int64 +openmp +shared', when='+petsc')

depends_on('suite-sparse@5.8.1: +pic +openmp +amd +camd +colamd +ccolamd +cholmod
↳+umfpack', when='+suite-sparse')
depends_on('suite-sparse +blas-no-underscore', when='%gcc +suite-sparse +essl')

trilinos_build_options = '~fortran +openmp +shared'
trilinos_tpls = '~boost ~glm ~gtest ~hdf5 ~hybre ~matio ~metis +mpi ~mumps ~
↳netcdf ~suite-sparse'
trilinos_packages = '+amesos +aztec +epetra +epetraext +ifpack +kokkos +ml +stk
↳+stratimikos +teuchos +tpetra ~amesos2 ~anasazi ~belos ~exodus ~ifpack2 ~muelu ~
↳sacado ~zoltan ~zoltan2'
depends_on('trilinos@12.18.1 ' + trilinos_build_options + trilinos_tpls +
↳trilinos_packages, when='+trilinos')
depends_on('trilinos +blas_lowercase_no_underscore', when='+trilinos +essl')
# depends_on('trilinos +force-new-lapack', when='+trilinos +essl')

depends_on('hybre@2.20.1 +shared +superlu-dist +mixedint +mpi +openmp', when=
↳'+hybre')
depends_on('hybre@2.20.1 +cuda +shared +superlu-dist +mpi +openmp +unified-memory
↳+cusparse', when='+hybre-cuda')

petsc_build_options = '+shared +mpi'
petsc_tpls = '+metis ~hdf5 ~hybre +superlu-dist +int64'
depends_on('petsc@3.13.0: ' + petsc_build_options + petsc_tpls, when='+petsc')

#
# Python
#
depends_on('python +shared +pic', when='+pygeosx')
depends_on('py-numpy@1.19: +blas +lapack +force-parallel-build', when='+pygeosx')
depends_on('py-scipy@1.5.2: +force-parallel-build', when='+pygeosx')
depends_on('py-mpi4py@3.0.3:', when='+pygeosx')
depends_on('py-pip', when='+pygeosx')

#
# Dev tools
#
depends_on('uncrustify@0.71:')

```

(continues on next page)

(continued from previous page)

```
#
# Documentation
#
depends_on('doxygen@1.8.13:', when='+docs', type='build')
depends_on('py-sphinx@1.6.3:', when='+docs', type='build')
```

Using the Spack spec syntax you can inturn specify variants for each of the dependencies of GEOSX. So for example if you could modify the spec above to build RAJA in debug by using `%gcc@8.3.1 ~caliper lai=petsc ^raja build_type=Debug`. When building with Uberenv Spack should print out a table containing the full spec for every dependency it will build. If you would like to look at the variants for say RAJA in more detail you can find the package file at `uberenv_libs/spack/var/spack/repos/builtin/packages/raja/package.py`.

8.4.2 Adding a Dependency (Advanced)

Adding a dependency to GEOSX is straight forward if the dependency already builds with Spack. If that is the case then all you need to do is add a `depends_on('cool-new-library')` to the GEOSX `package.py` file. If however the dependency doesn't have a Spack package, you will have to add one by creating a `cool-new-library/package.yaml` file in the `scripts/uberenv/packages` directory and adding the logic to build it there.

Oftentimes (unfortunately), even when a package already exists, it might not work out of the box for your system. In this case copy over the existing `package.py` file from the Spack repository into `scripts/uberenv/packages/cool-new-library/package.py`, as if you were adding a new package, and perform your modifications there. Once you have the package working, copy the package back into the Spack repository (running Uberenv should do this for you) and commit+push your changes to Spack.

8.5 Continuous Integration process

To save building time, the third party libraries (that do not change so often) and GEOSX are build separately.

Everytime a pull is requested in the TPL repository, docker images are generated and deployed on [dockerhub](#). The repository names (`ubuntu18.04-gcc8`, `centos7.7.1908-clang9.0.0`, `centos7.6.1810-gcc8.3.1-cuda10.1.243` etc.) obviously reflect the OS and the compiler flavour used. For each image, the unique `${TRAVIS_PULL_REQUEST}-${TRAVIS_BUILD_NUMBER}` tag is used so we can connect the related code source in a rather convenient way. Each docker contains the `org.opencontainers.image.created` and `org.opencontainers.image.revision` labels to provide additional information.

For the OSX builds, we construct a tarball of the TPLs and save them in a remote cloud storage. There is currently only one mac osx tested environment (xcode 11.2) and the same `${TRAVIS_PULL_REQUEST}-${TRAVIS_BUILD_NUMBER}` pattern is used as an identifier for the build. An important counterpart to using a tarball and not a docker image is that the tarball does not provide the whole system the precompiled binaries rely on. Problems may arise since we use the rolling release [Homebrew](#) (to install open-mpi in particular). To circumvent this potential issue, the brew version is fixed to a specific commit (see `BREW_HASH` variable in [third party's .travis.yml](#)) and stored as a metainformation of the tarball blob inside the cloud storage. It is therefore possible for GEOSX to recover this information and build against the same revision of brew packages. Note that the `TRAVIS_PULL_REQUEST`, `TRAVIS_BUILD_NUMBER` and `TRAVIS_COMMIT` are also stored as metainformation in the same way (have a look at the OSX build section of [GEOSX's .travis.yml](#) to see how to retrieve these informations).

There thus is only one unique identifier for both dockers and mac osx builds for one TPL code base. It is necessary to define the global environment `GEOSX_TPL_TAG` (e.g. something like 82-254) to build against one selected version of the TPL.

It must be mentioned that one and only one version of the compiled TPL tarball is stored per pull request (older ones are removed automatically). Therefore, a client building against a work in progress PR may experience a 404 error sooner or later.

8.5.1 Building docker images

Our continuous integration process builds the TPL and GEOSX against two operating systems (ubuntu and centos) and two compilers (clang and gcc). The docker files use [multi-stage builds](#) in order to minimise the sizes of the images.

- First stage installs and defines all the elements that are commons to both TPL and GEOSX (for example, MPI and c++ compiler, BLAS, LAPACK, path to the installation directory...).
- As a second stage, we install everything needed to build (*not run*) the TPLs. We keep nothing from this second step for GEOSX, except the compiled TPL themselves. For example, a fortran compiler is needed by the TPL but not by GEOSX: it shall be installed during this step, so GEOSX won't access a fortran compiler (it does not have to).
- Last stage copies the compiled TPL from second stage and installs the elements only required by GEOSX (there are few).

8.5.2 Docker images contract

GEOSX will find a compiled version of the third party libraries.

As part of the contract provided by the TPL, the docker images also defines several environment variables. The

```
GEOSX_TPL_DIR
```

variable contains the absolute path of the installation root directory of the third party libraries. GEOSX must use it when building.

Other variables are classical absolute path compiler variables.

```
CC
CXX
MPICC
MPICXX
```

And the absolute path the mpirun (or equivalent) command.

```
MPIEXEC
```

The following `openmpi` environment variables allow it to work properly in the docker container. But there should be no reason to access or use them explicitly.

```
OMPI_CC=$CC
OMPI_CXX=$CXX
```


9.1 Input Schema Definitions

XML Schema

9.2 Element: AcousticSEM

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
outputSismoTrace	localIndex	0	Flag that indicates if we write the sismo trace in a file .txt, 0 no output, 1 otherwise
receiverCoordinates	real64_array	2, required	Coordinates (x,y,z) of the receivers
rickerOrder	localIndex	2	Flag that indicates the order of the Ricker to be used 0, 1 or 2. Order 2 by default
sourceCoordinates	real64_array	2, required	Coordinates (x,y,z) of the sources
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
timeSourceFrequency	real64	required	Central frequency for the time source
LinearSolverParameters	node	unique	Element: LinearSolverParameters
NonlinearSolverParameters	node	unique	Element: NonlinearSolverParameters

9.3 Element: Benchmarks

Name	Type	Default	Description
lassen	node	unique	<i>Element: lassen</i>
quartz	node	unique	<i>Element: quartz</i>

9.4 Element: BiotPorosity

Name	Type	Default	Description
defaultReferencePorosity	real64	required	Default value of the reference porosity
grainBulkModulus	real64	required	Grain bulk modulus
name	string	required	A name is required for any non-unique nodes

9.5 Element: BlackOilFluid

Name	Type	Default	Description
componentMolarWeight	real64_array	required	Component molar weights
componentNames	string_array	{}	List of component names
hydrocarbonFormationVolumeFactorTableNames	string_array	{}	<p>List of formation volume factor TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
hydrocarbonViscosityTableNames	string_array	{}	<p>List of viscosity TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases
surfaceDensities	real64_array	required	List of surface mass densities for each phase
tableFiles	path_array	{}	List of filenames with input PVT tables (one per phase)
waterCompressibility	real64	0	Water compressibility
waterFormationVolumeFactor	real64	0	Water formation volume factor
waterReferencePressure	real64	0	Water reference pressure
waterViscosity	real64	0	Water viscosity

9.6 Element: Blueprint

Name	Type	De- fault	Description
childDirectory	string		Child directory path
name	string	re- quired	A name is required for any non-unique nodes
outputFullQuadratureData	integer	0	If true writes out data associated with every quadrature point.
parallelThreads	integer	1	Number of plot files.
plotLevel	geosx_dataRepository_PlotLevel		Determines which fields to write.

9.7 Element: BoundedPlane

Name	Type	De- fault	Description
dimensions	real64_array	re- quired	Length and width of the bounded plane
length-Vector	R1Tensor	re- quired	Tangent vector defining the orthonormal basis along with the normal.
name	string	re- quired	A name is required for any non-unique nodes
normal	R1Tensor	re- quired	Normal (n_x,n_y,n_z) to the plane (will be normalized automatically)
origin	R1Tensor	re- quired	Origin point (x,y,z) of the plane (basically, any point on the plane)
tolerance	real64	1e-05	Tolerance to determine if a point sits on the plane or not. It is relative to the maximum dimension of the plane.
widthVec- tor	R1Tensor	re- quired	Tangent vector defining the orthonormal basis along with the normal.

9.8 Element: Box

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
strike	real64	-90	The strike angle of the box
xMax	R1Tensor	required	Maximum (x,y,z) coordinates of the box
xMin	R1Tensor	required	Minimum (x,y,z) coordinates of the box

9.9 Element: BrooksCoreyBakerRelativePermeability

Name	Type	Default	Description
gasOilRelPermExponent	real64_array	{1}	Rel perm power law exponent for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasExp, oilExp }”, in that order
gasOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasMax, oilMax }”, in that order
name	string	required	A name is required for any non-unique nodes
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases
waterOilRelPermExponent	real64_array	{1}	Rel perm power law exponent for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterExp, oilExp }”, in that order
waterOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterMax, oilMax }”, in that order

9.10 Element: BrooksCoreyCapillaryPressure

Name	Type	De- fault	Description
capPressureEpsilon	real64	1e-06	Wetting-phase saturation at which the max cap. pressure is attained; used to avoid infinite cap. pressure values for saturations close to zero
name	string	re- quired	A name is required for any non-unique nodes
phaseCapPressureExponentInv	real64_array	{1}	Inverse of capillary power law exponent for each phase
phaseEntryPressure	real64_array	{1}	Entry pressure value for each phase
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	re- quired	List of fluid phases

9.11 Element: BrooksCoreyRelativePermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases
phaseRelPermExponent	real64_array	{1}	MinimumRel perm power law exponent for each phase
phaseRelPermMaxValue	real64_array	{0}	Maximum rel perm value for each phase

9.12 Element: CO2BrineFluid

Name	Type	De- fault	Description
componentMolarWeight	real64_array	{0}	Component molar weights
componentNames	string_array	{ }	List of component names
flashModelParaFile	path	re- quired	Name of the file defining the parameters of the flash model
name	string	re- quired	A name is required for any non-unique nodes
phaseNames	string_array	{ }	List of fluid phases
phasePVTParaFiles	path_array	re- quired	Names of the files defining the parameters of the viscosity and density models

9.13 Element: CarmanKozenyPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
particleDiameter	real64	required	Diameter of the spherical particles.
sphericity	real64	required	Sphericity of the particles.

9.14 Element: CellElementRegion

Name	Type	Default	Description
cellBlocks	string_array	{ }	(no description available)
coarseningRatio	real64	0	(no description available)
materialList	string_array	required	List of materials present in this region
name	string	required	A name is required for any non-unique nodes

9.15 Element: ChomboIO

Name	Type	Default	Description
beginCycle	real64	required	Cycle at which the coupling will commence.
childDirectory	string		Child directory path
inputPath	string	/IN- VALID_INPUT_PATH	Path at which the chombo to geosx file will be written.
name	string	required	A name is required for any non-unique nodes
outputPath	string	required	Path at which the geosx to chombo file will be written.
parallelThreads	integer	1	Number of plot files.
useChombo-Pressures	integer	0	True iff geosx should use the pressures chombo writes out.
waitForInput	integer	required	True iff geosx should wait for chombo to write out a file. When true the inputPath must be set.

9.16 Element: CompositeFunction

Name	Type	Default	Description
expression	string		Composite math expression
function-Names	string_array	{ }	List of source functions. The order must match the variableNames argument.
inputVar-Names	string_array	{ }	Name of fields are input to function.
name	string	required	A name is required for any non-unique nodes
variable-Names	string_array	{ }	List of variables in expression

9.17 Element: CompositionalMultiphaseFVM

Name	Type	De- fault	Description
allowLocalCompDensityChopping	integer	1	Flag indicating whether local (cell-wise) chopping of negative compositions is allowed
capPressureNames	string array		Name of the capillary pressure constitutive model to use
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
computeCFLNumbers	integer	0	Flag indicating whether CFL numbers are computed or not
discretization	string	required	Name of discretization object to use for this solver.
fluidNames	string array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
inputFluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
maxCompFractionChange	real64	1	Maximum (absolute) change in a component fraction between two Newton iterations
name	string	required	A name is required for any non-unique nodes
permeabilityNames	string array	required	Names of permeability constitutive models for each region.
relPermNames	string array	required	Name of the relative permeability constitutive model to use
solidNames	string array	required	Names of solid constitutive models for each region.
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
temperature	real64	required	Temperature
useMass	integer	0	Use mass formulation instead of molar
LinearSolverParameters	node	unique	Element: LinearSolverParameters
NonlinearSolverParameters	node	unique	Element: NonlinearSolverParameters

9.18 Element: CompositionalMultiphaseFluid

Name	Type	Default	Description
componentAcentricFactor	real64_array	required	Component acentric factors
componentBinaryCoeff	real64_array2d	{{0}}	Table of binary interaction coefficients
componentCriticalPressure	real64_array	required	Component critical pressures
componentCriticalTemperature	real64_array	required	Component critical temperatures
componentMolarWeight	real64_array	required	Component molar weights
componentNames	string_array	required	List of component names
componentVolumeShift	real64_array	{0}	Component volume shifts
equationsOfState	string_array	required	List of equation of state types for each phase
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases

9.19 Element: CompositionalMultiphaseHybridFVM

Name	Type	De- fault	Description
allowLocalCompDensityChopping	integer	1	Flag indicating whether local (cell-wise) chopping of negative compositions is allowed
capPressureNames	string array		Name of the capillary pressure constitutive model to use
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
computeCFLNumbers	integer	0	Flag indicating whether CFL numbers are computed or not
discretization	string	required	Name of discretization object to use for this solver.
fluidNames	string array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
inputFluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
maxCompFractionChange	real64	1	Maximum (absolute) change in a component fraction between two Newton iterations
maxRelativePressureChange	real64	1	Maximum (relative) change in (face) pressure between two Newton iterations
name	string	required	A name is required for any non-unique nodes
permeabilityNames	string array	required	Names of permeability constitutive models for each region.
relPermNames	string array	required	Name of the relative permeability constitutive model to use
solidNames	string array	required	Names of solid constitutive models for each region.
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
temperature	real64	required	Temperature
useMass	integer	0	Use mass formulation instead of molar
LinearSolverParameters	node	unique	Element: LinearSolverParameters
NonlinearSolverParameters	node	unique	Element: NonlinearSolverParameters

9.20 Element: CompositionalMultiphaseReservoir

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
flow-Solver-Name	string	required	Name of the flow solver to use in the reservoir-well system solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
well-Solver-Name	string	required	Name of the well solver to use in the reservoir-well system solver
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.21 Element: CompositionalMultiphaseWell

Name	Type	De- fault	Description
allowLocalCompDensityChopping	integer	1	Flag indicating whether local (cell-wise) chopping of negative compositions is allowed
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fluidNames	string_array	required	Name of fluid constitutive object to use for this solver.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
maxCompFractionChange	real64	1	Maximum (absolute) change in a component fraction between two Newton iterations
maxRelativePressureChange	real64	1	Maximum (relative) change in pressure between two Newton iterations (recommended with rate control)
name	string	required	A name is required for any non-unique nodes
relPermNames	string_array	required	Names of relative permeability constitutive models to use
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
useMass	integer	0	Use mass formulation instead of molar
wellTemperature	real64	required	Temperature
LinearSolverParameters	node	unique	Element: LinearSolverParameters
NonlinearSolverParameters	node	unique	Element: NonlinearSolverParameters
WellControls	node		Element: WellControls

9.22 Element: CompressibleSinglePhaseFluid

Name	Type	Default	Description
compressibility	real64	0	Fluid compressibility
defaultDensity	real64	required	Default value for density.
defaultViscosity	real64	required	Default value for viscosity.
densityModelType	geosx_constitutive_ExponentialApproximationType		Type of density model. Valid options: * exponential * linear * quadratic
name	string	required	A name is required for any non-unique nodes
referenceDensity	real64	1000	Reference fluid density
referencePressure	real64	0	Reference pressure
referenceViscosity	real64	0.001	Reference fluid viscosity
viscosibility	real64	0	Fluid viscosity exponential coefficient
viscosityModelType	geosx_constitutive_ExponentialApproximationType		Type of viscosity model. Valid options: * exponential * linear * quadratic

9.23 Element: CompressibleSolidCarmanKozenyPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.24 Element: CompressibleSolidConstantPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.25 Element: CompressibleSolidParallelPlatesPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.26 Element: ConstantPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityComponents	RTensor	required	xx, yy and zz components of a diagonal permeability tensor.

9.27 Element: Constitutive

Name	Type	Default	Description
BiotPorosity	node		<i>Element: BiotPorosity</i>
BlackOilFluid	node		<i>Element: BlackOilFluid</i>
BrooksCoreyBakerRelativePermeability	node		<i>Element: BrooksCoreyBakerRelativePermeability</i>
BrooksCoreyCapillaryPressure	node		<i>Element: BrooksCoreyCapillaryPressure</i>
BrooksCoreyRelativePermeability	node		<i>Element: BrooksCoreyRelativePermeability</i>
CO2BrineFluid	node		<i>Element: CO2BrineFluid</i>
CarmanKozenyPermeability	node		<i>Element: CarmanKozenyPermeability</i>
CompositionalMultiphaseFluid	node		<i>Element: CompositionalMultiphaseFluid</i>
CompressibleSinglePhaseFluid	node		<i>Element: CompressibleSinglePhaseFluid</i>
CompressibleSolidCarmanKozenyPermeability	node		<i>Element: CompressibleSolidCarmanKozenyPermeability</i>
CompressibleSolidConstantPermeability	node		<i>Element: CompressibleSolidConstantPermeability</i>
CompressibleSolidParallelPlatesPermeability	node		<i>Element: CompressibleSolidParallelPlatesPermeability</i>
ConstantPermeability	node		<i>Element: ConstantPermeability</i>
Contact	node		<i>Element: Contact</i>
Coulomb	node		<i>Element: Coulomb</i>
DamageElasticIsotropic	node		<i>Element: DamageElasticIsotropic</i>
DamageSpectralElasticIsotropic	node		<i>Element: DamageSpectralElasticIsotropic</i>
DamageVolDevElasticIsotropic	node		<i>Element: DamageVolDevElasticIsotropic</i>
DeadOilFluid	node		<i>Element: DeadOilFluid</i>
DruckerPrager	node		<i>Element: DruckerPrager</i>
ElasticIsotropic	node		<i>Element: ElasticIsotropic</i>
ElasticOrthotropic	node		<i>Element: ElasticOrthotropic</i>
ElasticTransverseIsotropic	node		<i>Element: ElasticTransverseIsotropic</i>
ExtendedDruckerPrager	node		<i>Element: ExtendedDruckerPrager</i>
NullModel	node		<i>Element: NullModel</i>
ParallelPlatesPermeability	node		<i>Element: ParallelPlatesPermeability</i>
ParticleFluid	node		<i>Element: ParticleFluid</i>
PermeabilityBase	node		<i>Element: PermeabilityBase</i>

Continued on next page

Table 9.1 – continued from previous page

Name	Type	Default	Description
PorousDruckerPrager	node		<i>Element: PorousDruckerPrager</i>
PorousElasticIsotropic	node		<i>Element: PorousElasticIsotropic</i>
PorousElasticOrthotropic	node		<i>Element: PorousElasticOrthotropic</i>
PorousElasticTransverseIsotropic	node		<i>Element: PorousElasticTransverseIsotropic</i>
PorousExtendedDruckerPrager	node		<i>Element: PorousExtendedDruckerPrager</i>
PressurePorosity	node		<i>Element: PressurePorosity</i>
ProppantSlurryFluid	node		<i>Element: ProppantSlurryFluid</i>
StrainDependentPermeability	node		<i>Element: StrainDependentPermeability</i>
TableRelativePermeability	node		<i>Element: TableRelativePermeability</i>
VanGenuchtenBakerRelativePermeability	node		<i>Element: VanGenuchtenBakerRelativePermeability</i>
VanGenuchtenCapillaryPressure	node		<i>Element: VanGenuchtenCapillaryPressure</i>

9.28 Element: Contact

Name	Type	De- fault	Description
aper- ture- Tol- er- ance	real64	1e-09	Value to be used to avoid floating point errors in expressions involving aperture. For example in the case of dividing by the actual aperture (not the effective aperture that results from the aperture function) this value may be used to avoid 1/0 errors. Note that this value may have some physical significance in its usage, as it may be used to smooth out highly nonlinear behavior associated with 1/0 in addition to avoiding the 1/0 error.
name	string	re- quired	A name is required for any non-unique nodes
penal- tyS- tiff- ness	real64	0	Value of the penetration penalty stiffness. Units of Pressure/length
Table- Func- tion	node		<i>Element: TableFunction</i>

9.29 Element: Coulomb

Name	Type	De- fault	Description
aperture-Tolerance	real64	1e-09	Value to be used to avoid floating point errors in expressions involving aperture. For example in the case of dividing by the actual aperture (not the effective aperture that results from the aperture function) this value may be used to avoid 1/0 errors. Note that this value may have some physical significance in its usage, as it may be used to smooth out highly nonlinear behavior associated with 1/0 in addition to avoiding the 1/0 error.
cohesion	real64	required	Cohesion
frictionAngle	real64	-1	Friction Angle (in radians)
frictionCoefficient	real64	-1	Friction Coefficient
name	string	required	A name is required for any non-unique nodes
penaltyStiffness	real64	0	Value of the penetration penalty stiffness. Units of Pressure/length
TableFunction	node		<i>Element: TableFunction</i>

9.30 Element: Cylinder

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
point1	R1Tensor	required	Center point of one (upper or lower) face of the cylinder
point2	R1Tensor	required	Center point of the other face of the cylinder
radius	real64	required	Radius of the cylinder

9.31 Element: DamageElasticIsotropic

Name	Type	Default	Description
criticalFractureEnergy	real64	required	Critical fracture energy
criticalStrainEnergy	real64	required	Critical stress in a 1d tension test
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultDensity	real64	required	Default Material Density
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
lengthScale	real64	required	Length scale l in the phase-field equation
name	string	required	A name is required for any non-unique nodes

9.32 Element: DamageSpectralElasticIsotropic

Name	Type	Default	Description
criticalFractureEnergy	real64	required	Critical fracture energy
criticalStrainEnergy	real64	required	Critical stress in a 1d tension test
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultDensity	real64	required	Default Material Density
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
lengthScale	real64	required	Length scale l in the phase-field equation
name	string	required	A name is required for any non-unique nodes

9.33 Element: DamageVolDevElasticIsotropic

Name	Type	Default	Description
criticalFractureEnergy	real64	required	Critical fracture energy
criticalStrainEnergy	real64	required	Critical stress in a 1d tension test
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultDensity	real64	required	Default Material Density
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
lengthScale	real64	required	Length scale l in the phase-field equation
name	string	required	A name is required for any non-unique nodes

9.34 Element: DeadOilFluid

Name	Type	Default	Description
componentMolarWeight	real64_array	required	Component molar weights
componentNames	string_array	{}	List of component names
hydrocarbonFormationVolumeFactorTableNames	string_array	{}	<p>List of formation volume factor TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
hydrocarbonViscosityTableNames	string_array	{}	<p>List of viscosity TableFunction names from the Functions block.</p> <p>The user must provide one TableFunction per hydrocarbon phase, in the order provided in “phaseNames”.</p> <p>For instance, if “oil” is before “gas” in “phaseNames”, the table order should be: oilTableName, gasTableName</p>
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases
surfaceDensities	real64_array	required	List of surface mass densities for each phase
tableFiles	path_array	{}	List of filenames with input PVT tables (one per phase)
waterCompressibility	real64	0	Water compressibility
waterFormationVolumeFactor	real64	0	Water formation volume factor
waterReferencePressure	real64	0	Water reference pressure
waterViscosity	real64	0	Water viscosity

9.35 Element: Dirichlet

Name	Type	Default	Description
bcApplicationTable-Name	string		Name of table that specifies the on/off application of the bc.
beginTime	real64	-1e+99	time at which BC will start being applied.
component	integer	-1	Component of field (if tensor) to apply boundary condition to
direction	R1Tensor	{0,0,0}	Direction to apply boundary condition to
endTime	real64	1e+99	time at which bc will stop being applied
fieldName	string		Name of field that boundary condition is applied to.
functionName	string		Name of function that specifies variation of the BC
initialCondition	integer	0	BC is applied as an initial condition.
name	string	required	A name is required for any non-unique nodes
objectPath	string		Path to the target field
scale	real64	0	Scale factor for value of BC.
setNames	string_array	required	Name of sets that boundary condition is applied to.

9.36 Element: DruckerPrager

Name	Type	Default	Description
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultCohesion	real64	0	Initial cohesion
defaultDensity	real64	required	Default Material Density
defaultDilationAngle	real64	30	Dilation angle (degrees)
defaultFrictionAngle	real64	30	Friction angle (degrees)
defaultHardeningRate	real64	0	Cohesion hardening/softening rate
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
name	string	required	A name is required for any non-unique nodes

9.37 Element: ElasticIsotropic

Name	Type	Default	Description
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultDensity	real64	required	Default Material Density
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
name	string	required	A name is required for any non-unique nodes

9.38 Element: ElasticOrthotropic

Name	Type	Default	Description
defaultC11	real64	-1	Default C11 Component of Voigt Stiffness Tensor
defaultC12	real64	-1	Default C12 Component of Voigt Stiffness Tensor
defaultC13	real64	-1	Default C13 Component of Voigt Stiffness Tensor
defaultC22	real64	-1	Default C22 Component of Voigt Stiffness Tensor
defaultC23	real64	-1	Default C23 Component of Voigt Stiffness Tensor
defaultC33	real64	-1	Default C33 Component of Voigt Stiffness Tensor
defaultC44	real64	-1	Default C44 Component of Voigt Stiffness Tensor
defaultC55	real64	-1	Default C55 Component of Voigt Stiffness Tensor
defaultC66	real64	-1	Default C66 Component of Voigt Stiffness Tensor
defaultDensity	real64	required	Default Material Density
defaultE1	real64	-1	Default Young's Modulus E1
defaultE2	real64	-1	Default Young's Modulus E2
defaultE3	real64	-1	Default Young's Modulus E3
defaultG12	real64	-1	Default Shear Modulus G12
defaultG13	real64	-1	Default Shear Modulus G13
defaultG23	real64	-1	Default Shear Modulus G23
defaultNu12	real64	-1	Default Poission's Ratio Nu12
defaultNu13	real64	-1	Default Poission's Ratio Nu13
defaultNu23	real64	-1	Default Poission's Ratio Nu23
name	string	required	A name is required for any non-unique nodes

9.39 Element: ElasticTransverselsotropic

Name	Type	Default	Description
defaultC11	real64	-1	Default Stiffness Parameter C11
defaultC13	real64	-1	Default Stiffness Parameter C13
defaultC33	real64	-1	Default Stiffness Parameter C33
defaultC44	real64	-1	Default Stiffness Parameter C44
defaultC66	real64	-1	Default Stiffness Parameter C66
defaultDensity	real64	required	Default Material Density
defaultPoissonRatioAxialTransverse	real64	-1	Default Axial-Transverse Poisson's Ratio
defaultPoissonRatioTransverse	real64	-1	Default Transverse Poisson's Ratio
defaultShearModulusAxialTransverse	real64	-1	Default Axial-Transverse Shear Modulus
defaultYoungModulusAxial	real64	-1	Default Axial Young's Modulus
defaultYoungModulusTransverse	real64	-1	Default Transverse Young's Modulus
name	string	required	A name is required for any non-unique nodes

9.40 Element: ElementRegions

Name	Type	Default	Description
CellElementRegion	node		<i>Element: CellElementRegion</i>
SurfaceElementRegion	node		<i>Element: SurfaceElementRegion</i>
WellElementRegion	node		<i>Element: WellElementRegion</i>

9.41 Element: EmbeddedSurfaceGenerator

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fractureRegion	string	FractureRegion	(no description available)
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
mpiCommOrder	integer	0	Flag to enable MPI consistent communication ordering
name	string	required	A name is required for any non-unique nodes
solidMaterialNames	string array	required	Name of the solid material used in solid mechanic solver
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
LinearSolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
NonlinearSolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.42 Element: Events

Name	Type	Default	Description
logLevel	integer	0	Log level
maxCycle	integer	2147483647	Maximum simulation cycle for the global event loop.
maxTime	real64	1.79769e+308	Maximum simulation time for the global event loop.
HaltEvent	node		<i>Element: HaltEvent</i>
PeriodicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

9.43 Element: ExtendedDruckerPrager

Name	Type	Default	Description
defaultBulkModulus	real64	-1	Default Bulk Modulus Parameter
defaultCohesion	real64	0	Initial cohesion
defaultDensity	real64	re- quired	Default Material Density
defaultDilationRatio	real64	1	Dilation ratio [0,1] (ratio = $\tan \text{dilationAngle} / \tan \text{frictionAngle}$)
defaultHardening	real64	0	Hardening parameter (hardening rate is faster for smaller values)
defaultInitialFrictionAngle	real64	30	Initial friction angle (degrees)
defaultPoissonRatio	real64	-1	Default Poisson's Ratio
defaultResidualFrictionAngle	real64	30	Residual friction angle (degrees)
defaultShearModulus	real64	-1	Default Shear Modulus Parameter
defaultYoungModulus	real64	-1	Default Young's Modulus
name	string	re- quired	A name is required for any non-unique nodes

9.44 Element: FieldSpecification

Name	Type	Default	Description
bcApplicationTableName	string		Name of table that specifies the on/off application of the bc.
beginTime	real64	-1e+99	time at which BC will start being applied.
component	integer	-1	Component of field (if tensor) to apply boundary condition to
direction	R1Tensor	{0,0,0}	Direction to apply boundary condition to
endTime	real64	1e+99	time at which bc will stop being applied
fieldName	string		Name of field that boundary condition is applied to.
functionName	string		Name of function that specifies variation of the BC
initialCondition	integer	0	BC is applied as an initial condition.
name	string	re- quired	A name is required for any non-unique nodes
objectPath	string		Path to the target field
scale	real64	0	Scale factor for value of BC.
setNames	string_array	re- quired	Name of sets that boundary condition is applied to.

9.45 Element: FieldSpecifications

Name	Type	Default	Description
Dirichlet	node		<i>Element: Dirichlet</i>
FieldSpecification	node		<i>Element: FieldSpecification</i>
SourceFlux	node		<i>Element: SourceFlux</i>
Traction	node		<i>Element: Traction</i>

9.46 Element: File

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

9.47 Element: FiniteElementSpace

Name	Type	Default	Description
for-mulation	string	de-fault	Specifier to indicate any specialized formulations. For instance, one of the many enhanced assumed strain methods of the Hexahedron parent shape would be indicated here
name	string	re-quired	A name is required for any non-unique nodes
order	in-te-ger	re-quired	The order of the finite element basis.

9.48 Element: FiniteElements

Name	Type	Default	Description
FiniteElementSpace	node		<i>Element: FiniteElementSpace</i>
LinearSolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
NonlinearSolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.49 Element: FiniteVolume

Name	Type	Default	Description
HybridMimeticDiscretization	node		<i>Element: HybridMimeticDiscretization</i>
TwoPointFluxApproximation	node		<i>Element: TwoPointFluxApproximation</i>

9.50 Element: FlowProppantTransport

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
flow-Solver-Name	string	re- quired	Name of the flow solver to use in the flowProppantTransport solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	re- quired	A name is required for any non-unique nodes
proppant-Solver-Name	string	re- quired	Name of the proppant transport solver to use in the flowProppantTransport solver
targetRegions	string array	re- quired	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.51 Element: Functions

Name	Type	Default	Description
CompositeFunction	node		<i>Element: CompositeFunction</i>
SymbolicFunction	node		<i>Element: SymbolicFunction</i>
TableFunction	node		<i>Element: TableFunction</i>

9.52 Element: Geometry

Name	Type	Default	Description
BoundedPlane	node		<i>Element: BoundedPlane</i>
Box	node		<i>Element: Box</i>
Cylinder	node		<i>Element: Cylinder</i>
ThickPlane	node		<i>Element: ThickPlane</i>

9.53 Element: HaltEvent

Name	Type	De- fault	Description
beginTime	real64	0	Start time of this event.
endTime	real64	1e+100	End time of this event.
finalDtStretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the endTime exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
logLevel	integer	0	Log level
maxEventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
maxRuntime	real64	required	The maximum allowable runtime for the job.
name	string	required	A name is required for any non-unique nodes
target	string		Name of the object to be executed when the event criteria are met.
targetExact-StartStop	integer	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
HaltEvent	node		<i>Element: HaltEvent</i>
PeriodicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

9.54 Element: HybridMimeticDiscretization

Name	Type	Default	Description
coefficientName	string	required	Name of coefficient field
innerProductType	string	required	Type of inner product used in the hybrid FVM solver
name	string	required	A name is required for any non-unique nodes

9.55 Element: Hydrofracture

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRelationName	string	required	Name of contact relation to enforce constraints on fracture boundary.
couplingTypeOption	geosx_HydrofractureSolver_CouplingTypeOption	CouplingTypeOption	Coupling method. Valid options: * FIM * SIM_FixedStress
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
fluidSolverName	string	required	Name of the fluid mechanics solver to use in the poromechanics solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
maxNumResolves	integer	10	Value to indicate how many resolves may be executed to perform surface generation after the execution of flow and mechanics solver.
name	string	required	A name is required for any non-unique nodes
porousMaterialNames	string_array	required	The name of the material that should be used in the constitutive updates
solidSolverName	string	required	Name of the solid mechanics solver to use in the poromechanics solver
surfaceGeneratorName	string	required	Name of the surface generator to use in the hydrofracture solver
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will

9.56 Element: Included

Name	Type	Default	Description
File	node		<i>Element: File</i>

9.57 Element: InternalMesh

Name	Type	Default	Description
cellBlock-Names	string_array	required	Names of each mesh block
element-Types	string_array	required	Element types of each mesh block
name	string	required	A name is required for any non-unique nodes
nx	integer_array	required	Number of elements in the x-direction within each mesh block
ny	integer_array	required	Number of elements in the y-direction within each mesh block
nz	integer_array	required	Number of elements in the z-direction within each mesh block
positionTolerance	real64	1e-10	A position tolerance to verify if a node belong to a nodeset
trianglePattern	integer	0	Pattern by which to decompose the hex mesh into prisms (more explanation required)
xBias	real64_array{1}		Bias of element sizes in the x-direction within each mesh block ($dx_left=(1+b)*L/N$, $dx_right=(1-b)*L/N$)
xCords	real64_array	required	x-coordinates of each mesh block vertex
yBias	real64_array{1}		Bias of element sizes in the y-direction within each mesh block ($dy_left=(1+b)*L/N$, $dy_right=(1-b)*L/N$)
yCords	real64_array	required	y-coordinates of each mesh block vertex
zBias	real64_array{1}		Bias of element sizes in the z-direction within each mesh block ($dz_left=(1+b)*L/N$, $dz_right=(1-b)*L/N$)
zCords	real64_array	required	z-coordinates of each mesh block vertex

9.58 Element: InternalWell

Name	Type	Default	Description
logLevel	integer	0	Log level
meshName	string	re- quired	Name of the reservoir mesh associated with this well
name	string	re- quired	A name is required for any non-unique nodes
numElementsPerSegment	integer	re- quired	Number of well elements per polyline segment
polylineNodeCoords	real64_array2d	re- quired	Physical coordinates of the well polyline nodes
polylineSegmentConn	globalIndex_array2d	re- quired	Connectivity of the polyline segments
radius	real64	re- quired	Radius of the well
wellControlsName	string	re- quired	Name of the set of constraints associated with this well
wellRegionName	string	re- quired	Name of the well element region
Perforation	node		<i>Element: Perforation</i>

9.59 Element: InternalWellbore

Name	Type	De- fault	Description
autoSpaceRadialElems	integer_array	{0}	Automatically set number and spacing of elements in the radial direction. This overrides the values of nr
cellBlockNames	string_array	required	Names of each mesh block
elementTypes	string_array	required	Element types of each mesh block
name	string	required	A name is required for any non-unique nodes
nr	integer_array	required	Number of elements in the radial direction
nt	integer_array	required	Number of elements in the tangent direction
nz	integer_array	required	Number of elements in the z-direction within each mesh block
positionTolerance	real64	1e-10	A position tolerance to verify if a node belong to a nodeset
rBias	real64_array	{-0.8}	Bias of element sizes in the radial direction
radius	real64_array	required	Wellbore radius
theta	real64_array	required	Tangent angle defining geometry size: 90 for quarter, 180 for half and 360 for full wellbore geometry
trajectory	real64_array	{0}	Coordinates defining the wellbore trajectory
trianglePattern	integer	0	Pattern by which to decompose the hex mesh into prisms (more explanation required)
useCartesianOuterBoundary	integer	1000000	Enforce a Cartesian aligned outer boundary on the outer block starting with the radial block specified in this value
xBias	real64_array	{1}	Bias of element sizes in the x-direction within each mesh block ($dx_{left}=(1+b)*L/N$, $dx_{right}=(1-b)*L/N$)
yBias	real64_array	{1}	Bias of element sizes in the y-direction within each mesh block ($dy_{left}=(1+b)*L/N$, $dy_{right}=(1-b)*L/N$)
zBias	real64_array	{1}	Bias of element sizes in the z-direction within each mesh block ($dz_{left}=(1+b)*L/N$, $dz_{right}=(1-b)*L/N$)
zCoords	real64_array	required	z-coordinates of each mesh block vertex

9.60 Element: LagrangianContact

Name	Type	Default	Description
activeSet- MaxIter	integer	10	Maximum number of iteration for the active set strategy in the lagrangian contact solver
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRela- tion- Name	string	required	Name of the constitutive law used for fracture elements
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
solid- Solver- Name	string	required	Name of the solid mechanics solver to use in the lagrangian contact solver
stabiliza- tion- Name	string	required	Name of the stabilization to use in the lagrangian contact solver
targetRe- gions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear- Solver- Param- eters	node	unique	Element: LinearSolverParameters
Non- linear- Solver- Param- eters	node	unique	Element: NonlinearSolverParameters

9.61 Element: LaplaceFEM

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
fieldName	string	required	Name of field variable
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
timeIntegrationOption	geosx_LaplaceBaseH1_TimeIntegrationOption	required	Time integration method. Options are: * SteadyState * ImplicitTransient
LinearSolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
NonlinearSolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.62 Element: LaplaceVEM

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
fieldName	string	required	Name of field variable
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
timeIntegrationOption	geosx_LaplaceBaseH1_TimeIntegrationOption	required	Time integration method. Options are: * SteadyState * ImplicitTransient
LinearSolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
NonlinearSolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.63 Element: LinearSolverParameters

Name	Type	Default	Description
amgAggressiveCoarseningLevel	integer	0	AMG number levels for aggressive coarsening Available options are: TODO
amgCoarseSolver	geosx_LinearSolverParameters AMG_CoarseType	direct	AMG coarsest level solver/smoother type. Available options are: default\ jacobi\ llt
amgCoarseningType	string	HMIS	AMG coarsening algorithm Available options are: TODO
amgInterpolationType	integer	6	AMG interpolation algorithm Available options are: TODO
amgNullSpaceType	geosx_LinearSolverParameters AMG_NullSpaceType	constantModes	AMG near null space approximation. Available options are: constantModes\ rigidBodyModes
amgNumFunctions	integer	1	AMG number of functions Available options are: TODO
amgNumSweeps	integer	2	AMG smoother sweeps
amgSmootherType	geosx_LinearSolverParameters AMG_SmootherType	gs	AMG smoother type. Available options are: default\ jacobi\ llt
amgThreshold	real64	0	AMG strength-of-connection threshold
directCheckResidual	integer	0	Whether to check the linear system solution residual
directColPerm	geosx_LinearSolverParameters Direct_ColPerm	none	How to permute the columns. Available options are: none\ MMD_AtplusA\ MMD_AtA\ col
directEquil	integer	1	Whether to scale the rows and columns of the matrix
directIterRef	integer	1	Whether to perform iterative refinement
directParallel	integer	1	Whether to use a parallel solver (instead of a serial one)
directRepTinyPivot	integer	1	Whether to raise

9.64 Element: Mesh

Name	Type	Default	Description
InternalMesh	node		<i>Element: InternalMesh</i>
InternalWell	node		<i>Element: InternalWell</i>
InternalWellbore	node		<i>Element: InternalWellbore</i>
PAMELAMeshGenerator	node		<i>Element: PAMELAMeshGenerator</i>

9.65 Element: MultiphasePoromechanics

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
fluid-Solver-Name	string	required	Name of the fluid mechanics solver to use in the poroelastic solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
porous-Material-Names	string array	required	The name of the material that should be used in the constitutive updates
solid-Solver-Name	string	required	Name of the solid mechanics solver to use in the poroelastic solver
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.66 Element: NonlinearSolverParameters

Name	Type	Default	Description
allowNonConverged	integer	0	Allow non-converged solution to be accepted. (i.e. exit from the Newton loop without achieving the desired tolerance)
dtCutIterLimit	real64	0.7	Fraction of the Max Newton iterations above which the solver asks for the time-step to be cut for the next dt.
dtIncIterLimit	real64	0.4	Fraction of the Max Newton iterations below which the solver asks for the time-step to be doubled for the next dt.
lineSearchAction	geosx_NonlinearSolverParameters.lineSearchAction	Attempt	How the line search is to be used. Options are: <ul style="list-style-type: none"> * None - Do not use line search. * Attempt - Use line search. Allow exit from line search without achieving smaller residual than starting residual. * Require - Use line search. If smaller residual than starting residual is not achieved, cut time step.
lineSearchCutFactor	real64	0.5	Line search cut factor. For instance, a value of 0.5 will result in the effective application of the last solution by a factor of (0.5, 0.25, 0.125, ...)
lineSearchMaxCuts	integer	4	Maximum number of line search cuts.
logLevel	integer	0	Log level
maxSubSteps	integer	10	Maximum number of time sub-steps allowed for the solver
maxTimeStepCuts	integer	2	Max number of time step cuts
newtonMaxIter	integer	5	Maximum number of iterations that are allowed in a Newton loop.
newtonMinIter	integer	1	Minimum number of iterations that are required before exiting the Newton
9.66. Element: NonlinearSolverParameters			loop. 379
newtonTol	real64	1e-06	The required tolerance in order to exit the Newton iteration loop.
timeStepCutFactor	real64	0.5	Factor by which the time

9.67 Element: NullModel

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

9.68 Element: NumericalMethods

Name	Type	Default	Description
FiniteElements	node	unique	<i>Element: FiniteElements</i>
FiniteVolume	node	unique	<i>Element: FiniteVolume</i>

9.69 Element: Outputs

Name	Type	Default	Description
Blueprint	node		<i>Element: Blueprint</i>
ChomboIO	node		<i>Element: ChomboIO</i>
Python	node		<i>Element: Python</i>
Restart	node		<i>Element: Restart</i>
Silo	node		<i>Element: Silo</i>
TimeHistory	node		<i>Element: TimeHistory</i>
VTK	node		<i>Element: VTK</i>

9.70 Element: PAMELAMeshGenerator

Name	Type	Default	Description
fieldNamesInGEOSX	string_array	{ }	Name of the fields within GEOSX
fieldsToImport	string_array	{ }	Fields to be imported from the external mesh file
file	path	required	path to the mesh file
name	string	required	A name is required for any non-unique nodes
reverseZ	integer	0	0 : Z coordinate is upward, 1 : Z coordinate is downward
scale	real64	1	Scale the coordinates of the vertices

9.71 Element: PackCollection

Name	Type	De- fault	Description
fieldName	string	re- quired	The name of the (packable) field associated with the specified object to retrieve data from
name	string	re- quired	A name is required for any non-unique nodes
objectPath	string	re- quired	The name of the object from which to retrieve field values.
onlyOn- SetChange	localIn- dex	0	Whether or not to only collect when the collected sets of indices change in any way.
setNames	string_array	{ }	The set(s) for which to retrieve data.

9.72 Element: ParallelPlatesPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

9.73 Element: Parameter

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
value	string	required	Input parameter definition for the preprocessor

9.74 Element: Parameters

Name	Type	Default	Description
Parameter	node		<i>Element: Parameter</i>

9.75 Element: ParticleFluid

Name	Type	Default	Description
collisionAlpha	real64	1.27	Collision alpha coefficient
collisionBeta	real64	1.5	Collision beta coefficient
fluidViscosity	real64	0.001	Fluid viscosity
hinderedSettlingCoefficient	real64	5.9	Hindered settling coefficient
isCollisionalSlip	integer	0	Whether the collisional component of the slip velocity is considered
maxProppantConcentration	real64	0.6	Max proppant concentration
name	string	required	A name is required for any non-unique nodes
particleSettlingModel	geosx_constitutive_ParticleSettlingModel	required	Particle settling velocity model. Valid options: * Stokes * Intermediate * Turbulence
proppantDensity	real64	1400	Proppant density
proppantDiameter	real64	0.0002	Proppant diameter
slipConcentration	real64	0.1	Slip concentration
sphericity	real64	1	Sphericity

9.76 Element: Perforation

Name	Type	Default	Description
distanceFromHead	real64	required	Linear distance from well head to the perforation
name	string	required	A name is required for any non-unique nodes
transmissibility	real64	-1	Perforation transmissibility

9.77 Element: PeriodicEvent

Name	Type	De- fault	Description
begin-Time	real64	0	Start time of this event.
cycleFrequency	integer	1	Event application frequency (cycle, default)
endTime	real64	1e+100	End time of this event.
finalDt-Stretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the endTime exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
function	string		Name of an optional function to evaluate when the time/cycle criteria are met.If the result is greater than the specified eventThreshold, the function will continue to execute.
logLevel	integer	0	Log level
max-EventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
name	string	required	A name is required for any non-unique nodes
object	string		If the optional function requires an object as an input, specify its path here.
set	string		If the optional function is applied to an object, specify the setname to evaluate (default = everything).
stat	integer	0	If the optional function is applied to an object, specify the statistic to compare to the eventThreshold.The current options include: min, avg, and max.
target	string		Name of the object to be executed when the event criteria are met.
targetExactStart-Stop	integer	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
targetExact-Timestep	integer	1	If this option is set, the event will reduce its timestep requests to match the specified timeFrequency perfectly: $dt_request = \min(dt_request, t_last + time_frequency - time)$.
threshold	real64	0	If the optional function is used, the event will execute if the value returned by the function exceeds this threshold.
timeFrequency	real64	-1	Event application frequency (time). Note: if this value is specified, it will override any cycle-based behavior.
HaltEvent	node		<i>Element: HaltEvent</i>
PeriodicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

9.78 Element: PermeabilityBase

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

9.79 Element: PhaseFieldDamageFEM

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
field-Name	string	required	name of field variable
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
localDissipation	string	required	Type of local dissipation function. Can be Linear or Quadratic
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
solid-Material-Names	string_array	required	name of solid constitutive model
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
timeIntegrationOption	string	required	option for default time integration method
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

9.80 Element: PhaseFieldFracture

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
couplingTypeOption	geosx_PhaseFieldFractureSolverCouplingTypeOption	required	Coupling option. Valid options: * FixedStress * TightlyCoupled
damageSolverName	string	required	Name of the damage mechanics solver to use in the PhaseFieldFracture solver
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
solidSolverName	string	required	Name of the solid mechanics solver to use in the PhaseFieldFracture solver
subcycling	integer	required	turn on subcycling on each load step
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.

LinearSolverParameters	node	unique	Element: LinearSolverParameters
NonlinearSolverParameters	node	unique	Element: NonlinearSolverParameters

9.81 Element: PorousDruckerPrager

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.82 Element: PorousElasticIsotropic

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.83 Element: PorousElasticOrthotropic

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.84 Element: PorousElasticTransverselIsotropic

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.85 Element: PorousExtendedDruckerPrager

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
permeabilityModelName	string	required	Name of the permeability model.
porosityModelName	string	required	Name of the porosity model.
solidModelName	string	required	Name of the solid model.

9.86 Element: PressurePorosity

Name	Type	Default	Description
compressibility	real64	required	Solid compressibility
defaultReferencePorosity	real64	required	Default value of the reference porosity
name	string	required	A name is required for any non-unique nodes
referencePressure	real64	required	Reference pressure for solid compressibility

9.87 Element: Problem

Name	Type	Default	Description
Benchmarks	node	unique	Element: Benchmarks
Constitutive	node	unique	Element: Constitutive
ElementRegions	node	unique	Element: ElementRegions
Events	node	unique, required	Element: Events
FieldSpecifications	node	unique	Element: FieldSpecifications
Functions	node	unique	Element: Functions
Geometry	node	unique	Element: Geometry
Included	node	unique	Element: Included
Mesh	node	unique, required	Element: Mesh
NumericalMethods	node	unique	Element: NumericalMethods
Outputs	node	unique, required	Element: Outputs
Parameters	node	unique	Element: Parameters
Solvers	node	unique, required	Element: Solvers
Tasks	node	unique	Element: Tasks

9.88 Element: ProppantSlurryFluid

Name	Type	Default	Description
componentNames	string_array	{ }	List of fluid component names
compressibility	real64	0	Fluid compressibility
defaultCompressibility	real64_array	{0}	Default value for compressibility.
defaultDensity	real64_array	{0}	Default value for density.
defaultViscosity	real64_array	{0}	Default value for viscosity.
flowBehaviorIndex	real64_array	{0}	Flow behavior index
flowConsistencyIndex	real64_array	{0}	Flow consistency index
maxProppantConcentration	real64	0.6	Maximum proppant concentration
name	string	required	A name is required for any non-unique nodes
referenceDensity	real64	1000	Reference fluid density
referencePressure	real64	100000	Reference pressure
referenceProppantDensity	real64	1400	Reference proppant density
referenceViscosity	real64	0.001	Reference fluid viscosity

9.89 Element: ProppantTransport

Name	Type	De- fault	Description
bridg- ingFac- tor	real64	0	Bridging factor used for bridging/screen-out calculation
cflFac- tor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
critical- Shield- sNum- ber	real64	0	Critical Shields number
dis- cretiza- tion	string	re- quired	Name of discretization object to use for this solver.
fluid- Names	string_array	re- quired	Names of fluid constitutive models for each region.
fric- tionCo- efficient	real64	0.03	Friction coefficient
ini- tialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
input- FluxEs- timate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	in- te- ger	0	Log level
max- Prop- ant- Con- centra- tion	real64	0.6	Maximum proppant concentration
name	string	re- quired	A name is required for any non-unique nodes
perme- ability- Names	string_array	re- quired	Names of permeability constitutive models for each region.
prop- ant- Density	real64	2500	Proppant density
prop- antDi- ameter	real64	0.0004	Proppant diameter
prop- ant- Names	string_array	re- quired	Name of proppant constitutive object to use for this solver.
solid- Names	string_array	re- quired	Names of solid constitutive models for each region.
targe- tRe- gions	string_array	re- quired	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
updateProp- ant- Pack- ing	in- te- ger	0	Flag that enables/disables proppant-packing update

9.90 Element: Python

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.

9.91 Element: Restart

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.

9.92 Element: Run

Name	Type	De- fault	Description
args	string		Any extra command line arguments to pass to GEOSX.
autoParti- tion	string		May be 'Off' or 'On', if 'On' partitioning arguments are created automati- cally. Default is Off.
name	string	re- quired	The name of this benchmark.
nodes	integer	re- quired	The number of nodes needed to run the benchmark.
strongScal- ing	inte- ger_array	{0}	Repeat the benchmark N times, scaling the number of nodes in the bench- mark by these values.
tasksPerN- ode	integer	re- quired	The number of tasks per node to run the benchmark with.
threadsPer- Task	integer	0	The number of threads per task to run the benchmark with.
timeLimit	integer	0	The time limit of the benchmark.

9.93 Element: Silo

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.
plotFileRoot	string	plot	(no description available)
plotLevel	integer	1	(no description available)
writeCellElementMesh	integer	1	(no description available)
writeEdgeMesh	integer	0	(no description available)
writeFEMFaces	integer	0	(no description available)
writeFaceElementMesh	integer	1	(no description available)

9.94 Element: SinglePhaseFVM

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object to use for this solver.
fluid-Names	string_array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
input-FluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
permeability-Names	string_array	required	Names of permeability constitutive models for each region.
solid-Names	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

9.95 Element: SinglePhaseHybridFVM

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object to use for this solver.
fluid-Names	string_array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
input-FluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
permeability-Names	string_array	required	Names of permeability constitutive models for each region.
solid-Names	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.96 Element: SinglePhasePoromechanics

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object (defined in the Numerical Methods) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
fluid-Solver-Name	string	required	Name of the fluid mechanics solver to use in the poromechanics solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
porous-Material-Names	string_array	required	The name of the material that should be used in the constitutive updates
solid-Solver-Name	string	required	Name of the solid mechanics solver to use in the poromechanics solver
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

9.97 Element: SinglePhasePoromechanicsEmbeddedFractures

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fluid-Solver-Name	string	required	Name of the fluid mechanics solver to use in the poromechanics solver
fractures-Solver-Name	string	required	Name of the fractures solver to use in the fractured poroelastic solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
porous-Material-Names	string_array	required	The name of the material that should be used in the constitutive updates
solid-Solver-Name	string	required	Name of the solid mechanics solver to use in the poromechanics solver
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

9.98 Element: SinglePhaseProppantFVM

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
discretization	string	required	Name of discretization object to use for this solver.
fluid-Names	string_array	required	Names of fluid constitutive models for each region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
input-FluxEstimate	real64	1	Initial estimate of the input flux used only for residual scaling. This should be essentially equivalent to the input flux * dt.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
permeability-Names	string_array	required	Names of permeability constitutive models for each region.
solid-Names	string_array	required	Names of solid constitutive models for each region.
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	Element: LinearSolverParameters
Non-linear-Solver-Parameters	node	unique	Element: NonlinearSolverParameters

9.99 Element: SinglePhaseReservoir

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
flow-Solver-Name	string	required	Name of the flow solver to use in the reservoir-well system solver
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
well-Solver-Name	string	required	Name of the well solver to use in the reservoir-well system solver
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.100 Element: SinglePhaseWell

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fluid-Names	string_array	required	Name of fluid constitutive object to use for this solver.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	required	A name is required for any non-unique nodes
targetRegions	string_array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>
Well-Controls	node		<i>Element: WellControls</i>

9.101 Element: SolidMechanicsEmbeddedFractures

Name	Type	De- fault	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRelation-Name	string	re- quired	Name of contact relation to enforce constraints on fracture boundary.
fractureRegion-Name	string	re- quired	Name of the fracture region.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
name	string	re- quired	A name is required for any non-unique nodes
solid-Solver-Name	string	re- quired	Name of the solid mechanics solver in the rock matrix
targetRegions	string array	re- quired	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
Linear-Solver-Parameters	node	unique	<i>Element: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.102 Element: SolidMechanicsLagrangianSSLE

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRelationName	string	NOCONTACT	Name of contact relation to enforce constraints on fracture boundary.
discretization	string	required	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
massDamping	real64	0	Value of mass based damping coefficient.
maxNumResolves	integer	10	Value to indicate how many resolves may be executed after some other event is executed. For example, if a SurfaceGenerator is specified, it will be executed after the mechanics solve. However if a new surface is generated, then the mechanics solve must be executed again due to the change in topology.
name	string	required	A name is required for any non-unique nodes
newmarkBeta	real64	0.25	Value of β in the Newmark Method for Implicit Dynamic time integration option. This should be $\text{pow}(\text{newmarkGamma}+0.5, 2.0)/4.0$ unless you know what you are doing.
newmarkGamma	real64	0.5	Value of γ in the Newmark Method for Implicit Dynamic time integration option
402			Chapter 9: Datastructure index
solidMaterialNames	string_array	required	The name of the material that should be used in the constitutive updates

9.103 Element: SolidMechanics_LagrangianFEM

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
contactRelationName	string	NOCONTACT	Name of contact relation to enforce constraints on fracture boundary.
discretization	string	required	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
massDamping	real64	0	Value of mass based damping coefficient.
maxNumResolves	integer	10	Value to indicate how many resolves may be executed after some other event is executed. For example, if a SurfaceGenerator is specified, it will be executed after the mechanics solve. However if a new surface is generated, then the mechanics solve must be executed again due to the change in topology.
name	string	required	A name is required for any non-unique nodes
newmarkBeta	real64	0.25	Value of β in the Newmark Method for Implicit Dynamic time integration option. This should be $\text{pow}(\text{newmarkGamma}+0.5, 2.0)/4.0$ unless you know what you are doing.
newmarkGamma	real64	0.5	Value of γ in the Newmark Method for Implicit Dynamic time integration option
404			Chapter 9: Datastructure index
solidMaterialNames	string_array	required	The name of the material that should be used in the constitutive updates

9.104 Element: SoloEvent

Name	Type	De- fault	Description
beginTime	real64	0	Start time of this event.
endTime	real64	1e+100	End time of this event.
finalDt- Stretch	real64	0.001	Allow the final dt request for this event to grow by this percentage to match the end-Time exactly.
forceDt	real64	-1	While active, this event will request this timestep value (ignoring any children/targets requests).
logLevel	in- te- ger	0	Log level
max- EventDt	real64	-1	While active, this event will request a timestep \leq this value (depending upon any child/target requests).
name	string	re- quired	A name is required for any non-unique nodes
target	string		Name of the object to be executed when the event criteria are met.
targetCy- cle	in- te- ger	-1	Targeted cycle to execute the event.
targetEx- actStart- Stop	in- te- ger	1	If this option is set, the event will reduce its timestep requests to match any specified beginTime/endTimes exactly.
targe- tExact- Timestep	in- te- ger	1	If this option is set, the event will reduce its timestep requests to match the specified execution time exactly: $dt_request = \min(dt_request, t_target - time)$.
targetTime	real64	-1	Targeted time to execute the event.
HaltEvent	node		<i>Element: HaltEvent</i>
Peri- odicEvent	node		<i>Element: PeriodicEvent</i>
SoloEvent	node		<i>Element: SoloEvent</i>

9.105 Element: Solvers

Name	Type	Default	Description
gravityVector	R1Tensor	{0,0,-9.81}	Gravity vector used in the physics solvers
AcousticSEM	node		<i>Element: AcousticSEM</i>
CompositionalMultiphaseFVM	node		<i>Element: CompositionalMultiphaseFVM</i>
CompositionalMultiphaseHybridFVM	node		<i>Element: CompositionalMultiphaseHybrid-FVM</i>
CompositionalMultiphaseReservoir	node		<i>Element: CompositionalMultiphaseReservoir</i>
CompositionalMultiphaseWell	node		<i>Element: CompositionalMultiphaseWell</i>
EmbeddedSurfaceGenerator	node		<i>Element: EmbeddedSurfaceGenerator</i>
FlowProppantTransport	node		<i>Element: FlowProppantTransport</i>
Hydrofracture	node		<i>Element: Hydrofracture</i>
LagrangianContact	node		<i>Element: LagrangianContact</i>
LaplaceFEM	node		<i>Element: LaplaceFEM</i>
LaplaceVEM	node		<i>Element: LaplaceVEM</i>
MultiphasePoromechanics	node		<i>Element: MultiphasePoromechanics</i>
PhaseFieldDamageFEM	node		<i>Element: PhaseFieldDamageFEM</i>
PhaseFieldFracture	node		<i>Element: PhaseFieldFracture</i>
ProppantTransport	node		<i>Element: ProppantTransport</i>
SinglePhaseFVM	node		<i>Element: SinglePhaseFVM</i>
SinglePhaseHybridFVM	node		<i>Element: SinglePhaseHybridFVM</i>
SinglePhasePoromechanics	node		<i>Element: SinglePhasePoromechanics</i>
SinglePhasePoromechanicsEmbedded-Fractures	node		<i>Element: SinglePhasePoromechanicsEmbeddedFractures</i>
SinglePhaseProppantFVM	node		<i>Element: SinglePhaseProppantFVM</i>
SinglePhaseReservoir	node		<i>Element: SinglePhaseReservoir</i>
SinglePhaseWell	node		<i>Element: SinglePhaseWell</i>
SolidMechanicsEmbeddedFractures	node		<i>Element: SolidMechanicsEmbeddedFractures</i>
SolidMechanicsLagrangianSSLE	node		<i>Element: SolidMechanicsLagrangianSSLE</i>
SolidMechanics_LagrangianFEM	node		<i>Element: SolidMechanics_LagrangianFEM</i>
SurfaceGenerator	node		<i>Element: SurfaceGenerator</i>

9.106 Element: SourceFlux

Name	Type	Default	Description
bcApplicationTable-Name	string		Name of table that specifies the on/off application of the bc.
beginTime	real64	-1e+99	time at which BC will start being applied.
component	integer	-1	Component of field (if tensor) to apply boundary condition to
direction	R1Tensor	{0,0,0}	Direction to apply boundary condition to
endTime	real64	1e+99	time at which bc will stop being applied
fieldName	string		Name of field that boundary condition is applied to.
functionName	string		Name of function that specifies variation of the BC
initialCondition	integer	0	BC is applied as an initial condition.
name	string	required	A name is required for any non-unique nodes
objectPath	string		Path to the target field
scale	real64	0	Scale factor for value of BC.
setNames	string_array	required	Name of sets that boundary condition is applied to.

9.107 Element: StrainDependentPermeability

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes

9.108 Element: SurfaceElementRegion

Name	Type	Default	Description
defaultAperture	real64	required	The default aperture of newly formed surface elements.
materialList	string_array	required	List of materials present in this region
name	string	required	A name is required for any non-unique nodes
subRegionType	geosx_SurfaceElementRegionSubRegionType	faceElement	Type of surface element subregion. Valid options: * faceElement * embeddedElement

9.109 Element: SurfaceGenerator

Name	Type	Default	Description
cflFactor	real64	0.5	Factor to apply to the CFL condition when calculating the maximum allowable time step. Values should be in the interval (0,1]
fractureRegion	string	Fracture	(no description available)
initialDt	real64	1e+99	Initial time-step value required by the solver to the event manager.
logLevel	integer	0	Log level
mpiCommOrder	integer	0	Flag to enable MPI consistent communication ordering
name	string	required	A name is required for any non-unique nodes
nodeBasedSIF	integer	0	Rock toughness of the solid material
rockToughness	real64	required	Rock toughness of the solid material
solidMaterialNames	string array	required	Name of the solid material used in solid mechanic solver
targetRegions	string array	required	Allowable regions that the solver may be applied to. Note that this does not indicate that the solver will be applied to these regions, only that allocation will occur such that the solver may be applied to these regions. The decision about what regions this solver will be applied to rests in the EventManager.
LinearSolverParameters	node	unique	<i>Element: LinearSolverParameters</i>
NonlinearSolverParameters	node	unique	<i>Element: NonlinearSolverParameters</i>

9.110 Element: SymbolicFunction

Name	Type	De- fault	Description
expression	string	re- quired	Symbolic math expression
inputVar- Names	string_array	{ }	Name of fields are input to function.
name	string	re- quired	A name is required for any non-unique nodes
variable- Names	string_array	re- quired	List of variables in expression. The order must match the evaluate argument

9.111 Element: TableFunction

Name	Type	Default	Description
coordinateFiles	path_array	{ }	List of coordinate file names for ND Table
coordinates	real64_array	{0}	Coordinates inputs for 1D tables
inputVarNames	string_array	{ }	Name of fields are input to function.
interpolation	geosx_TableFunction_InterpolationType		Interpolation method. Valid options: * linear * nearest * upper * lower
name	string	required	A name is required for any non-unique nodes
values	real64_array	{0}	Values for 1D tables
voxelFile	path		Voxel file name for ND Table

9.112 Element: TableRelativePermeability

Name	Type	Default	Description
gasOilRelPermTableNames	string_array	{ }	List of relative permeability tables for the pair (gas phase, oil phase) The expected format is “{ gasPermTableName, oilPermTableName }”, in that order
name	string	required	A name is required for any non-unique nodes
phaseNames	string_array	required	List of fluid phases
waterOilRelPermTableNames	string_array	{ }	List of relative permeability tables for the pair (water phase, oil phase) The expected format is “{ waterPermTableName, oilPermTableName }”, in that order

9.113 Element: Tasks

Name	Type	Default	Description
PackCollection	node		<i>Element: PackCollection</i>
TriaxialDriver	node		<i>Element: TriaxialDriver</i>

9.114 Element: ThickPlane

Name	Type	Default	Description
name	string	required	A name is required for any non-unique nodes
normal	R1Tensor	required	Normal (n_x, n_y, n_z) to the plane (will be normalized automatically)
origin	R1Tensor	required	Origin point (x, y, z) of the plane (basically, any point on the plane)
thickness	real64	required	The total thickness of the plane (with half to each side)

9.115 Element: TimeHistory

Name	Type	Default	Description
childDirectory	string		Child directory path
filename	string	TimeHistory	The filename to which to write time history output.
format	string	hdf	The output file format for time history output.
name	string	required	A name is required for any non-unique nodes
parallelThreads	integer	1	Number of plot files.
sources	string_array	required	A list of collectors from which to collect and output time history information.

9.116 Element: Traction

Name	Type	Default	Description
bcApplicationTableName	string		Name of table that specifies the on/off application of the bc.
beginTime	real64	-1e+99	time at which BC will start being applied.
direction	R1Tensor	{0,0,0}	Direction to apply boundary condition to
endTime	real64	1e+99	time at which bc will stop being applied
functionName	string		Name of function that specifies variation of the BC
initialCondition	integer	0	BC is applied as an initial condition.
inputStress	R2SymTensor	{0,0,0,0,0,0}	Input stress for traction-Type = stress
name	string	required	A name is required for any non-unique nodes
objectPath	string		Path to the target field
scale	real64	0	Scale factor for value of BC.
setNames	string_array	required	Name of sets that boundary condition is applied to.
tractionType	geosx_TractionBoundaryCondition	vector_TractionType	Type of traction boundary condition. Options are: vector - traction is applied to the faces as specified from the scale and direction, normal - traction is applied to the faces as a pressure specified from the product of scale and the outward face normal, stress - traction is applied to the faces as specified by the inner product of input stress and face normal.

9.117 Element: TriaxialDriver

Name	Type	Default	Description
baseline	path	none	Baseline file
logLevel	integer	0	Log level
material	string	required	Solid material to test
mode	string	required	Test mode [triaxial, volumetric, oedometer]
name	string	required	A name is required for any non-unique nodes
output	string	none	Output file
steps	integer	required	Number of load steps to take
strainFunction	string	required	Function controlling strain loading (role depends on test mode)
stressFunction	string	required	Function controlling stress loading (role depends on test mode)

9.118 Element: TwoPointFluxApproximation

Name	Type	De- fault	Description
areaRelTol	real64	1e-08	Relative tolerance for area calculations.
coefficientModel-Names	string_array	{ }	List of constitutive models that contain the coefficient used to build the stencil
coefficientName	string	re- quired	Name of coefficient field
fieldName	string	re- quired	Name of primary solution field
meanPermCoeffi- cient	real64	1	(no description available)
name	string	re- quired	A name is required for any non-unique nodes
targetRegions	string_array	{ }	List of regions to build the stencil for

9.119 Element: VTK

Name	Type	Default	Description
childDirectory	string		Child directory path
name	string	re- quired	A name is required for any non-unique nodes
parallelThreads	inte- ger	1	Number of plot files.
plotFileRoot	string	VTK	Name of the root file for this output.
plotLevel	inte- ger	1	Level detail plot. Only fields with lower of equal plot level will be output.
writeBinary- Data	inte- ger	1	Output the data in binary format
writeFEMFaces	inte- ger	0	(no description available)

9.120 Element: VanGenuchtenBakerRelativePermeability

Name	Type	Default	Description
gasOilRelPermExponentInv	real64_array	{0.5}	Rel perm power law exponent inverse for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasExp, oilExp }”, in that order
gasOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (gas phase, oil phase) at residual water saturation The expected format is “{ gasMax, oilMax }”, in that order
name	string	required	A name is required for any non-unique nodes
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	required	List of fluid phases
waterOilRelPermExponentInv	real64_array	{0.5}	Rel perm power law exponent inverse for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterExp, oilExp }”, in that order
waterOilRelPermMaxValue	real64_array	{0}	Maximum rel perm value for the pair (water phase, oil phase) at residual gas saturation The expected format is “{ waterMax, oilMax }”, in that order

9.121 Element: VanGenuchtenCapillaryPressure

Name	Type	De- fault	Description
capPressureEpsilon	real64	1e-06	Saturation at which the extremum capillary pressure is attained; used to avoid infinite capillary pressure values for saturations close to 0 and 1
name	string	re- quired	A name is required for any non-unique nodes
phaseCapPressureExponentInv	real64_array	{0.5}	Inverse of capillary power law exponent for each phase
phaseCapPressureMultiplier	real64_array	{1}	Entry pressure value for each phase
phaseMinVolumeFraction	real64_array	{0}	Minimum volume fraction value for each phase
phaseNames	string_array	re- quired	List of fluid phases

9.122 Element: WellControls

Name	Type	Default	Description
control	geosx_WellControls_Control	required	Well control. Valid options: * BHP * phaseVolRate * totalVolRate
injectionStream	real64_array	{-1}	Global component densities for the injection stream
name	string	required	A name is required for any non-unique nodes
referenceElevation	real64	required	Reference elevation where BHP control is enforced
surfacePressure	real64	0	Surface pressure used to compute volumetric rates when surface conditions are used
surfaceTemperature	real64	0	Surface temperature used to compute volumetric rates when surface conditions are used
targetBHP	real64	-1	Target bottom-hole pressure
targetBHPTablename	string		Name of the BHP table when the rate is a time dependent function
targetPhaseName	string		Name of the target phase
targetPhaseRate	real64	0	Target phase volumetric rate
targetPhaseRateTablename	string		Name of the phase rate table when the rate is a time dependent function
targetTotalRate	real64	0	Target total volumetric rate
targetTotalRateTablename	string		Name of the total rate table when the rate is a time dependent function
type	geosx_WellControls_Type	required	Well type. Valid options: * producer * injector
useSurfaceConditions	integer	0	Flag to specify whether rates are checked at surface or reservoir conditions.
9.122. Element: WellControls			Equal to 1 for surface conditions, and to 0 for reservoir conditions 417

9.123 Element: WellElementRegion

Name	Type	Default	Description
materialList	string_array	required	List of materials present in this region
name	string	required	A name is required for any non-unique nodes

9.124 Element: lassen

Name	Type	Default	Description
Run	node	unique	<i>Element: Run</i>

9.125 Element: quartz

Name	Type	Default	Description
Run	node	unique	<i>Element: Run</i>

9.125.1 Datastructure Definitions

9.126 Datastructure: AcousticSEM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
pressureNp1AtReceivers	real64_array		Pressure value at each receiver for each timestep
receiverIsLocal	localIndex_array		Flag that indicates whether the receiver is local to this MPI rank
receiverNodeIds	localIndex_array2d		Indices of the nodes (in the right order) for each receiver point
sourceConstants	real64_array2d		Constant part of the receiver for the nodes listed in m_receiverNodeIds
sourceIsLocal	localIndex_array		Flag that indicates whether the source is local to this MPI rank
sourceNodeIds	localIndex_array2d		Indices of the nodes (in the right order) for each source point
dampingVector	real64_array	<i>Datastructure: nodeManager</i>	Diagonal of the Damping Matrix.
freeSurfaceFaceIndicator	localIndex_array	<i>Datastructure: FaceManager</i>	Free surface indicator, 1 if a face is on free surface 0 otherwise.
freeSurfaceNodeIndicator	localIndex_array	<i>Datastructure: nodeManager</i>	Free surface indicator, 1 if a node is on free surface 0 otherwise.
massVector	real64_array	<i>Datastructure: nodeManager</i>	Diagonal of the Mass Matrix.
pressure_n	real64_array	<i>Datastructure: nodeManager</i>	Scalar pressure at time n.
pressure_nm1	real64_array	<i>Datastructure: nodeManager</i>	Scalar pressure at time n-1.
pressure_np1	real64_array	<i>Datastructure: nodeManager</i>	Scalar pressure at time n+1.
rhs	real64_array	<i>Datastructure: nodeManager</i>	RHS
stiffnessVector	real64_array	<i>Datastructure: nodeManager</i>	Stiffness vector contains $R_h * \text{Pressure}_n$.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.127 Datastructure: Benchmarks

Name	Type	Description
lassen	node	<i>Datastructure: lassen</i>
quartz	node	<i>Datastructure: quartz</i>

9.128 Datastructure: BiotPorosity

Name	Type	Description
biotCoefficient	real64_array	Biot coefficient.
dPorosity_dPressure	real64_array2d	(no description available)
oldPorosity	real64_array2d	(no description available)
porosity	real64_array2d	(no description available)
referencePorosity	real64_array	(no description available)

9.129 Datastructure: BlackOilFluid

Name	Type
PVTO	geosx_constitutive_PVTOData
dPhaseCompFraction_dGlobalCompFraction	LvArray_Array< double, 5, camp_int_seq< long, 0l, 1l, 2l, 3l, 4l >, long, LvArray_CL
dPhaseCompFraction_dPressure	real64_array4d
dPhaseCompFraction_dTemperature	real64_array4d
dPhaseDensity_dGlobalCompFraction	real64_array4d
dPhaseDensity_dPressure	real64_array3d
dPhaseDensity_dTemperature	real64_array3d
dPhaseFraction_dGlobalCompFraction	real64_array4d
dPhaseFraction_dPressure	real64_array3d
dPhaseFraction_dTemperature	real64_array3d
dPhaseMassDensity_dGlobalCompFraction	real64_array4d
dPhaseMassDensity_dPressure	real64_array3d
dPhaseMassDensity_dTemperature	real64_array3d
dPhaseViscosity_dGlobalCompFraction	real64_array4d
dPhaseViscosity_dPressure	real64_array3d
dPhaseViscosity_dTemperature	real64_array3d
dTotalDensity_dGlobalCompFraction	real64_array3d
dTotalDensity_dPressure	real64_array2d
dTotalDensity_dTemperature	real64_array2d
formationVolFactorTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >,
hydrocarbonPhaseOrder	integer_array
phaseCompFraction	real64_array4d
phaseDensity	real64_array3d
phaseFraction	real64_array3d
phaseMassDensity	real64_array3d
phaseOrder	integer_array
phaseTypes	integer_array
phaseViscosity	real64_array3d
totalDensity	real64_array2d
useMass	integer
viscosityTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >,

9.130 Datastructure: Blueprint

Name	Type	Description

9.131 Datastructure: BoundedPlane

Name	Type	Description

9.132 Datastructure: Box

Name	Type	Description
center	R1Tensor	(no description available)
cosStrike	real64	(no description available)
sinStrike	real64	(no description available)

9.133 Datastructure: BrooksCoreyBakerRelativePermeability

Name	Type	Description
dPhaseRelPerm_dPhaseVolFraction	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseRelPerm	real64_array	(no description available)
phaseTypes	integer_array	(no description available)
volFracScale	real64	Factor used to scale the phase capillary pressure, defined as: one minus the sum of the phase minimum volume fractions.

9.134 Datastructure: BrooksCoreyCapillaryPressure

Name	Type	Description
dPhaseCapPressure_dPhaseVolFraction	real64_array	(no description available)
phaseCapPressure	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseTypes	integer_array	(no description available)
volFracScale	real64	Factor used to scale the phase capillary pressure, defined as: one minus the sum of the phase minimum volume fractions.

9.135 Datastructure: BrooksCoreyRelativePermeability

Name	Type	Description
dPhaseRelPerm_dPhaseVolFrac	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseRelPerm	real64_array	(no description available)
phaseTypes	integer_array	(no description available)
volFracScale	real64	Factor used to scale the phase capillary pressure, defined as: one minus the sum of the phase minimum volume fractions.

9.136 Datastructure: CO2BrineFluid

Name	Type	Description
dPhaseCompFrac- tion_dGlobalCompFraction	LvArray_Array< double, 5, camp_int_seq< long, 0l, 1l, 2l, 3l, 4l >, long, LvArray_ChaiBuffer >	(no description available)
dPhaseCompFrac- tion_dPressure	real64_array4d	(no description available)
dPhaseCompFrac- tion_dTemperature	real64_array4d	(no description available)
dPhaseDen- sity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseDensity_dPressure	real64_array3d	(no description available)
dPhaseDen- sity_dTemperature	real64_array3d	(no description available)
dPhaseFrac- tion_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseFraction_dPressure	real64_array3d	(no description available)
dPhaseFrac- tion_dTemperature	real64_array3d	(no description available)
dPhaseMassDen- sity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseMassDen- sity_dPressure	real64_array3d	(no description available)
dPhaseMassDen- sity_dTemperature	real64_array3d	(no description available)
dPhaseViscos- ity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseViscosity_dPressure	real64_array3d	(no description available)
dPhaseViscos- ity_dTemperature	real64_array3d	(no description available)
dTotalDen- sity_dGlobalCompFraction	real64_array3d	(no description available)
dTotalDensity_dPressure	real64_array2d	(no description available)
dTotalDensity_dTemperature	real64_array2d	(no description available)
phaseCompFraction	real64_array4d	(no description available)
phaseDensity	real64_array3d	(no description available)
phaseFraction	real64_array3d	(no description available)
phaseMassDensity	real64_array3d	(no description available)
phaseViscosity	real64_array3d	(no description available)
totalDensity	real64_array2d	(no description available)
useMass	integer	(no description available)

9.137 Datastructure: CarmanKozenyPermeability

Name	Type	Description
dPerm_dPorosity	real64_array3d	(no description available)
dPerm_dPressure	real64_array3d	dPerm_dPressure of the rock.
permeability	real64_array3d	permeability of the rock.

9.138 Datastructure: CellElementRegion

Name	Type	Description
domainBoundaryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlobalMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
elementSubRegions	node	<i>Datastructure: elementSubRegions</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.139 Datastructure: ChomboLO

Name	Type	Description

9.140 Datastructure: CompositeFunction

Name	Type	Description

9.141 Datastructure: CompositionalMultiphaseFVM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
facePressure	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the pressures at the faces.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.142 Datastructure: CompositionalMultiphaseFluid

Name	Type	Description
dPhaseCompFrac- tion_dGlobalCompFraction	LvArray_Array< double, 5, camp_int_seq< long, 0l, 1l, 2l, 3l, 4l >, long, LvArray_ChaiBuffer >	(no description available)
dPhaseCompFrac- tion_dPressure	real64_array4d	(no description available)
dPhaseCompFrac- tion_dTemperature	real64_array4d	(no description available)
dPhaseDen- sity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseDensity_dPressure	real64_array3d	(no description available)
dPhaseDen- sity_dTemperature	real64_array3d	(no description available)
dPhaseFrac- tion_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseFraction_dPressure	real64_array3d	(no description available)
dPhaseFrac- tion_dTemperature	real64_array3d	(no description available)
dPhaseMassDen- sity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseMassDen- sity_dPressure	real64_array3d	(no description available)
dPhaseMassDen- sity_dTemperature	real64_array3d	(no description available)
dPhaseViscos- ity_dGlobalCompFraction	real64_array4d	(no description available)
dPhaseViscosity_dPressure	real64_array3d	(no description available)
dPhaseViscos- ity_dTemperature	real64_array3d	(no description available)
dTotalDen- sity_dGlobalCompFraction	real64_array3d	(no description available)
dTotalDensity_dPressure	real64_array2d	(no description available)
dTotalDensity_dTemperature	real64_array2d	(no description available)
phaseCompFraction	real64_array4d	(no description available)
phaseDensity	real64_array3d	(no description available)
phaseFraction	real64_array3d	(no description available)
phaseMassDensity	real64_array3d	(no description available)
phaseViscosity	real64_array3d	(no description available)
totalDensity	real64_array2d	(no description available)
useMass	integer	(no description available)

9.143 Datastructure: CompositionalMultiphaseHybridFVM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
facePressure	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the pressures at the faces.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.144 Datastructure: CompositionalMultiphaseReservoir

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.145 Datastructure: CompositionalMultiphaseWell

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>
Well-Controls	node	<i>Datastructure: WellControls</i>

9.146 Datastructure: CompressibleSinglePhaseFluid

Name	Type	Description
dDensity_dPressure	real64_array2d	(no description available)
dViscosity_dPressure	real64_array2d	(no description available)
density	real64_array2d	(no description available)
viscosity	real64_array2d	(no description available)

9.147 Datastructure: CompressibleSolidCarmanKozenyPermeability

Name	Type	Description

9.148 Datastructure: CompressibleSolidConstantPermeability

Name	Type	Description

9.149 Datastructure: CompressibleSolidParallelPlatesPermeability

Name	Type	Description

9.150 Datastructure: ConstantPermeability

Name	Type	Description
dPerm_dPressure	real64_array3d	dPerm_dPressure of the rock.
permeability	real64_array3d	permeability of the rock.

9.151 Datastructure: Constitutive

Name	Type	Description
BiotPorosity	node	<i>Datastructure: BiotPorosity</i>
BlackOilFluid	node	<i>Datastructure: BlackOilFluid</i>
BrooksCoreyBakerRelativePermeability	node	<i>Datastructure: BrooksCoreyBakerRelativePermeability</i>
BrooksCoreyCapillaryPressure	node	<i>Datastructure: BrooksCoreyCapillaryPressure</i>
BrooksCoreyRelativePermeability	node	<i>Datastructure: BrooksCoreyRelativePermeability</i>
CO2BrineFluid	node	<i>Datastructure: CO2BrineFluid</i>

Continued on next page

Table 9.3 – continued from previous page

Name	Type	Description
CarmanKozenyPermeability	node	<i>Datastructure: CarmanKozenyPermeability</i>
CompositionalMultiphaseFluid	node	<i>Datastructure: CompositionalMultiphaseFluid</i>
CompressibleSinglePhaseFluid	node	<i>Datastructure: CompressibleSinglePhaseFluid</i>
CompressibleSolidCarmanKozenyPermeability	node	<i>Datastructure: CompressibleSolidCarmanKozenyPermeability</i>
CompressibleSolidConstantPermeability	node	<i>Datastructure: CompressibleSolidConstantPermeability</i>
CompressibleSolidParallelPlatesPermeability	node	<i>Datastructure: CompressibleSolidParallelPlatesPermeability</i>
ConstantPermeability	node	<i>Datastructure: ConstantPermeability</i>
Contact	node	<i>Datastructure: Contact</i>
Coulomb	node	<i>Datastructure: Coulomb</i>
DamageElasticIsotropic	node	<i>Datastructure: DamageElasticIsotropic</i>
DamageSpectralElasticIsotropic	node	<i>Datastructure: DamageSpectralElasticIsotropic</i>
DamageVolDevElasticIsotropic	node	<i>Datastructure: DamageVolDevElasticIsotropic</i>
DeadOilFluid	node	<i>Datastructure: DeadOilFluid</i>
DruckerPrager	node	<i>Datastructure: DruckerPrager</i>
ElasticIsotropic	node	<i>Datastructure: ElasticIsotropic</i>
ElasticOrthotropic	node	<i>Datastructure: ElasticOrthotropic</i>
ElasticTransverseIsotropic	node	<i>Datastructure: ElasticTransverseIsotropic</i>
ExtendedDruckerPrager	node	<i>Datastructure: ExtendedDruckerPrager</i>
NullModel	node	<i>Datastructure: NullModel</i>
ParallelPlatesPermeability	node	<i>Datastructure: ParallelPlatesPermeability</i>
ParticleFluid	node	<i>Datastructure: ParticleFluid</i>
PermeabilityBase	node	<i>Datastructure: PermeabilityBase</i>
PorousDruckerPrager	node	<i>Datastructure: PorousDruckerPrager</i>
PorousElasticIsotropic	node	<i>Datastructure: PorousElasticIsotropic</i>
PorousElasticOrthotropic	node	<i>Datastructure: PorousElasticOrthotropic</i>
PorousElasticTransverseIsotropic	node	<i>Datastructure: PorousElasticTransverseIsotropic</i>
PorousExtendedDruckerPrager	node	<i>Datastructure: PorousExtendedDruckerPrager</i>
PressurePorosity	node	<i>Datastructure: PressurePorosity</i>
ProppantSlurryFluid	node	<i>Datastructure: ProppantSlurryFluid</i>
StrainDependentPermeability	node	<i>Datastructure: StrainDependentPermeability</i>
TableRelativePermeability	node	<i>Datastructure: TableRelativePermeability</i>
VanGenuchtenBakerRelativePermeability	node	<i>Datastructure: VanGenuchtenBakerRelativePermeability</i>
VanGenuchtenCapillaryPressure	node	<i>Datastructure: VanGenuchtenCapillaryPressure</i>

9.152 Datastructure: ConstitutiveModels

Name	Type	Description

9.153 Datastructure: Contact

Name	Type	Description
TableFunction	node	<i>Datastructure: TableFunction</i>

9.154 Datastructure: Coulomb

Name	Type	Description
TableFunction	node	<i>Datastructure: TableFunction</i>

9.155 Datastructure: Cylinder

Name	Type	Description

9.156 Datastructure: DamageElasticIsotropic

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
damage	real64_array2d	Material Damage Variable
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
shearModulus	real64_array	Elastic Shear Modulus Field
strainEnergyDensity	real64_array2d	Strain Energy Density
stress	real64_array3d	Current Material Stress

9.157 Datastructure: DamageSpectralElasticIsotropic

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
damage	real64_array2d	Material Damage Variable
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
shearModulus	real64_array	Elastic Shear Modulus Field
strainEnergyDensity	real64_array2d	Strain Energy Density
stress	real64_array3d	Current Material Stress

9.158 Datastructure: DamageVolDevElasticIsotropic

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
damage	real64_array2d	Material Damage Variable
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
shearModulus	real64_array	Elastic Shear Modulus Field
strainEnergyDensity	real64_array2d	Strain Energy Density
stress	real64_array3d	Current Material Stress

9.159 Datastructure: DeadOilFluid

Name	Type
dPhaseCompFraction_dGlobalCompFraction	LvArray_Array< double, 5, camp_int_seq< long, 0l, 1l, 2l, 3l, 4l >, long, LvArray_Cl
dPhaseCompFraction_dPressure	real64_array4d
dPhaseCompFraction_dTemperature	real64_array4d
dPhaseDensity_dGlobalCompFraction	real64_array4d
dPhaseDensity_dPressure	real64_array3d
dPhaseDensity_dTemperature	real64_array3d
dPhaseFraction_dGlobalCompFraction	real64_array4d
dPhaseFraction_dPressure	real64_array3d
dPhaseFraction_dTemperature	real64_array3d
dPhaseMassDensity_dGlobalCompFraction	real64_array4d
dPhaseMassDensity_dPressure	real64_array3d
dPhaseMassDensity_dTemperature	real64_array3d
dPhaseViscosity_dGlobalCompFraction	real64_array4d
dPhaseViscosity_dPressure	real64_array3d
dPhaseViscosity_dTemperature	real64_array3d
dTotalDensity_dGlobalCompFraction	real64_array3d
dTotalDensity_dPressure	real64_array2d
dTotalDensity_dTemperature	real64_array2d
formationVolFactorTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >,
hydrocarbonPhaseOrder	integer_array
phaseCompFraction	real64_array4d
phaseDensity	real64_array3d
phaseFraction	real64_array3d
phaseMassDensity	real64_array3d
phaseOrder	integer_array
phaseTypes	integer_array
phaseViscosity	real64_array3d
totalDensity	real64_array2d
useMass	integer
viscosityTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >,

9.160 Datastructure: Dirichlet

Name	Type	Description

9.161 Datastructure: DruckerPrager

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
cohesion	real64_array2d	New cohesion state
density	real64_array2d	Material Density
dilation	real64_array	Plastic potential slope
friction	real64_array	Yield surface slope
hardening	real64_array	Hardening rate
oldCohesion	real64_array2d	Old cohesion state
oldStress	real64_array3d	Previous Material Stress
shearModulus	real64_array	Elastic Shear Modulus Field
stress	real64_array3d	Current Material Stress

9.162 Datastructure: ElasticIsotropic

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
shearModulus	real64_array	Elastic Shear Modulus Field
stress	real64_array3d	Current Material Stress

9.163 Datastructure: ElasticOrthotropic

Name	Type	Description
c11	real64_array	Elastic Stiffness Field C11
c12	real64_array	Elastic Stiffness Field C12
c13	real64_array	Elastic Stiffness Field C13
c22	real64_array	Elastic Stiffness Field C22
c23	real64_array	Elastic Stiffness Field C23
c33	real64_array	Elastic Stiffness Field C33
c44	real64_array	Elastic Stiffness Field C44
c55	real64_array	Elastic Stiffness Field C55
c66	real64_array	Elastic Stiffness Field C66
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
stress	real64_array3d	Current Material Stress

9.164 Datastructure: ElasticTransverselsotropic

Name	Type	Description
c11	real64_array	Elastic Stiffness Field C11
c13	real64_array	Elastic Stiffness Field C13
c33	real64_array	Elastic Stiffness Field C33
c44	real64_array	Elastic Stiffness Field C44
c66	real64_array	Elastic Stiffness Field C66
density	real64_array2d	Material Density
oldStress	real64_array3d	Previous Material Stress
stress	real64_array3d	Current Material Stress

9.165 Datastructure: ElementRegions

Name	Type	Description
domainBoundaryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlobalMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
CellElementRegion	node	<i>Datastructure: CellElementRegion</i>
SurfaceElementRegion	node	<i>Datastructure: SurfaceElementRegion</i>
WellElementRegion	node	<i>Datastructure: WellElementRegion</i>
elementRegionsGroup	node	<i>Datastructure: elementRegionsGroup</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.166 Datastructure: EmbeddedSurfaceGenerator

Name	Type	Registered On	Description
discretization	string		Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
parent-EdgeIndex	localIndex_array	<i>Datastructure: embeddedSurfacesNodeManager</i>	Index of parent edge within the mesh object it is registered on.
Linear-Solver-Parameters	node		<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.167 Datastructure: Events

Name	Type	Description
currentSubEvent	integer	Index of the current subevent.
cycle	integer	Current simulation cycle number.
dt	real64	Current simulation timestep.
time	real64	Current simulation time.
HaltEvent	node	<i>Datastructure: HaltEvent</i>
PeriodicEvent	node	<i>Datastructure: PeriodicEvent</i>
SoloEvent	node	<i>Datastructure: SoloEvent</i>

9.168 Datastructure: ExtendedDruckerPrager

Name	Type	Description
bulkModulus	real64_array	Elastic Bulk Modulus Field
density	real64_array2d	Material Density
dilationRatio	real64_array	Plastic potential slope ratio
hardening	real64_array	Hardening parameter
initialFriction	real64_array	Initial yield surface slope
oldStateVariable	real64_array2d	Old equivalent plastic shear strain
oldStress	real64_array3d	Previous Material Stress
pressureIntercept	real64_array	Pressure point at cone vertex
residualFriction	real64_array	Residual yield surface slope
shearModulus	real64_array	Elastic Shear Modulus Field
stateVariable	real64_array2d	New equivalent plastic shear strain
stress	real64_array3d	Current Material Stress

9.169 Datastructure: FaceManager

Name	Type	Registered By	Description
deltaFacePressure	real64_array		(no description available)
domainBoundaryIndicator	integer_array		(no description available)
edgeList	geosx_InterObjectRelation<LvAr-ray_ArrayOfArrays<long, long, LvAr-ray_ChaiBuffer > >		(no description available)
elemList	localIndex_array2d		(no description available)
elemRegionList	localIndex_array2d		(no description available)
elemSubRegionList	localIndex_array2d		(no description available)
faceArea	real64_array		(no description available)
faceCenter	real64_array2d		(no description available)
faceNormal	real64_array2d		(no description available)
ghostRank	integer_array		(no description available)
globalToLocalMap	geosx_mapBase<long long, long, std_integral_constant<bool, false > >		(no description available)
gravityCoefficient	real64_array		(no description available)
isExternal	integer_array		(no description available)
localToGlobalMap	globalIndex_array		Array that contains a map from localIndex to globalIndex.
mimGravityCoefficient	real64_array		(no description available)
nodeList	geosx_InterObjectRelation<LvAr-ray_ArrayOfArrays<long, long, LvAr-ray_ChaiBuffer > >		(no description available)

Continued on next page

Table 9.5 – continued from previous page

Name	Type	Registered By	Description
K_IC	real64_array2d	<i>Datastructure: Surface-Generator</i>	Critical Stress Intensity Factor K_{IC} in the plane of the face.
SIFonFace	real64_array	<i>Datastructure: Surface-Generator</i>	Calculated Stress Intensity Factor on the face.
TransMultiplier	real64_array	<i>Datastructure: Hybrid-MimeticDiscretization, Datastructure: TwoPoint-FluxApproximation</i>	An array that holds the transmissibility multipliers
childIndex	localIndex_array	<i>Datastructure: Surface-Generator</i>	Index of child within the mesh object it is registered on.
degreeFromCrackTip	integer_array	<i>Datastructure: Surface-Generator</i>	Distance to the crack tip in terms of topological distance. (i.e. how many nodes are along the path to the closest node that is on the crack surface.
facePressure	real64_array	<i>Datastructure: CompositionalMultiphaseFVM, Datastructure: CompositionalMultiphaseHybridFVM, Datastructure: SinglePhaseFVM, Datastructure: SinglePhaseHybridFVM, Datastructure: SinglePhaseProppant-FVM</i>	An array that holds the pressures at the faces.
freeSurfaceFaceIndicator	localIndex_array	<i>Datastructure: Acoustic-SEM</i>	Free surface indicator, 1 if a face is on free surface 0 otherwise.
isFaceSeparable	integer_array	<i>Datastructure: Surface-Generator</i>	A flag to mark if the face is separable.
parentIndex	localIndex_array	<i>Datastructure: Surface-Generator</i>	Index of parent within the mesh object it is registered on.
primaryCandidateFace	localIndex_array	<i>Datastructure: Surface-Generator</i>	??
ruptureState	integer_array	<i>Datastructure: Surface-Generator</i>	Rupture state of the face: 0=not ready for rupture 1=ready for rupture 2=ruptured.
ruptureTime	real64_array	<i>Datastructure: Surface-Generator</i>	Time that the object was ruptured/split.
neighborData	node		<i>Datastructure: neighbor-Data</i>
sets	node		<i>Datastructure: sets</i>

9.170 Datastructure: FieldSpecification

Name	Type	Description

9.171 Datastructure: FieldSpecifications

Name	Type	Description
Dirichlet	node	<i>Datastructure: Dirichlet</i>
FieldSpecification	node	<i>Datastructure: FieldSpecification</i>
SourceFlux	node	<i>Datastructure: SourceFlux</i>
Traction	node	<i>Datastructure: Traction</i>

9.172 Datastructure: File

Name	Type	Description

9.173 Datastructure: FiniteElementSpace

Name	Type	Description

9.174 Datastructure: FiniteElements

Name	Type	Description
FiniteElementSpace	node	<i>Datastructure: FiniteElementSpace</i>
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.175 Datastructure: FiniteVolume

Name	Type	Description
HybridMimeticDiscretization	node	<i>Datastructure: HybridMimeticDiscretization</i>
TwoPointFluxApproximation	node	<i>Datastructure: TwoPointFluxApproximation</i>

9.176 Datastructure: FlowProppantTransport

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.177 Datastructure: Functions

Name	Type	Description
CompositeFunction	node	<i>Datastructure: CompositeFunction</i>
SymbolicFunction	node	<i>Datastructure: SymbolicFunction</i>
TableFunction	node	<i>Datastructure: TableFunction</i>

9.178 Datastructure: Geometry

Name	Type	Description
BoundedPlane	node	<i>Datastructure: BoundedPlane</i>
Box	node	<i>Datastructure: Box</i>
Cylinder	node	<i>Datastructure: Cylinder</i>
ThickPlane	node	<i>Datastructure: ThickPlane</i>

9.179 Datastructure: HaltEvent

Name	Type	Description
currentSubEvent	integer	Index of the current subevent
eventForecast	integer	Indicates when the event is expected to execute
isTargetExecuting	integer	Index of the current subevent
lastCycle	integer	Last event occurrence (cycle)
lastTime	real64	Last event occurrence (time)
HaltEvent	node	<i>Datastructure: HaltEvent</i>
PeriodicEvent	node	<i>Datastructure: PeriodicEvent</i>
SoloEvent	node	<i>Datastructure: SoloEvent</i>

9.180 Datastructure: HybridMimeticDiscretization

Name	Type	Registered On	Description
TransMultiplier	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the transmissibility multipliers

9.181 Datastructure: Hydrofracture

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.182 Datastructure: Included

Name	Type	Description
File	node	<i>Datastructure: File</i>

9.183 Datastructure: InternalMesh

Name	Type	Description
meshLevels	integer	(no description available)
Level0	node	<i>Datastructure: Level0</i>

9.184 Datastructure: InternalWell

Name	Type	Description
meshLevels	integer	(no description available)
Level0	node	<i>Datastructure: Level0</i>
Perforation	node	<i>Datastructure: Perforation</i>

9.185 Datastructure: InternalWellbore

Name	Type	Description
meshLevels	integer	(no description available)
nx	integer_array	Number of elements in the x-direction within each mesh block
ny	integer_array	Number of elements in the y-direction within each mesh block
xCoords	real64_array	x-coordinates of each mesh block vertex
yCoords	real64_array	y-coordinates of each mesh block vertex
Level0	node	<i>Datastructure: Level0</i>

9.186 Datastructure: LagrangianContact

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.187 Datastructure: LaplaceFEM

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.188 Datastructure: LaplaceVEM

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.189 Datastructure: Level0

Name	Type	Description
meshLevel	integer	(no description available)
ElementRegions	node	<i>Datastructure: ElementRegions</i>
FaceManager	node	<i>Datastructure: FaceManager</i>
edgeManager	node	<i>Datastructure: edgeManager</i>
embeddedSurfacesEdgeManager	node	<i>Datastructure: embeddedSurfacesEdgeManager</i>
embeddedSurfacesNodeManager	node	<i>Datastructure: embeddedSurfacesNodeManager</i>
finiteVolumeStencils	node	<i>Datastructure: finiteVolumeStencils</i>
nodeManager	node	<i>Datastructure: nodeManager</i>

9.190 Datastructure: LinearSolverParameters

Name	Type	Description

9.191 Datastructure: Mesh

Name	Type	Description
InternalMesh	node	<i>Datastructure: InternalMesh</i>
InternalWell	node	<i>Datastructure: InternalWell</i>
InternalWellbore	node	<i>Datastructure: InternalWellbore</i>
PAMELAMeshGenerator	node	<i>Datastructure: PAMELAMeshGenerator</i>

9.192 Datastructure: MeshBodies

Name	Type	Description
InternalMesh	node	<i>Datastructure: InternalMesh</i>
InternalWell	node	<i>Datastructure: InternalWell</i>
InternalWellbore	node	<i>Datastructure: InternalWellbore</i>
PAMELAMeshGenerator	node	<i>Datastructure: PAMELAMeshGenerator</i>

9.193 Datastructure: MultiphasePoromechanics

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.194 Datastructure: NonlinearSolverParameters

Name	Type	Description
newtonNumberOfIterations	integer	Number of Newton's iterations.

9.195 Datastructure: NullModel

Name	Type	Description

9.196 Datastructure: NumericalMethods

Name	Type	Description
FiniteElements	node	<i>Datastructure: FiniteElements</i>
FiniteVolume	node	<i>Datastructure: FiniteVolume</i>

9.197 Datastructure: Outputs

Name	Type	Description
Blueprint	node	<i>Datastructure: Blueprint</i>
ChomboIO	node	<i>Datastructure: ChomboIO</i>
Python	node	<i>Datastructure: Python</i>
Restart	node	<i>Datastructure: Restart</i>
Silo	node	<i>Datastructure: Silo</i>
TimeHistory	node	<i>Datastructure: TimeHistory</i>
VTk	node	<i>Datastructure: VTK</i>

9.198 Datastructure: PAMELAMeshGenerator

Name	Type	Description
meshLevels	integer	(no description available)
Level0	node	<i>Datastructure: Level0</i>

9.199 Datastructure: PackCollection

Name	Type	Description

9.200 Datastructure: ParallelPlatesPermeability

Name	Type	Description
dPerm_dAperture	real64_array3d	(no description available)
dPerm_dPressure	real64_array3d	dPerm_dPressure of the rock.
permeability	real64_array3d	permeability of the rock.

9.201 Datastructure: Parameter

Name	Type	Description

9.202 Datastructure: Parameters

Name	Type	Description
Parameter	node	<i>Datastructure: Parameter</i>

9.203 Datastructure: ParticleFluid

Name	Type	Description
collisionFactor	real64_array	(no description available)
dCollisionFactor_dProppantConcentration	real64_array	(no description available)
dSettlingFactor_dComponentConcentration	real64_array2d	(no description available)
dSettlingFactor_dPressure	real64_array	(no description available)
dSettlingFactor_dProppantConcentration	real64_array	(no description available)
proppantPackPermeability	real64_array	(no description available)
settlingFactor	real64_array	(no description available)

9.204 Datastructure: Perforation

Name	Type	Description

9.205 Datastructure: PeriodicEvent

Name	Type	Description
currentSubEvent	integer	Index of the current subevent
eventForecast	integer	Indicates when the event is expected to execute
isTargetExecuting	integer	Index of the current subevent
lastCycle	integer	Last event occurrence (cycle)
lastTime	real64	Last event occurrence (time)
HaltEvent	node	<i>Datastructure: HaltEvent</i>
PeriodicEvent	node	<i>Datastructure: PeriodicEvent</i>
SoloEvent	node	<i>Datastructure: SoloEvent</i>

9.206 Datastructure: PermeabilityBase

Name	Type	Description
dPerm_dPressure	real64_array3d	dPerm_dPressure of the rock.
permeability	real64_array3d	permeability of the rock.

9.207 Datastructure: PhaseFieldDamageFEM

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.208 Datastructure: PhaseFieldFracture

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.209 Datastructure: PorousDruckerPrager

Name	Type	Description

9.210 Datastructure: PorousElasticIsotropic

Name	Type	Description

9.211 Datastructure: PorousElasticOrthotropic

Name	Type	Description

9.212 Datastructure: PorousElasticTransverselyIsotropic

Name	Type	Description

9.213 Datastructure: PorousExtendedDruckerPrager

Name	Type	Description

9.214 Datastructure: PressurePorosity

Name	Type	Description
dPorosity_dPressure	real64_array2d	(no description available)
oldPorosity	real64_array2d	(no description available)
porosity	real64_array2d	(no description available)
referencePorosity	real64_array	(no description available)

9.215 Datastructure: Problem

Name	Type	Description
Benchmarks	node	<i>Datastructure: Benchmarks</i>
Constitutive	node	<i>Datastructure: Constitutive</i>
ElementRegions	node	<i>Datastructure: ElementRegions</i>
Events	node	<i>Datastructure: Events</i>
FieldSpecifications	node	<i>Datastructure: FieldSpecifications</i>
Functions	node	<i>Datastructure: Functions</i>
Geometry	node	<i>Datastructure: Geometry</i>
Included	node	<i>Datastructure: Included</i>
Mesh	node	<i>Datastructure: Mesh</i>
NumericalMethods	node	<i>Datastructure: NumericalMethods</i>
Outputs	node	<i>Datastructure: Outputs</i>
Parameters	node	<i>Datastructure: Parameters</i>
Solvers	node	<i>Datastructure: Solvers</i>
Tasks	node	<i>Datastructure: Tasks</i>
commandLine	node	<i>Datastructure: commandLine</i>
domain	node	<i>Datastructure: domain</i>

9.216 Datastructure: ProppantSlurryFluid

Name	Type	Description
FluidDensity	real64_array2d	(no description available)
FluidViscosity	real64_array2d	(no description available)
componentDensity	real64_array3d	(no description available)
dCompDens_dCompConc	real64_array4d	(no description available)
dCompDens_dPres	real64_array3d	(no description available)
dDens_dCompConc	real64_array3d	(no description available)
dDens_dPres	real64_array2d	(no description available)
dDens_dProppantConc	real64_array2d	(no description available)
dFluidDens_dCompConc	real64_array3d	(no description available)
dFluidDens_dPres	real64_array2d	(no description available)
dFluidVisc_dCompConc	real64_array3d	(no description available)
dFluidVisc_dPres	real64_array2d	(no description available)
dVisc_dCompConc	real64_array3d	(no description available)
dVisc_dPres	real64_array2d	(no description available)
dVisc_dProppantConc	real64_array2d	(no description available)
density	real64_array2d	(no description available)
viscosity	real64_array2d	(no description available)

9.217 Datastructure: ProppantTransport

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.218 Datastructure: Python

Name	Type	Description

9.219 Datastructure: Restart

Name	Type	Description

9.220 Datastructure: Run

Name	Type	Description

9.221 Datastructure: Silo

Name	Type	Description

9.222 Datastructure: SinglePhaseFVM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
facePressure	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the pressures at the faces.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.223 Datastructure: SinglePhaseHybridFVM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
facePressure	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the pressures at the faces.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.224 Datastructure: SinglePhasePoromechanics

Name	Type	Description
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
LinearSolverParameters	node	<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.225 Datastructure: SinglePhasePoromechanicsEmbeddedFractures

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.226 Datastructure: SinglePhaseProppantFVM

Name	Type	Registered On	Description
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
facePressure	real64_array	<i>Datastructure: FaceManager</i>	An array that holds the pressures at the faces.
LinearSolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
NonlinearSolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.227 Datastructure: SinglePhaseReservoir

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.228 Datastructure: SinglePhaseWell

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>
Well-Controls	node	<i>Datastructure: WellControls</i>

9.229 Datastructure: SolidMechanicsEmbeddedFractures

Name	Type	Description
discretization	string	Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
maxStableDt	real64	Value of the Maximum Stable Timestep for this solver.
Linear-Solver-Parameters	node	<i>Datastructure: LinearSolverParameters</i>
Non-linear-Solver-Parameters	node	<i>Datastructure: NonlinearSolverParameters</i>

9.230 Datastructure: SolidMechanicsLagrangianSSLE

Name	Type	Registered On	Description
maxForce	real64		The maximum force contribution in the problem domain.
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
Acceleration	real64_array2ds	structure: nodeManager	An array that holds the current acceleration on the nodes. This array also is used to hold the summation of nodal forces resulting from the governing equations.
IncrementalDisplacement	real64_array2ds	structure: nodeManager	An array that holds the incremental displacements for the current time step on the nodes.
Mass	real64_array	structure: nodeManager	An array that holds the mass on the nodes.
TotalDisplacement	real64_array2ds	structure: nodeManager	An array that holds the total displacements on the nodes.
Velocity	real64_array2ds	structure: nodeManager	An array that holds the current velocity on the nodes.
contact-Force	real64_array2ds	structure: nodeManager	An array that holds the contact force.
external-Force	real64_array2ds	structure: nodeManager	An array that holds the external forces on the nodes. This includes any boundary conditions as well as coupling forces such as hydraulic forces.
uhatTilde	real64_array2ds	structure: nodeManager	An array that holds the incremental displacement predictors on the nodes.
velocityTilde	real64_array2ds	structure: nodeManager	An array that holds the velocity predictors on the nodes.
Linear-SolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
Nonlinear-SolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.231 Datastructure: SolidMechanics_LagrangianFEM

Name	Type	Registered On	Description
maxForce	real64		The maximum force contribution in the problem domain.
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
Acceleration	real64_array2ds- structure: nodeManager		An array that holds the current acceleration on the nodes. This array also is used to hold the summation of nodal forces resulting from the governing equations.
IncrementalDisplacement	real64_array2ds- structure: nodeManager		An array that holds the incremental displacements for the current time step on the nodes.
Mass	real64_array2ds- structure: nodeManager		An array that holds the mass on the nodes.
TotalDisplacement	real64_array2ds- structure: nodeManager		An array that holds the total displacements on the nodes.
Velocity	real64_array2ds- structure: nodeManager		An array that holds the current velocity on the nodes.
contact-Force	real64_array2ds- structure: nodeManager		An array that holds the contact force.
external-Force	real64_array2ds- structure: nodeManager		An array that holds the external forces on the nodes. This includes any boundary conditions as well as coupling forces such as hydraulic forces.
uhatTilde	real64_array2ds- structure: nodeManager		An array that holds the incremental displacement predictors on the nodes.
velocityTilde	real64_array2ds- structure: nodeManager		An array that holds the velocity predictors on the nodes.
Linear-SolverParameters	node		<i>Datastructure: LinearSolverParameters</i>
Nonlinear-SolverParameters	node		<i>Datastructure: NonlinearSolverParameters</i>

9.232 Datastructure: SoloEvent

Name	Type	Description
currentSubEvent	integer	Index of the current subevent
eventForecast	integer	Indicates when the event is expected to execute
isTargetExecuting	integer	Index of the current subevent
lastCycle	integer	Last event occurrence (cycle)
lastTime	real64	Last event occurrence (time)
HaltEvent	node	<i>Datastructure: HaltEvent</i>
PeriodicEvent	node	<i>Datastructure: PeriodicEvent</i>
SoloEvent	node	<i>Datastructure: SoloEvent</i>

9.233 Datastructure: Solvers

Name	Type	Description
AcousticSEM	node	<i>Datastructure: AcousticSEM</i>
CompositionalMultiphaseFVM	node	<i>Datastructure: CompositionalMultiphaseFVM</i>
CompositionalMultiphaseHybridFVM	node	<i>Datastructure: CompositionalMultiphaseHybridFVM</i>
CompositionalMultiphaseReservoir	node	<i>Datastructure: CompositionalMultiphaseReservoir</i>
CompositionalMultiphaseWell	node	<i>Datastructure: CompositionalMultiphaseWell</i>
EmbeddedSurfaceGenerator	node	<i>Datastructure: EmbeddedSurfaceGenerator</i>
FlowProppantTransport	node	<i>Datastructure: FlowProppantTransport</i>
Hydrofracture	node	<i>Datastructure: Hydrofracture</i>
LagrangianContact	node	<i>Datastructure: LagrangianContact</i>
LaplaceFEM	node	<i>Datastructure: LaplaceFEM</i>
LaplaceVEM	node	<i>Datastructure: LaplaceVEM</i>
MultiphasePoromechanics	node	<i>Datastructure: MultiphasePoromechanics</i>
PhaseFieldDamageFEM	node	<i>Datastructure: PhaseFieldDamageFEM</i>
PhaseFieldFracture	node	<i>Datastructure: PhaseFieldFracture</i>
ProppantTransport	node	<i>Datastructure: ProppantTransport</i>
SinglePhaseFVM	node	<i>Datastructure: SinglePhaseFVM</i>
SinglePhaseHybridFVM	node	<i>Datastructure: SinglePhaseHybridFVM</i>
SinglePhasePoromechanics	node	<i>Datastructure: SinglePhasePoromechanics</i>
SinglePhasePoromechanicsEmbeddedFractures	node	<i>Datastructure: SinglePhasePoromechanicsEmbeddedFractures</i>
SinglePhaseProppantFVM	node	<i>Datastructure: SinglePhaseProppantFVM</i>
SinglePhaseReservoir	node	<i>Datastructure: SinglePhaseReservoir</i>
SinglePhaseWell	node	<i>Datastructure: SinglePhaseWell</i>
SolidMechanicsEmbeddedFractures	node	<i>Datastructure: SolidMechanicsEmbeddedFractures</i>
SolidMechanicsLagrangianSSLE	node	<i>Datastructure: SolidMechanicsLagrangianSSLE</i>
SolidMechanics_LagrangianFEM	node	<i>Datastructure: SolidMechanics_LagrangianFEM</i>
SurfaceGenerator	node	<i>Datastructure: SurfaceGenerator</i>

9.234 Datastructure: SourceFlux

Name	Type	Description

9.235 Datastructure: StrainDependentPermeability

Name	Type	Description
dPerm_dPressure	real64_array3d	dPerm_dPressure of the rock.
permeability	real64_array3d	permeability of the rock.

9.236 Datastructure: SurfaceElementRegion

Name	Type	Description
domainBound- aryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLo- calMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlob- alMap	globalIndex_array	Array that contains a map from localIn- dex to globalIndex.
elementSubRe- gions	node	<i>Datastructure: elementSubRegions</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.237 Datastructure: SurfaceGenerator

Name	Type	Registered On	Description
discretization	string		Name of discretization object (defined in the <i>Numerical Methods</i>) to use for this solver. For instance, if this is a Finite Element Solver, the name of a FiniteElement should be specified. If this is a Finite Volume Method, the name of a FiniteVolume discretization should be specified.
failCriterion	integer		(no description available)
maxStableDt	real64		Value of the Maximum Stable Timestep for this solver.
tipEdges	LvArray_SortedArray< long, long, LvArray_ChaiBuffer >		Set containing all the tip edges
tipFaces	LvArray_SortedArray< long, long, LvArray_ChaiBuffer >		Set containing all the tip faces
tipNodes	LvArray_SortedArray< long, long, LvArray_ChaiBuffer >		Set containing all the nodes at the fracture tip
trailingFaces	LvArray_SortedArray< long, long, LvArray_ChaiBuffer >		Set containing all the trailing faces
K_IC	real64_array2d	<i>Datastructure: FaceManager</i>	Critical Stress Intensity Factor K_{IC} in the plane of the face.
SIFNode	real64_array	<i>Datastructure: nodeManager</i>	Calculated Stress Intensity Factor on the node.
SIF_I	real64_array	<i>Datastructure: edgeManager</i>	Calculated mode 1 Stress Intensity Factor on the node.
SIF_II	real64_array	<i>Datastructure: edgeManager</i>	Calculated mode 2 Stress Intensity Factor on the node.
SIF_III	real64_array	<i>Datastructure: edgeManager</i>	Calculated mode 3 Stress Intensity Factor on the node.
SIFonFace	real64_array	<i>Datastructure: FaceManager</i>	Calculated Stress Intensity Factor on the face.
childIndex	localIndex_array	<i>Datastructure: edgeManager</i>	Index of child within the mesh object it is registered on.
degreeFromCrack	integer_array	<i>Datastructure: nodeManager</i>	Distance to the crack in terms of topological distance. (i.e. how many nodes are along the path to
458		Chapter 9: Datastructure Index	
degreeFromCrackTip	integer_array	<i>Datastructure: nodeManager</i>	Distance to the crack tip in terms of topological distance. (i.e. how many nodes are along the path to

9.238 Datastructure: SymbolicFunction

Name	Type	Description

9.239 Datastructure: TableFunction

Name	Type	Description

9.240 Datastructure: TableRelativePermeability

Name	Type	Description
dPhaseRelPerm_dPhaseVolumeFraction4d		(no description available)
gasOil-RelPermTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >, long, LvArray_ChaiBuffer >	(no description available)
phaseMinVolume-Fraction	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseRelPerm	real64_array3d	(no description available)
phaseTypes	integer_array	(no description available)
waterOil-RelPermTableWrappers	LvArray_Array< geosx_TableFunction_KernelWrapper, 1, camp_int_seq< long, 0l >, long, LvArray_ChaiBuffer >	(no description available)

9.241 Datastructure: Tasks

Name	Type	Description
PackCollection	node	<i>Datastructure: PackCollection</i>
TriaxialDriver	node	<i>Datastructure: TriaxialDriver</i>

9.242 Datastructure: ThickPlane

Name	Type	Description

9.243 Datastructure: TimeHistory

Name	Type	Description
restart	integer	The current history record to be written, on restart from an earlier time allows use to remove invalid future history.

9.244 Datastructure: Traction

Name	Type	Description
component	integer	Component of field (if tensor) to apply boundary condition to
fieldName	string	Name of field that boundary condition is applied to.

9.245 Datastructure: TriaxialDriver

Name	Type	Description

9.246 Datastructure: TwoPointFluxApproximation

Name	Type	Registered On	Description
cellStencil	geosx_CellElementStencilTPFA		(no description available)
edfmStencil	geosx_EmbeddedSurfaceToCellStencil		(no description available)
faceElementTo-CellStencil	geosx_EmbeddedSurfaceToCellStencil		(no description available)
fractureStencil	geosx_SurfaceElementStencil		(no description available)
TransMultiplier	real64_array	<i>Datastructure: Face-Manager</i>	An array that holds the transmissibility multipliers

9.247 Datastructure: VTK

Name	Type	Description

9.248 Datastructure: VanGenuchtenBakerRelativePermeability

Name	Type	Description
dPhaseRelPerm_dPhaseVolFraction	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseRelPerm	real64_array	(no description available)
phaseTypes	integer_array	(no description available)
volFracScale	real64	Factor used to scale the phase capillary pressure, defined as: one minus the sum of the phase minimum volume fractions.

9.249 Datastructure: VanGenuchtenCapillaryPressure

Name	Type	Description
dPhaseCapPressure_dPhaseVolFraction	real64_array	(no description available)
phaseCapPressure	real64_array	(no description available)
phaseOrder	integer_array	(no description available)
phaseTypes	integer_array	(no description available)
volFracScale	real64	Factor used to scale the phase capillary pressure, defined as: one minus the sum of the phase minimum volume fractions.

9.250 Datastructure: WellControls

Name	Type	Description

9.251 Datastructure: WellElementRegion

Name	Type	Description
domainBoundaryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlobalMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
wellControlsName	string	(no description available)
wellGeneratorName	string	(no description available)
elementSubRegions	node	<i>Datastructure: elementSubRegions</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.252 Datastructure: WellElementRegionuniqueSubRegion

Name	Type	Description
domainBoundaryIndicator	integer_array	(no description available)
elementCenter	real64_array2d	(no description available)
elementVolume	real64_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlobalMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
nextWellElementIndex	localIndex_array	(no description available)
nextWellElementIndex-Global	localIndex_array	(no description available)
nodeList	geosx_InterObjectRelation< LvArray_Array< long, 2, camp_int_seq< long, 0l, 1l >, long, LvArray_ChaiBuffer > >	(no description available)
numEdges-PerElement	localIndex	(no description available)
numFaces-PerElement	localIndex	(no description available)
numNodes-PerElement	localIndex	(no description available)
radius	real64_array	(no description available)
topRank	integer	(no description available)
topWellElementIndex	localIndex	(no description available)
wellControlsName	string	(no description available)
Constitutive-Models	node	<i>Datastructure: Constitutive-Models</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>
wellElementSubRegion	node	<i>Datastructure: wellElementSubRegion</i>

9.253 Datastructure: cellBlocks

Name	Type	Description

9.254 Datastructure: cellManager

Name	Type	Description
domainBoundaryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlobalMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
cellBlocks	node	<i>Datastructure: cellBlocks</i>
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.255 Datastructure: commandLine

Name	Type	Description
beginFromRestart	integer	Flag to indicate restart run.
inputFileName	string	Name of the input xml file.
outputDirectory	string	Directory in which to put the output files, if not specified defaults to the current directory.
overridePartitionNumbers	integer	Flag to indicate partition number override
problemName	string	Used in writing the output files, if not specified defaults to the name of the input file.
restartFileName	string	Name of the restart file.
schemaFileName	string	Name of the output schema
suppressPinned	integer	Whether to disallow using pinned memory allocations for MPI communication buffers.
useNonblockingMPI	integer	Whether to prefer using non-blocking MPI communication where implemented (results in non-deterministic DOF numbering).
xPartitionsOverride	integer	Number of partitions in the x-direction
yPartitionsOverride	integer	Number of partitions in the y-direction
zPartitionsOverride	integer	Number of partitions in the z-direction

9.256 Datastructure: domain

Name	Type	Description
Neighbors	<code>std_vector< geosx_NeighborCommunicator, std_allocator< geosx_NeighborCommunicator > ></code>	(no description available)
partition-Manager	<code>geosx_PartitionBase</code>	(no description available)
Constitutive	node	<i>Datastructure: Constitutive</i>
MeshBodies	node	<i>Datastructure: MeshBodies</i>
cellManager	node	<i>Datastructure: cell-Manager</i>

9.257 Datastructure: edgeManager

Name	Type	Registered By	Description
domain-BoundaryIndicator	integer_array		(no description available)
edgesToFractureConnectors	geosx_mapBase< long, long, std_integral_constant< bool, true > >		A map of edge local indices to the fracture connector local indices.
faceList	geosx_InterObjectRelation< LvArray_ArrayOfSets< long, long, LvArray_ChaiBuffer > >		(no description available)
fracture-ConnectorsToEdges	localIndex_array		A map of fracture connector local indices to edge local indices.
fractureConnectorsToElementIndex	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >		A map of fracture connector local indices face element local indices
ghostRank	integer_array		(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >		(no description available)
isExternal	integer_array		(no description available)
localToGlobalMap	globalIndex_array		Array that contains a map from localIndex to globalIndex.
nodeList	geosx_InterObjectRelation< LvArray_Array< long, 2, camp_int_seq< long, 0l, 1l >, long, LvArray_ChaiBuffer > >		(no description available)
SIF_I	real64_array	<i>Datastructure: SurfaceGenerator</i>	Calculated mode 1 Stress Intensity Factor on the node.
SIF_II	real64_array	<i>Datastructure: SurfaceGenerator</i>	Calculated mode 2 Stress Intensity Factor on the node.
SIF_III	real64_array	<i>Datastructure: SurfaceGenerator</i>	Calculated mode 3 Stress Intensity Factor on the node.
childIndex	localIndex_array	<i>Datastructure: SurfaceGenerator</i>	Index of child within the mesh object it is registered on.
parentIndex	localIndex_array	<i>Datastructure: SurfaceGenerator</i>	Index of parent within the mesh object it is registered on.
neighborData	node		<i>Datastructure: neighbor-Data</i>
sets	node		<i>Datastructure: sets</i>

9.258 Datastructure: elementRegionsGroup

Name	Type	Description

9.259 Datastructure: elementSubRegions

Name	Type	Description
WellElementRegionuniqueSubRegion	node	<i>Datastructure: WellElementRegionuniqueSubRegion</i>

9.260 Datastructure: embeddedSurfacesEdgeManager

Name	Type	Description
domainBound-aryIndicator	integer_array	(no description available)
edgesToFracture-Connectors	geosx_mapBase< long, long, std_integral_constant< bool, true > >	A map of edge local indices to the fracture connector local indices.
faceList	geosx_InterObjectRelation< LvArray_ArrayOfSets< long, long, LvArray_ChaiBuffer > >	(no description available)
fractureConnec-torsToEdges	localIndex_array	A map of fracture connector local indices to edge local indices.
fractureConnec-torsToElementIn-dex	LvArray_ArrayOfArrays< long, long, LvAr-ray_ChaiBuffer >	A map of fracture connector local indices face element local indices
ghostRank	integer_array	(no description available)
globalToLo-calMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlob-alMap	globalIndex_array	Array that contains a map from localIndex to globalIndex.
nodeList	geosx_InterObjectRelation< LvArray_Array< long, 2, camp_int_seq< long, 0l, 1l >, long, LvArray_ChaiBuffer > >	(no description available)
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

9.261 Datastructure: embeddedSurfacesNodeManager

Name	Type	Registered By	Description
domain-Bound-aryIndicator	integer_array		(no description available)
edgeList	geosx_InterObjectRelation< LvArray_ArrayOfSets< long, long, LvArray_ChaiBuffer > >		(no description available)
elemList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >		(no description available)
elemRegionList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >		(no description available)
elemSubRegionList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >		(no description available)
ghostRank	integer_array		(no description available)
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >		(no description available)
isExternal	integer_array		(no description available)
local-ToGlobalMap	globalIndex_array		Array that contains a map from localIndex to globalIndex.
parentEdge-GlobalIndex	globalIndex_array		(no description available)
reference-Position	real64_array2d		(no description available)
parent-EdgeIndex	localIndex_array	<i>Datastructure: EmbeddedSurface-Generator</i>	Index of parent edge within the mesh object it is registered on.
neighbor-Data	node		<i>Datastructure: neighbor-Data</i>
sets	node		<i>Datastructure: sets</i>

9.262 Datastructure: finiteVolumeStencils

Name	Type	Description
TwoPointFluxApproximation	node	<i>Datastructure: TwoPointFluxApproximation</i>

9.263 Datastructure: lassen

Name	Type	Description
Run	node	<i>Datastructure: Run</i>

9.264 Datastructure: neighborData

Name	Type	Description

9.265 Datastructure: nodeManager

Name	Type	Registered
ReferencePosition	real64_array2d	
domainBoundaryIndicator	integer_array	
edgeList	geosx_InterObjectRelation< LvArray_ArrayOfSets< long, long, LvArray_ChaiBuffer > >	
elemList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >	
elemRegionList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >	
elemSubRegionList	LvArray_ArrayOfArrays< long, long, LvArray_ChaiBuffer >	
faceList	geosx_InterObjectRelation< LvArray_ArrayOfSets< long, long, LvArray_ChaiBuffer > >	
ghostRank	integer_array	
globalToLocalMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	
isExternal	integer_array	
localToGlobalMap	globalIndex_array	
primaryField	real64_array	
Acceleration	real64_array2d	Datastructure
IncrementalDisplacement	real64_array2d	Datastructure
Mass	real64_array	Datastructure
SIFNode	real64_array	Datastructure
TotalDisplacement	real64_array2d	Datastructure
Velocity	real64_array2d	Datastructure
childIndex	localIndex_array	Datastructure
contactForce	real64_array2d	Datastructure
dampingVector	real64_array	Datastructure
degreeFromCrack	integer_array	Datastructure
degreeFromCrackTip	integer_array	Datastructure
externalForce	real64_array2d	Datastructure
freeSurfaceNodeIndicator	localIndex_array	Datastructure
massVector	real64_array	Datastructure
parentIndex	localIndex_array	Datastructure
pressure_n	real64_array	Datastructure
pressure_nm1	real64_array	Datastructure
pressure_np1	real64_array	Datastructure
rhs	real64_array	Datastructure
ruptureTime	real64_array	Datastructure
stiffnessVector	real64_array	Datastructure
uhatTilde	real64_array2d	Datastructure
velocityTilde	real64_array2d	Datastructure
neighborData	node	
sets	node	

9.266 Datastructure: quartz

Name	Type	Description
Run	node	<i>Datastructure: Run</i>

9.267 Datastructure: sets

Name	Type	Description
externalSet	LvArray_SortedArray< long, long, LvArray_ChaiBuffer >	(no description available)

9.268 Datastructure: wellElementSubRegion

Name	Type	Description
domainBound-aryIndicator	integer_array	(no description available)
ghostRank	integer_array	(no description available)
globalToLo-calMap	geosx_mapBase< long long, long, std_integral_constant< bool, false > >	(no description available)
isExternal	integer_array	(no description available)
localToGlob-alMap	globalIndex_array	Array that contains a map from localIn-dex to globalIndex.
location	real64_array2d	(no description available)
numPerforations-Global	globalIndex	(no description available)
reservoirEle-mentIndex	localIndex_array	(no description available)
reservoirElemen-tRegion	localIndex_array	(no description available)
reservoirEle-mentSubregion	localIndex_array	(no description available)
wellElementIndex	localIndex_array	(no description available)
wellTransmissibil-ity	real64_array	(no description available)
neighborData	node	<i>Datastructure: neighborData</i>
sets	node	<i>Datastructure: sets</i>

CHAPTER 10

Contributors

An up-to-date list of all GEOSX contributors can be found on our Github page:

GEOSX Contributors

The following is the list of GEOSX contributors as of January 2019:

Author Name	Affiliation
Nicola Castelletto	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Benjamin Corbett	<i>Applications, Simulations, and Quality Division, Lawrence Livermore National Laboratory</i>
Matthias Cremon	<i>Department of Energy Resources Engineering, Stanford University</i>
Pengcheng Fu	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Hervé Gross	<i>Total</i>
François Hamon	<i>Total</i>
Jixiang Huang	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Sergey Klevtsov	<i>Department of Energy Resources Engineering, Stanford University</i>
Alexandre Lapene	<i>Total</i>
Antoine Mazuyer	<i>Department of Energy Resources Engineering, Stanford University</i>
Shabnam Semnani	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Randolph Settgaest	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Christopher Sherman	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Arturo Vargas	<i>Applications, Simulations, and Quality Division, Lawrence Livermore National Laboratory</i>
Joshua A. White	<i>Atmospheric, Earth, and Energy Division, Lawrence Livermore National Laboratory</i>
Christopher White	<i>Applications, Simulations, and Quality Division, Lawrence Livermore National Laboratory</i>

CHAPTER 11

Publications

Last updated 23-January-2021

11.1 Preprints and Early-Views

Multigrid reduction preconditioning framework for coupled processes in porous and fractured media

QM Bui, FP Hamon, N Castelletto, D Osei-Kuffuor, RR Settgaest, JA White

arXiv preprint

[arXiv:2101.11649](https://arxiv.org/abs/2101.11649)

Hybrid mimetic finite-difference and virtual element formulation for coupled poromechanics

A Borio, FP Hamon, N Castelletto, JA White, RR Settgaest

arXiv preprint

[arXiv:2010.15470](https://arxiv.org/abs/2010.15470)

Preconditioners for multiphase poromechanics with strong capillarity

JT Camargo, JA White, N Castelletto, RI Borja

International Journal for Numerical and Analytical Methods in Geomechanics

[doi.org:10.1002/nag.3192](https://doi.org/10.1002/nag.3192)

An anisotropic viscoplasticity model for shale based on layered microstructure homogenization

J Choo, SJ Semnani, JA White

International Journal for Numerical and Analytical Methods in Geomechanics

[doi.org:10.1002/nag.3167](https://doi.org/10.1002/nag.3167)

Efficient solvers for hybridized three-field mixed finite element coupled poromechanics

M Frigo, N Castelletto, M Ferronato, JA White

Computers and Mathematics with Applications

[doi.org:10.1016/j.camwa.2020.07.010](https://doi.org/10.1016/j.camwa.2020.07.010)

Simulation of coupled multiphase flow and geomechanics in porous media with embedded discrete fractures

M Cusini, JA White, N Castelletto, RR Settgaest

International Journal for Numerical and Analytical Methods in Geomechanics

[doi:10.1002/nag.3168](https://doi.org/10.1002/nag.3168)

A macroelement stabilization for mixed finite element/finite volume discretizations of multiphase poromechanics

JT Camargo, JA White, RI Borja

Computational Geosciences

[doi:10.1007/s10596-020-09964-3](https://doi.org/10.1007/s10596-020-09964-3)

Approximate inverse-based block preconditioners in poroelasticity

11.1. Preprints and Early-Views
A Franceschini, N Castelletto, M Ferronato

Computational Geosciences

[doi:10.1007/s10596-020-09981-2](https://doi.org/10.1007/s10596-020-09981-2)

11.2 2021

Enhanced multiscale restriction-smoothed basis (MsRSB) preconditioning with applications to porous media flow and geomechanics

SBM Bosma, S Klevtsov, O Møyner, N Castelletto

Journal of Computational Physics

[doi.org:10.1016/j.jcp.2020.109934](https://doi.org/10.1016/j.jcp.2020.109934)

11.3 2020

Nonlinear multigrid based on local spectral coarsening for heterogeneous diffusion problems

CS Lee, F Hamon, N Castelletto, PS Vassilevski, JA White
 Computer Methods in Applied Mechanics and Engineering
[doi:10.1016/j.cma.2020.113432](https://doi.org/10.1016/j.cma.2020.113432)

An inelastic homogenization framework for layered materials with planes of weakness

SJ Semnani, JA White
 Computer Methods in Applied Mechanics and Engineering
[doi:10.1016/j.cma.2020.113221](https://doi.org/10.1016/j.cma.2020.113221)

Algebraically stabilized Lagrange multiplier method for frictional contact mechanics with hydraulically active fractures.

A Franceschini, N Castelletto, JA White, HA Tchelepi
 Computer Methods in Applied Mechanics and Engineering
[doi:10.1016/j.cma.2020.113161](https://doi.org/10.1016/j.cma.2020.113161)

Multi-stage preconditioners for thermal–compositional–reactive flow in porous media.

MA Cremon, N Castelletto, JA White
 Journal of Computational Physics
[doi:10.1016/j.jcp.2020.109607](https://doi.org/10.1016/j.jcp.2020.109607)

Scalable multigrid reduction framework for multiphase poromechanics of heterogeneous media

QM Bui, D Osei-Kuffuor, N Castelletto, JA White
 SIAM Journal on Scientific Computing
[doi:10.1137/19M1256117](https://doi.org/10.1137/19M1256117)

Fully implicit multidimensional hybrid upwind scheme for coupled flow and transport

F Hamon, B Mallison
 Computer Methods in Applied Mechanics and Engineering
[doi:10.1016/j.cma.2019.112606](https://doi.org/10.1016/j.cma.2019.112606)

11.4 2019

A two-stage preconditioner for multiphase poromechanics in reservoir simulation

JA White, N Castelletto, S Klevtsov, QM Bui, D Osei-Kuffuor, HA Tchelepi

Computer Methods in Applied Mechanics and Engineering

[doi:10.1016/j.cma.2019.112575](https://doi.org/10.1016/j.cma.2019.112575)**Multiscale two-stage solver for Biot's poroelasticity equations in subsurface media**

N Castelletto, S Klevtsov, H Hajibeygi, HA Tchelepi

Computational Geosciences

[doi:10.1007/s10596-018-9791-z](https://doi.org/10.1007/s10596-018-9791-z)**Block preconditioning for fault/fracture mechanics saddle-point problems**

A Franceschini, N Castelletto, M Ferronato

Computer Methods in Applied Mechanics and Engineering

[doi:10.1016/j.cma.2018.09.039](https://doi.org/10.1016/j.cma.2018.09.039)**A relaxed physical factorization preconditioner for mixed finite element coupled poromechanics**

M Frigo, N Castelletto, M Ferronato

SIAM Journal on Scientific Computing

[doi:10.1137/18M120645X](https://doi.org/10.1137/18M120645X)

CHAPTER 12

Acknowledgements

GEOSX was developed with supporting funds from a number of organizations, including

- Lawrence Livermore National Laboratory
- U.S. Department of Energy, Office of Science
- Total, S.A.

This support is gratefully acknowledged.

CHAPTER 13

Indices and tables

- `genindex`
- `modindex`
- `search`

G

`get_group()` (*pygeosx.Group method*), 236
`get_wrapper()` (*pygeosx.Group method*), 236
`groups()` (*pygeosx.Group method*), 236

P

`pygeosx.apply_initial_conditions()`
 (*built-in function*), 235
`pygeosx.COMPLETED` (*built-in variable*), 236
`pygeosx.finalize()` (*built-in function*), 235
`pygeosx.Group` (*built-in class*), 236
`pygeosx.initialize()` (*built-in function*), 235
`pygeosx.INITIALIZED` (*built-in variable*), 236
`pygeosx.READY_TO_RUN` (*built-in variable*), 236
`pygeosx.reinit()` (*built-in function*), 235
`pygeosx.run()` (*built-in function*), 236
`pygeosx.UNINITIALIZED` (*built-in variable*), 236
`pygeosx.Wrapper` (*built-in class*), 236

R

`register()` (*pygeosx.Group method*), 236

V

`value()` (*pygeosx.Wrapper method*), 236

W

`wrappers()` (*pygeosx.Group method*), 236