Vaango Developers Manual

Version 20 .7 .16

July 16, 2020
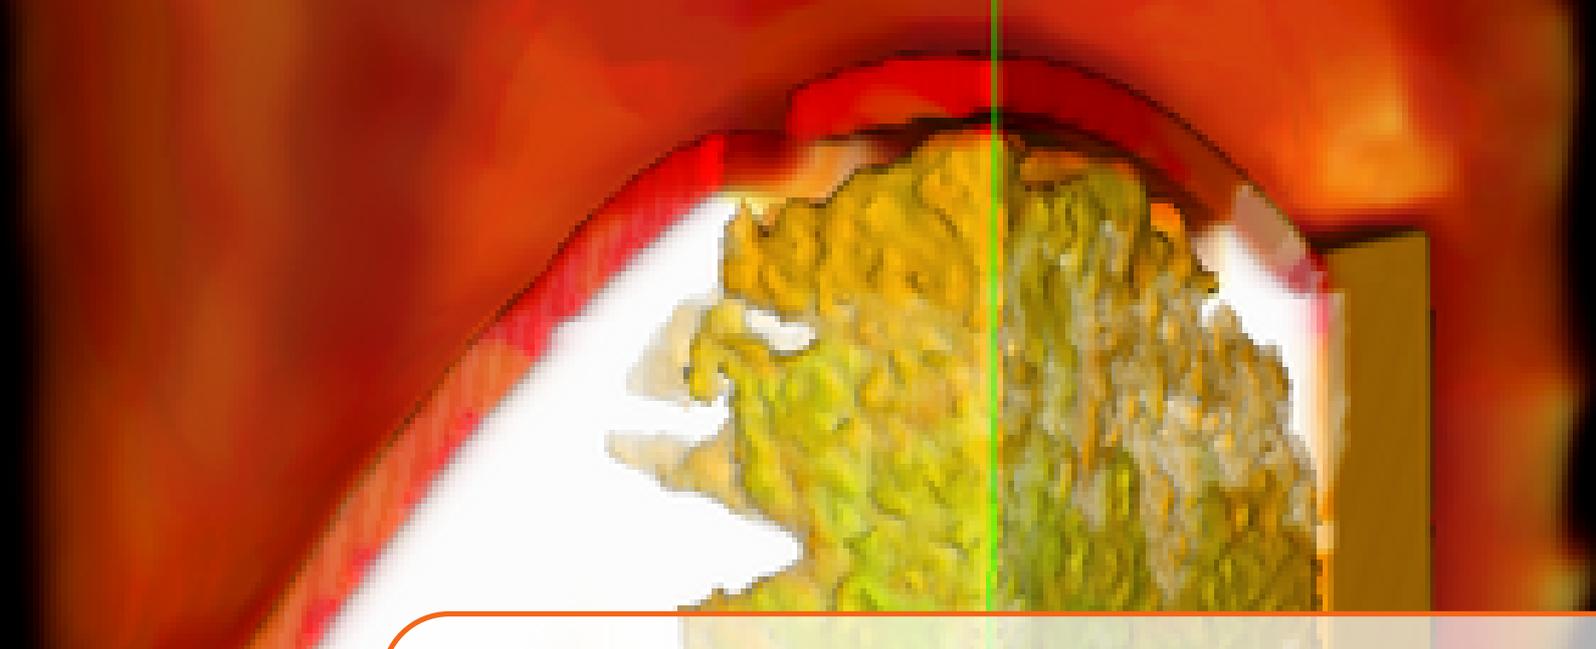
The Utah Uintah team

and

Biswajit Banerjee

Contents

1 — The Vaango framework

1.1 Historical information

Vaango is a fork of the Uintah Computational Framework (Uintah) created in 2012 to allow for the development of tools to solve solid mechanics problems in mechanical and civil engineering. The orginal Uintah code continues to be actively developed, but the focus of that code is multiphysics problems, particularly computational fluid dynamics (CFD) and chemical engineering. Around once a year, the underlying parallel computing infrastructure of Vaango is updated to keep up with developments in Uintah.

The Vaango framework, like Uintah, consists of a set of software components and libraries that facilitate the solution of Partial Differential Equations (PDEs) on Structured AMR (SAMR) grids using hundreds to thousands of processors. However, unlike the CFD problems that Uintah was designed for, the use of the grid is often only incidental to the solution of the governing PDEs of solid mechanics.

1.2 Overview

One of the challenges in designing a parallel, component-based multi-physics application is determining how to efficiently decompose the problem domain. Components, by definition, make local decisions. Yet parallel efficiency is only obtained through a globally optimal domain decomposition and scheduling of computational tasks. Typical techniques include allocating disjoint sets of processing resources to each component, or defining a single domain decomposition that is a compromise between the ideal load balance of multiple components. However, neither of these techniques will achieve maximum efficiency for complex multi-physics problems.

Vaango uses a non-traditional approach to achieving parallelism, employing an abstract taskgraph representation to describe computation and communication. The taskgraph is an explicit representation of the computation and communication that occur in the coarse of a single iteration of the simulation (typically a timestep or nonlinear solver iteration) see figure 1.1. Vaango components delegate decisions about parallelism to a scheduler component, using variable dependencies to describe communication patterns and characterizing computational workloads to facilitate a global resource optimization. The taskgraph representation has a number of advantages, including efficient fine-grained coupling of multi-physics components, flexible load balancing mechanisms and a separation of application concerns from parallelism concerns. However, it creates a challenge for scalability which we overcome by creating an implicit definition of this graph and representing it in a distributed fashion.
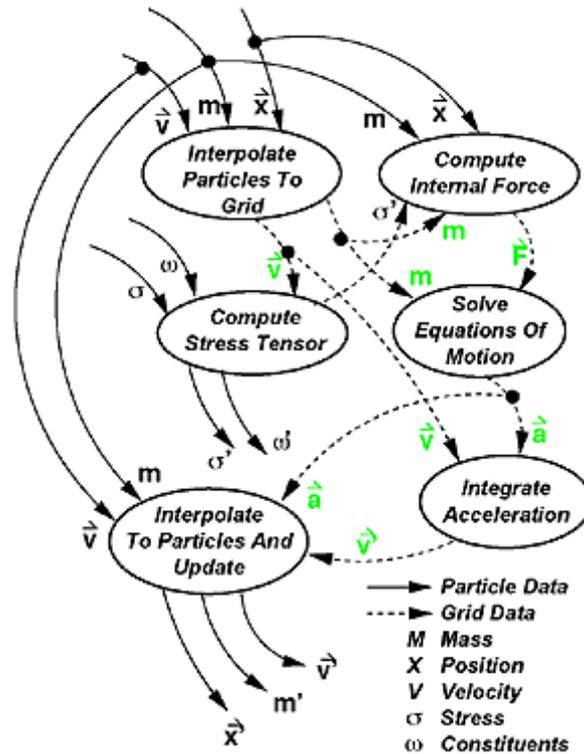
Figure 1.1: Example Task Graph

The primary advantage of a component-based approach is that it facilitates the separate development of simulation algorithms, models, and infrastructure. Components of the simulation can evolve independently. The component-based architecture allows pieces of the system to be implemented in a rudimentary form at first and then evolve as the technologies mature. Most importantly, VAANGO allows the aspects of parallelism (schedulers, load-balancers, parallel input/output, and so forth) to evolve independently of the simulation components. Furthermore, components enable replacement of computation pieces without complex decision logic in the code itself.

1.3   An example

The sequence of steps performed in a simulation can be seen in the abbreviated example below. A complete version of this example can be found in the unit test for TabularPlasticity located at src/CCA/Components/MPM/ConstitutiveModel/UnitTests/testTabularPlasticity.cc .

```
try {
    // Read the input file
    ProblemSpecP ups = VaangoEnv::createInput();
    ups->getNode()->_private = (void *) ups_file.c_str();

    // Create the MPI/threading environment
    const ProcessorGroup* world = Uintah::Parallel::getRootProcessorGroup();

    // Create the simulation controller
    SimulationController* ctl = scinew AMRSimulationController(world, false, ups);

    // Create a regridder if needed
    RegridderCommon* reg = 0;

    // Create an implicit solver if needed
    SolverInterface* solve = SolverFactory::create(ups, world, "");

    // Create the component for the simulation (MPM/MPMICE/Peridynamics)
```

```cpp
    UintahParallelComponent* comp = ComponentFactory::create(ups, world, false, "");

    // Create the simulation base object
    SimulationInterface* sim = dynamic_cast<SimulationInterface*>(comp);
    ctl->attachPort("sim", sim);
    comp->attachPort("solver", solve);
    comp->attachPort("regridder", reg);

    // Create a load balancer
    LoadBalancerCommon* lbc = LoadBalancerFactory::create(ups, world);
    lbc->attachPort("sim", sim);

    // Create a data archiver
    DataArchiver* dataarchiver = scinew DataArchiver(world, -1);
    Output* output = dataarchiver;
    ctl->attachPort("output", dataarchiver);
    dataarchiver->attachPort("load balancer", lbc);
    comp->attachPort("output", dataarchiver);
    dataarchiver->attachPort("sim", sim);

    // Create a task scheduler
    SchedulerCommon* sched = SchedulerFactory::create(ups, world, output);
    sched->attachPort("load balancer", lbc);
    ctl->attachPort("scheduler", sched);
    lbc->attachPort("scheduler", sched);
    comp->attachPort("scheduler", sched);
    sched->setStartAddr( start_addr );
    sched->addReference();

    // Run the simulation
    ctl->run();

    // Clean up after the simulation is complete
    delete ctl;
    sched->removeReference();
    delete sched;
    delete lbc;
    delete sim;
    delete solve;
    delete output;

} catch (ProblemSetupException& e) {
  std::cout << e.message() << std::endl;
  thrownException = true;
} catch (Exception& e) {
  std::cout << e.message() << std::endl;
  thrownException = true;
} catch (...) {
  std::cout << "**ERROR** Unknown exception" << std::endl;
  thrownException = true;
}
```
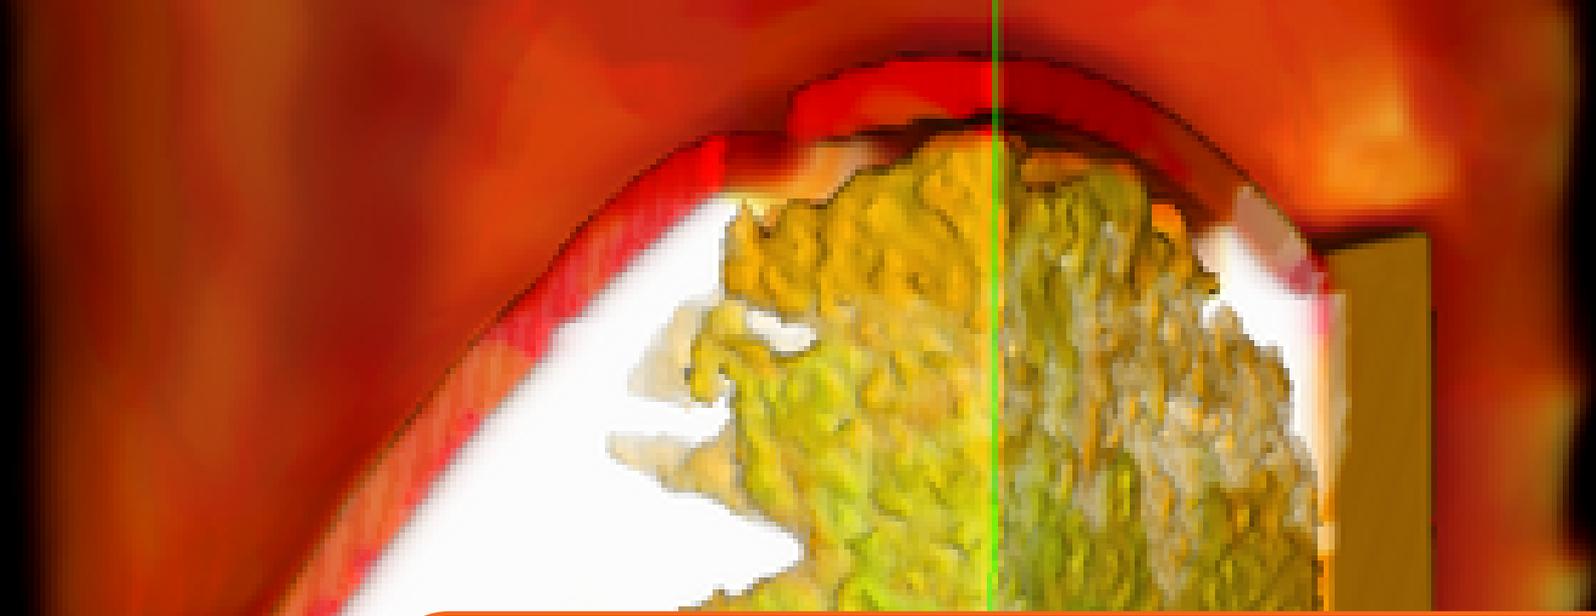
2 — General concepts

This chapter discusses the main concepts used in a Vaango simulation. These concepts are independent of the type of problem being considered and are core to the computational framework. Special treatment is needed for components such as MPM or Peridynamics.

2.1 Scheduler

The Scheduler in Vaango is responsible for determining the order of tasks and ensuring that the correct inter-processor data is made available when necessary. Each software component passes a set of tasks to the scheduler. Each task is responsible for computing some subset of variables, and may require previously computed variables, possibly from different processors. The scheduler will then compile this task information into a task graph, and the task graph will contain a sorted order of tasks, along with any information necessary to perform inter-process communication via MPI or threading. Then, when the scheduler is executed, the tasks will execute in the pre-determined order.

2.1.1 needRecompile()

Each component has a needRecompile() function that is called once per timestep. If, for whatever reason, a Component determines that the list of tasks it had previously scheduled is no longer valid, then the Component must return 'true' when its needRecompile() function is called. This will cause the scheduler to rebuild the task graph (by asking each component to re-specify tasks). Note, rebuilding the taskgraph is a relatively expensive operation, so only should be done if necessary.

2.2 Tasks

A task contains two essential components: a pointer to a function that performs the actual computations, and the data inputs and outputs, i.e. the data dependencies required by the function. When a task requests a previously computed variable from the data warehouse, the number of ghost cells are also specified. The Unitah framework uses the ghost cell information to excecute inter-process communication to retrieve the necessary ghost cell data.

An example of a task description is presented showing the essential features that are commonly used by the application developer when implementing an algorithm within the Vaango framework. The task component is assigned a name and in this particular example, it is called taskexample and a func-

tion pointer, &Example::taskexample . Following the instantiation of the task itself, the dependency information is assigned to the tasks. In the following example, the task requires data from the previous timestep (Task::OldDW) associated with the name variable1_label and requires one ghost node (Ghost::AroundNodes,1) level of information which will be retrieved from another processor via MPI. In addition, the task will compute two new pieces of data each associated with different variables, i.e. variable1_label , and variable2_label . Finally, the task is added to the scheduler component with specifications about what patches and materials are associated with the actual computation.

```
Task* task = scinew Task("Example::taskexample",this, &Example::taskexample);
task->requires(Task::OldDW, variable1_label, Ghost::AroundNodes, 1);
task->computes(variable1_label);
task->computes(variable2_label);
sched->addTask(task, level->eachPatch(), sharedState_->allMaterials());
```

For more complex problems involving multiple materials and multi-physics calculations, a subset of the materials may only be used in the calculation of particular tasks. The Vaango framework allows for the independent scheduling and computation of multi-material within a multi-physics calculation.

2.3 Simulation Component Class Description

Each Vaango component can be described as a C++ class that is derived from two other classes: UintahParallelComponent and a SimulationInterface . The new derived class must provide the following virtual methods: problemSetup , scheduleInitialize , scheduleComputeStableTimestep , and scheduleTimeAdvance . Here is an example of the typical *.h file that needs to be created for a new component.

```
class Example : public UintahParallelComponent, public SimulationInterface {
  public:

    virtual void problemSetup(const ProblemSpecP& params, const ProblemSpecP&
        restart_prob_spec, GridP& grid, SimulationStateP&);

    virtual void scheduleInitialize(const LevelP& level,SchedulerP& sched);

    virtual void scheduleComputeStableTimestep(const LevelP& level, SchedulerP&);

    virtual void scheduleTimeAdvance(const LevelP& level, SchedulerP&);

  private:
    Example(const ProcessorGroup* myworld);
    virtual ~Example();

    void initialize(const ProcessorGroup*, const PatchSubset* patches, const
        MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw);

    void computeStableTimestep(const ProcessorGroup*, const PatchSubset* patches,
        const MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw);

    void timeAdvance(const ProcessorGroup*, const PatchSubset* patches, const
        MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw);
}
```

Each new component inherits from the classes UintahParallelComponent and SimulationInterface . The component overrides default implementations of various methods. The above methods are the essential functions that a new component must implement. Additional methods to do AMR will be described as more complex examples are presented.

The roles of each of the scheduling methods are described below. Each scheduling method, i.e. scheduleInitialize , scheduleComputeStableTimestep , and scheduleTimeAdvance describe

2.3.1 ProblemSetup

The purpose of this method is to read a problem specification which requires a minimum of information about the grid used, time information, i.e. time step size, length of time for simulation, etc, and where and what data is actually saved. Depending on the problem that is solved, the input file can be rather complex, and this method would evolve to establish any and all parameters needed to initially setup the problem.

2.3.2 ScheduleInitialize

The purpose of this method is to initialize the grid data with values read in from the problemSetup and to define what variables are actually computed in the TimeAdvance stage of the simulation. A task is defined which references a function pointer called initialize .

2.3.3 ScheduleComputeStableTimestep

The purpose of this method is to compute the next timestep in the simulation. A task is defined which references a function pointer called computeStableTimestep .

2.3.4 ScheduleTimeAdvance

The purpose of this method is to schedule the actual algorithmic implementation. For simple algorithms, there is only one task defined with a minimal set of data dependencies specified. However, for more complicated algorithms, the best way to schedule the algorithm is to break it down into individual tasks. Each task of the algorithm will have its own data dependencies and function pointers that reference individual computational methods.

2.4 Data Storage Concepts

During the course of the simulation, data is computed and stored in a data structure called the DataWarehouse. The DataWarehouse is an abstraction for storage, retrieval, and access of Vaango simulation data. The Data warehouse presents a localized view of the global data distribution.

2.4.1 Data Archiver

The Data Archiver is a component of the framework that allows for reading, saving, and accessing simulation data. It presents a global shared memory abstraction to the simulation data. The component developer does not have to worry about retrieving data that has been produced on a remote processor or sending data from a local processor to another processor for use. The framework takes care of these tasks implicitly for the component.

2.4.2 Two Data Warehouses

During each time step (assuming a single level problem), there are two data warehouses associated with the simulations. The `new_dw` and the `old_dw` (as they are commonly referred to in the code). The `old_dw` contains data that was generated in the previous timestep, and may not be modified. Data that is generated during the current timestep will be placed in the `new_dw` . At the end of a timestep, the current `old_dw` is deleted, and then replaced with the current `new_dw` . Then a new (empty) `new_dw` is created.

At the end of a time step, some data is automatically transferred from the current time step's `new_dw` into the next time step's `old_dw` .
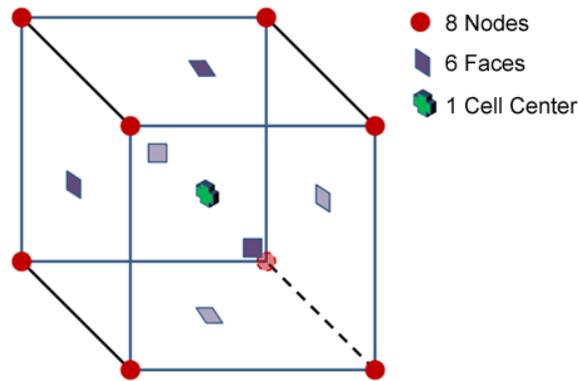
Figure 2.1: Variable locations with respect to a single cell.

2.4.3   Storing and Retrieving Variables

In order to store data to the data warehouse, or to retrieve data, several things are required. First, the task (which procedure wishes access to the data) must have registered this information with the Scheduler during task creation time. Then, inside of the task, the following call is used to pull the data from the data warehouse:

```
SFCXVariable<double> uVelocity;
new_dw->get( uVelocity, d_lab->d_uVelocitySPBCLabel, matlIndex, patch, Ghost::
    AroundFaces, Arches::ONEGHOSTCELL );
```

Similarly, to put data into the data warehouse, use this call:

```
PerPatch<CellInformationP> cellInfoP;
new_dw->put( cellInfoP, d_lab->d_cellInfoLabel, matlIndex, patch );
```

2.4.4   Variables

There are three (general) types of variables in VAANGO: Particle, Grid, and Reduction. Each type of variable is discussed below. Remember, in order to interact with the Data Warehouse, each variable must have a corresponding label.

Particle variables contain information about particles.

A grid variable is a representation of data across (usually) the entire computational domain. It can be used to represent temperature, volume, velocity, etc. It is implemented using a 3D array.

A reduction, in terms of distributed software (using, for example, MPI) is a point in which all (or some defined set of) processors all communicate with each other. It usually is an expensive operation (because it requires all processors to synchronize with each other and pass data). However, many times this operation cannot be avoided. VAANGO provides a "Reduction Variable" to facilitate this operation. Common reductions operations include finding the min or max of a single number on all processors, or creating the sum of a number from every processor, and returning the result to all processors.

Variables can be stored at several different locations. Variables are defined with respect to how they relate to a single cell in the computational domain. They may be located on the faces (FC) or nodes (NC) of the cell, or in the center (CC) of the cell (see Figure 2.1).

2.5   Grid data

Data (Variables) used by the framework are stored on the Grid. The Grid is comprised of one or more (in the case of AMR) Levels. Each level contains one or more patches; and each patch contains a number of
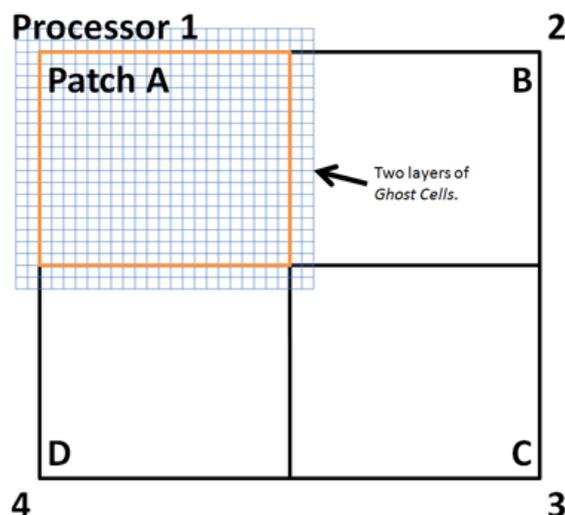
Figure 2.2: Schematic of ghost cell locations laid out with respect to patches.

cells. The simulation data itself exists in (at) each cell (and is stored in a 'variable' which is implemented as a 3D array of the data in each cell across the patch).

2.5.1 Cells

Data in (at) each cell can be specified in several ways. Specifically the data can be Node centered (NC), Cell centered (CC), or Face centered (FC). For CC data, there is one value associated with each cell; for NC data, 4 values; and for FC data, 6 values.

2.5.2 Patches

For calculations to take place on a patch, information from bordering patches is required. This information is stored in boundary cells.

2.5.3 Boundary Cells

Boundary cells (Figure 2.2) represent portions of the computational domain outside the boundaries of the data assigned to a given processor. Specifically, they are almost always used as boundary cells to a patch that are necessary for computation, but are not "owned" by the patch that is currently being computed. As can be seen in Figure 2.2, patch A owns all the cells within the orange rectangle, but also has two layers of ghost cells. Portions of these cells are actually owned by patches B, C, and D. (The data found in these cells is automatically transfered (by the framework) from processors 2, 3, and 4 to processor 1 in order for the cell data to be available for use in computations on patch A.)

Many operations will require a stencil consisting of several cells in order to perform calculations at a given cell. However, at the border of a patch, there are no cells belonging to that patch that contain the required information - that information is "owned" by another processor. In this situation, data that belongs to another patch must be accessed. This is the purpose of ghost cells. The Data Warehouse takes care of moving the required ghost data from the "owner" processor to the neighbor processor so that the individual task can assume the required data is available.

In summary, ghost cells exist between patches on the same level. They are cells that are owned by the neighboring patch but are required for computation due to the stencil width.

Extra cells exist at the edge of the computational domain and are used for boundary conditions. Extra

cells exist on the boundaries of the level where no neighboring patches exist. This can either be the edge of a domain or on a coarse-fine interface.

2.5.4   Indexing grid cells

VAANGO uses a straightforward indexing scheme. Each cell on a given level is uniquely identified by an x,y,z coordinate (stored in an `IntVector`). An IntVector is a vector of 3 integers that represent an X,Y,Z coordinate. See Core/Grid/Level.cc/h for more information. The following pseudo code shows how indices across levels are mapped to each other.

```
IntVector
Level::mapCellToCoarser(const IntVector& idx) const
{
  IntVector ratio = idx/d_refinementRatio;

  return ratio;
}

IntVector
Level::mapCellToFiner(const IntVector& idx) const
{
  IntVector r_ratio = grid->getLevel(d_index+1)->d_refinementRatio;
  IntVector fineCell = idx*r_ratio;

  return fineCell;
}
```

2.6   Saving Simulation Data

The computational framework automatically saves data as specified in the .ups file. This works for all data that is stored in the Data Warehouse. However, there are times when a component will need to save its own data. (It is preferable that the Data Warehouse be updated to manage this Component data, but sometimes this is not expedient...) In these cases, here is an outline on how to save that data:

- Create a function that will save your data:

  ```
  void
  Component::saveMyData( const PatchSubset * patches )
  {
    ...
  }
  ```

- From the last Task in your algorithm, call the function:

  ```
  saveMyData( patches );
  ```

- In the function, you need to do several things: Make sure that it is an output timestep. (Your component must have saved a pointer to the Data Archiver. Most components already do this... if not, you can do it in Component::problemSetup().)

  ```
  if( dataArchiver_->isOutputTimestep() ) {
    // Dump your data... (see below)
  }
  ```

- You need to create a directory to put the data in:

  ```
  const int     & timestep = d_sharedState->getCurrentTopLevelTimeStep();
  ostringstream   data_dir;

  // This gives you ".../simulation.uda.###/t#####/mydata/
  data_dir << dataArchiver()->getOutputLocation() << "/t" << setw(5) << setfill('
      0') << timestep << "/mydata";

  int result = MKDIR( data_dir.c_str(), 0777 );  // Make sure you #include <Core/
      OS/Dir.h>
  ```

```
if( result != 0 ) {
  ostringstream error;
  error << "Component::saveMyData(): couldn't create directory '") + data_dir.
      str() + "' (" << result << ").";
  throw InternalError( error.str(), __FILE__, __LINE__ );
}
```

- Loop over the patches and save your data... (Note, depending on how much data you have, you need to determine whether you want to write the files as a binary file, or as an ascii file. Ascii files are easier for a human to look at for errors, but take a lot more time to read and space to write.)

```
for( int pIndex = 0; pIndex < patches->size(); pIndex++ ){
  const Patch * patch = patches->get( pIndex );
  ostringstream file_name;
  file_name << data_dir << "/p" << setw(5) << setfill('0') << pIndex;

  ofstream output( ceFileName, ios::out | ios::binary);
  if( !output ) {
    throw InternalError( string( "Component::saveMyData(): couldn't open file '
        ") + file_name.str() + "'.", __FILE__, __LINE__ );
  }

  // Output data
  for( loop over data ) {
    output << data...
  }

  output.close();
} // end for pIndex
```
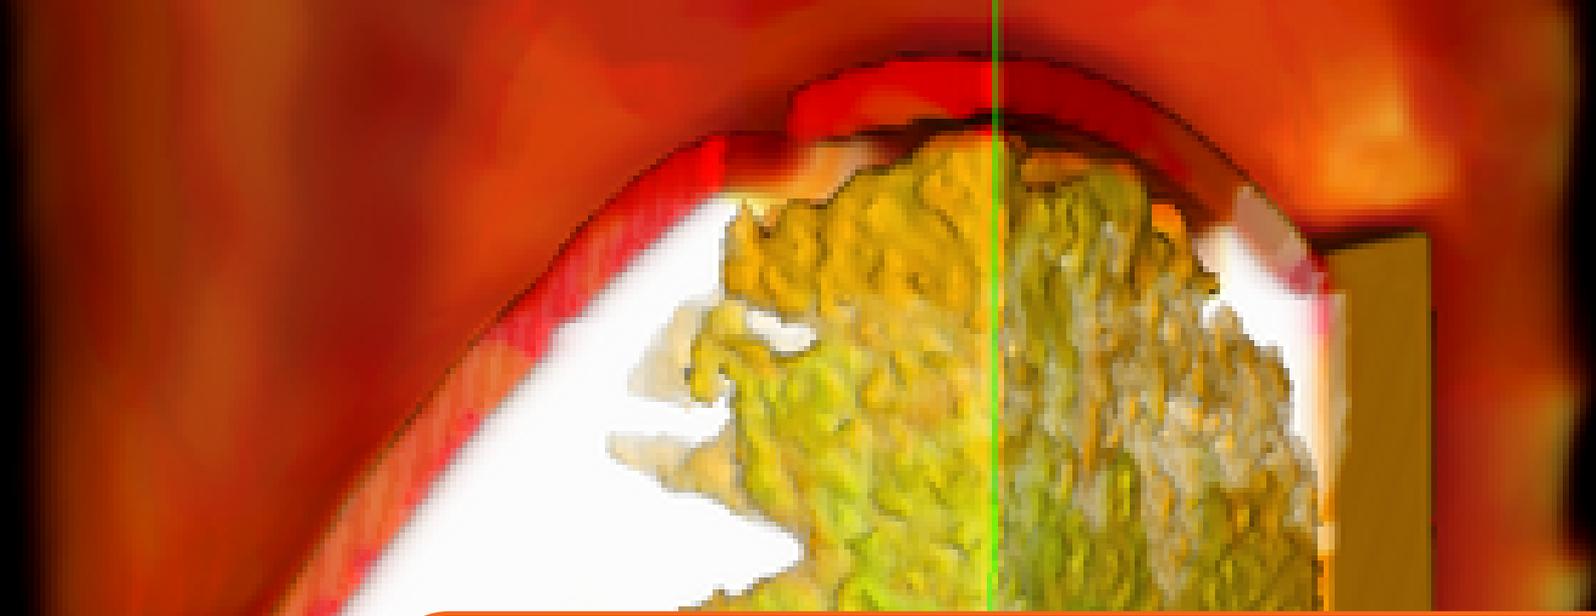
3 — PDE Examples

This chapter will describe a set of example problems showing various stages of algorithm complexity and how the Vaango framework is used to solve the discretized form of the solutions. Emphasis will not be on the most efficient or fast algorithms, but intead will demonstrate straightforward implementations of well known algorithms within the Vaango Framework. Several examples will be given that show an increasing level of complexity.

All examples described are found in the directory src/CCA/Components/Examples .

3.1   Poisson1

Poisson1 solves Poisson's equation on a grid using Jacobi iteration. Since this is not a time dependent problem and Vaango is fundamentally designed for time dependent problems, each Jacobi iteration is considered to be a timestep. The timestep specified and computed is a fixed value obtained from the input file and has no bearing on the actual computation.

The following equation is discretized and solved using an iterative method. Each timestep is one iteration. At the end of the timestep, we the residual is computed showing the convergence of the solution and the next iteration is computed until.

The following shows a simplified form of the Poisson1 of the .h and .cc files found in the Examples directory. The argument list for some of the methods are eliminated for readibility purposes. Please refer to the actual source for a complete description of the arguments required for each method.

```
class Poisson1 : public UintahParallelComponent , public SimulationInterface {
public:
  Poisson1(const ProcessorGroup* myworld);
  virtual ~Poisson1();
  virtual void problemSetup(const ProblemSpecP& params , const ProblemSpecP&
      restart_prob_spec , GridP& grid , SimulationStateP&);
  virtual void scheduleInitialize(const LevelP& level ,SchedulerP& sched);
  virtual void scheduleComputeStableTimestep(const LevelP& level ,SchedulerP&);
  virtual void scheduleTimeAdvance(const LevelP& level ,SchedulerP&);

private:
  void initialize(const ProcessorGroup*, const PatchSubset* patches , const
      MaterialSubset* matls , DataWarehouse* old_dw , DataWarehouse* new_dw);
  void computeStableTimestep(const ProcessorGroup*,const PatchSubset* patches , const
      MaterialSubset* matls ,DataWarehouse* old_dw , DataWarehouse* new_dw);
  void timeAdvance(const ProcessorGroup ,const PatchSubset* patches , const
      MaterialSubset* matls ,DataWarehouse* old_dw , DataWarehouse* new_dw);
```

```
    SimulationStateP  sharedState_;
    double  delt_;
    const  VarLabel*  phi_label;
    const  VarLabel*  residual_label;
    SimpleMaterial*  mymat_;

    Poisson1(const  Poisson1&);
    Poisson1&  operator=(const  Poisson1&);
  };
```

The private methods and data shown are the functions that are function pointers referred to in the task descriptions. The VarLabel data type stores the names of the various data that can be referenced uniquely by the data warehouse. The SimulationStateP data type is essentially a global variable that stores information about the materials that are needed by other internal Vaango framework components. SimpleMaterial is a data type that refers to the material properties.

Within each schedule function, i.e. sheduleInitialize, scheduleComputeStableTimestep, and scheduleTimeAdvance, a task is specified that has a function pointer associated with it. The function pointers point to the actual implementation of the specific task and have a different argument list than the associated schedule method.

The typical task implementation, i.e. timeAdvance() contains the following arguments: ProcessorGroup, PatchSubset, MaterialSubset, and two DataWarehouse objects. The purpose of the ProcessorGroup is to hold various MPI information such as the MPI_Communicator, the rank of the process and the number of processes that are actually being used.

3.1.1 Description of Scheduling Functions

The actual implementation with descriptions are presented following the code snippets.

```
Poisson1::Poisson1(const  ProcessorGroup*  myworld)
 :  UintahParallelComponent(myworld)
{
  phi_label  =  VarLabel::create("phi",
  NCVariable<double>::getTypeDescription());
  residual_label  =  VarLabel::create("residual",
  sum_vartype::getTypeDescription());

}

Poisson1::~Poisson1()
{
  VarLabel::destroy(phi_label);
  VarLabel::destroy(residual_label);
}
```

Typical constructor and destructor for simple examples where the data label names (phi and residual) are created for data wharehouse storage and retrieval.

```
void  Poisson1::problemSetup(const  ProblemSpecP&  params,  const  ProblemSpecP&
    restart_prob_spec,  GridP&  /*grid*/,  SimulationStateP&  sharedState)
{
  sharedState_  =  sharedState;
  ProblemSpecP  poisson  =  params->findBlock("Poisson");

  poisson->require("delt",  delt_);

  mymat_  =  scinew  SimpleMaterial();

  sharedState->registerSimpleMaterial(mymat_);
}
```

The problemSetup is based in a xml description of the input file. The input file is parsed and the delt tag is set. The sharedState is assigned and is used to register a material and store it for later use by the

Vaango internals.

```
void Poisson1::scheduleInitialize(const LevelP& level, SchedulerP& sched)
{
  Task* task = scinew Task("Poisson1::initialize",
  this, &Poisson1::initialize);

  task->computes(phi_label);
  task->computes(residual_label);
  sched->addTask(task, level->eachPatch(), sharedState_->allMaterials());
}
```

A task is defined which contains a name and a function pointer, i.e. initialize which is described later in Poisson1.cc The task defines two variables that are computed in the initialize function, phi and residual. The task is then added to the scheduler. This task is only computed once at the beginning of the simulation.

```
void Poisson1::scheduleComputeStableTimestep(const LevelP& level, SchedulerP& sched)
{
  Task* task = scinew Task("Poisson1::computeStableTimestep",
  this, &Poisson1::computeStableTimestep);

  task->requires(Task::NewDW, residual_label);
  task->computes(sharedState_->get_delt_label());
  sched->addTask(task, level->eachPatch(), sharedState_->allMaterials());
}
```

A task is defined for the computing the stable timestep and uses the function pointer, computeStable-Timestep defined later in Poisson1.cc. This requires data from the New DataWarehouse, and the next timestep size is computed and stored.

```
void
Poisson1::scheduleTimeAdvance( const LevelP& level, SchedulerP& sched)
{
  Task* task = scinew Task("Poisson1::timeAdvance", this, &Poisson1::timeAdvance);

  task->requires(Task::OldDW, phi_label, Ghost::AroundNodes, 1);
  task->computes(phi_label);
  task->computes(residual_label);
  sched->addTask(task, level->eachPatch(), sharedState_->allMaterials());
}
```

The timeAdvance function is the main function that describes the computational algorithm. For simple examples, the entire algorithm is usually defined by one task with a small set of data dependencies. However, for more complicated algorithms, it is best to break the algorithm down into a set of tasks with each task describing its own set of data dependencies.

For this example, a single task is described and the timeAdvance function pointer is specified. Data from the previous timestep (OldDW) is required for the current timestep. For a simple seven (7) point stencil, only one level of ghost cells is required. The algorithm is set up for nodal values, the ghost cells are specified by the the Ghost::AroundNodes syntax. The task computes both the new data values for phi and a residual.

3.1.2 Description of Computational Functions

```
void Poisson1::computeStableTimestep(const ProcessorGroup* pg, const PatchSubset* /*
    patches*/, const MaterialSubset* /*matls*/, DataWarehouse*, DataWarehouse* new_dw)
{
  if(pg->myrank() == 0){
    sum_vartype residual;
    new_dw->get(residual, residual_label);
    cerr << "Residual=" << residual << '\n';
  }
  new_dw->put(delt_vartype(delt_), sharedState_->get_delt_label());
}
```

In this particular example, no timestep is actually computed, instead the original timestep specified in the input file is used and stored in the data warehouse (new_dw→put(delt_vartype(delt_) , sharedState_→get_delt_label())). The residual computed in the main timeAdvance function is retrieved from the data warehouse and printed out to standard error for only the processor with a rank of 0.

```
void Poisson1::initialize(const ProcessorGroup*, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse* /*old_dw*/, DataWarehouse* new_dw)
{
  int matl = 0;
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);
```

The node centered variable (NCVariable<double>phi) has space reserved in the DataWarehouse for the given patch and the given material (int matl = 0;). The phi variable is initialized to 0 for every grid node on the patch.

```
    NCVariable<double> phi;
    new_dw->allocateAndPut(phi, phi_label, matl, patch);
    phi.initialize(0.);
```

The boundary faces on the xminus face of the computational domain are specified and set to a value of 1. All other boundary values are set to a value of 0 as well as the internal nodes via the phi.initialize(0.) construct. VAANGO provides helper functions for determining which nodes are on the boundaries. In addition, there are convenient looping constructs such as NodeIterator that alleviate the need to specify triply nested loops for visiting each node in the domain.

```
    if(patch->getBCType(Patch::xminus) != Patch::Neighbor){
      IntVector l,h;
      patch->getFaceNodes(Patch::xminus, 0, l, h);

      for(NodeIterator iter(l,h); !iter.done(); iter++){
        phi[*iter]=1;
      }
    }
    new_dw->put(sum_vartype(-1), residual_label);
  }
}
```

The initial residual value of -1 is stored at the beginning of the simulation.

The main computational algorithm is defined in the timeAdvance function. The overall algorithm is based on a simple Jacobi iteration step.

```
void Poisson1::timeAdvance(const ProcessorGroup*, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw)
{
  int matl = 0;
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);
    constNCVariable<double> phi;
```

Data from the previous timestep is retrieved from the data warehouse and copied to the current timestep's phi variable (newphi).

```
    old_dw->get(phi, phi_label, matl, patch, Ghost::AroundNodes, 1);
    NCVariable<double> newphi;

    new_dw->allocateAndPut(newphi, phi_label, matl, patch);
    newphi.copyPatch(phi, newphi.getLowIndex(), newphi.getHighIndex());
```

The indices for the patch are obtained and altered depending on whether or not the patch's internal boundaries are on the coincident with the grid domain. If the patch boundaries are the same as the grid domain, the boundary values are not overwritten since the lower and upper indices are modified to only specify internal nodal grid points.

```
    double residual=0;
    IntVector l = patch->getNodeLowIndex__New();
    IntVector h = patch->getNodeHighIndex__New();

    l += IntVector(patch->getBCType(Patch::xminus) == Patch::Neighbor?0:1,
    patch->getBCType(Patch::yminus) == Patch::Neighbor?0:1,
    patch->getBCType(Patch::zminus) == Patch::Neighbor?0:1);
    h -= IntVector(patch->getBCType(Patch::xplus)  == Patch::Neighbor?0:1,
    patch->getBCType(Patch::yplus)  == Patch::Neighbor?0:1,
    patch->getBCType(Patch::zplus)  == Patch::Neighbor?0:1);
```

The Jacobi iteration step is applied at each internal grid node. The residual is computed based on the old and new values and stored as a reduction variable (sum_vartype) in the data warehouse.

```
      //_____
      //  Stencil
      for(NodeIterator iter(l, h);!iter.done(); iter++){
        IntVector n = *iter;

        newphi[n]=(1./6)*(
        phi[n+IntVector(1,0,0)] + phi[n+IntVector(-1,0,0)] +
        phi[n+IntVector(0,1,0)] + phi[n+IntVector(0,-1,0)] +
        phi[n+IntVector(0,0,1)] + phi[n+IntVector(0,0,-1)]);

        double diff = newphi[n] - phi[n];
        residual += diff * diff;
      }
      new_dw->put(sum_vartype(residual), residual_label);
    }
  }
```

3.1.3 Input file

The input file that is used to run this example is given below and is given in SCIRun/src/Packages/Uintah/StandAlone/inputs/Examples/poisson1.ups . Relevant sections of the input file that can be modified are found in the <Time> section, and the <Grid> section, specifically, the number of patches and the grid resolution.

```
<Uintah_specification>
  <Meta>
    <title>Poisson1 test</title>
  </Meta>
  <SimulationComponent>
    <type> poisson1 </type>
  </SimulationComponent>
  <Time>
    <maxTime>       1.0          </maxTime>
    <initTime>      0.0          </initTime>
    <delt_min>      0.00001  </delt_min>
    <delt_max>      1            </delt_max>
    <max_Timesteps> 100          </max_Timesteps>
    <timestep_multiplier>  1  </timestep_multiplier>
  </Time>
  <DataArchiver>
    <filebase>poisson.uda</filebase>
    <outputTimestepInterval>1</outputTimestepInterval>
    <save label = "phi"/>
    <save label = "residual"/>
    <checkpoint cycle = "2" timestepInterval = "1"/>
  </DataArchiver>
  <Poisson>
    <delt>.01</delt>
    <maxresidual>.01</maxresidual>
  </Poisson>
  <Grid>
    <Level>
      <Box label = "1">
        <lower>     [0,0,0]          </lower>
```

```
        <upper>      [1.0,1.0,1.0]   </upper>
        <resolution>[50,50,50]      </resolution>
        <patches>    [2,1,1]         </patches>
      </Box>
    </Level>
  </Grid>
</Uintah_specification>
```

Running the Poisson1 Example

To run the poisson1.ups example,

```
cd ~/SCIRun/dbg/Packages/Uintah/StandAlone/
```

create a symbolic link to the inputs directory:

```
ln -s ~/SCIRun/src/Packages/Uintah/StandAlone/inputs
```

For a single processor run type the following:

```
vaango inputs/Examples/poisson1.ups
```

For a two processor run, type the following:

```
mpirun -np 2 vaango -mpi inputs/Examples/poisson1.ups
```

Changing the number of patches in the poisson1.ups and the resolution, enables you to run a more refined problem on more processors.

ADVICE: For non-AMR problems, it is advised to have at least the same number of patches as processors. You can always have more patches than processors, but you cannot have fewer patches than processors.

3.2   Poisson2

The next example also solves the Poisson's equation but instead of iterating in time, the subscheduler feature iterates within a given timestep, thus solving the problem in one timestep. The use of the subscheduler is important for implementing algorithms which solve non-linear problems which require iterating on a solution for each timestep.

The majority of the schedule and computational functions are similar to the Poisson1 example and are not repeated. Only the revised code is presented with explanations about the new features of Uintah.

```
void Poisson2::scheduleTimeAdvance( const LevelP& level, SchedulerP& sched)
{
  Task* task = scinew Task("timeAdvance", this, &Poisson2::timeAdvance, level, sched.
      get_rep());
  task->hasSubScheduler();
  task->requires(Task::OldDW, phi_label, Ghost::AroundNodes, 1);
  task->computes(phi_label);
  LoadBalancer* lb = sched->getLoadBalancer();
  const PatchSet* perproc_patches = lb->getPerProcessorPatchSet(level);
  sched->addTask(task, perproc_patches, sharedState_->allMaterials());
}
```

Within this function, the task is specified with two additional arguments, the level and the scheduler sched.get_rep(). The task must also set the flag that a subscheduler will be used within the scheduling of the various tasks. Similar code to the Poisson1 example is used to specify what data is required and computed during the actual task execution. In addition, a loadbalancer component is required to query the patch distribution for each level of the grid. The task is then added to the top level scheduler with the requisite information, i.e. patches and materials.

The actual implementation of the timeAdvance function is also different from the Poisson1 example. The code is specified below with text explaining the use of the subscheduler. The new feature of the

subscheduler shows the creation of a the iterate task within the subscheduler. This task will perform the actual Jacobi iteration for a given timestep.

```
void Poisson2::timeAdvance(const ProcessorGroup* pg, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw, LevelP level
    , Scheduler* sched)
{
```

The subscheduler is instantiated and initialized.

```
SchedulerP subsched = sched->createSubScheduler();
subsched->initialize();
GridP grid = level->getGrid();
```

An iterate task is created and added to the subscheduler. The typical computes and requires are specified for a 7 point stencil used in Jacobi iteration scheme with one layer of ghost cells. The new task is added to the subscheduler. A residual variable is only computed within the subscheduler and not passed back to the main scheduler. This is in contrast to the phi variable which was specified in scheduleTimeAdvance in the computes, as well as being specified in the computes for the subscheduler. Any variables that are only computed and used in an iterative step of an algorithm do not need to be added to the dependency specification for the top level task.

```
// Create the tasks
Task* task = scinew Task("iterate", this, &Poisson2::iterate);
task->requires(Task::OldDW, phi_label, Ghost::AroundNodes, 1);
task->computes(phi_label);
task->computes(residual_label);
subsched->addTask(task, level->eachPatch(), sharedState_->allMaterials());
```

The subscheduler has its own data wharehouse that is separate from the top level's scheduler's data warehouse. This data warehouse must be initialized and any data from the top level's data warehouse must be passed to the subscheduler's version. This data resides in the data warehouse position NewDW.

```
// Compile the scheduler
subsched->advanceDataWarehouse(grid);
subsched->compile();

int count = 0;
double residual;
subsched->get_dw(1)->transferFrom(old_dw, phi_label, patches, matls);
```

Within each iteration, the following must occur for the subscheduler: the data warehouse's new data must be moved to the OldDW position, since any new values will be stored in NewDW and the old values cannot be overwritten. The OldDW is referred to in the subscheduler via the subsched->get_dw(0) and the NewDW is referred to in the subscheduler via subsched->get_dw(1). Once the iteration is deemed to have met the tolerance, the data from the subscheduler is transferred to the scheduler's data warehouse.

```
// Iterate
do {
  subsched->advanceDataWarehouse(grid);
  subsched->get_dw(0)->setScrubbing(DataWarehouse::ScrubComplete);
  subsched->get_dw(1)->setScrubbing(DataWarehouse::ScrubNonPermanent);
  subsched->execute();

  sum_vartype residual_var;
  subsched->get_dw(1)->get(residual_var, residual_label);
  residual = residual_var;

  if(pg->myrank() == 0)
    cerr << "Iteration " << count++ << ", residual=" << residual << '\n';
} while(residual > maxresidual_);

new_dw->transferFrom(subsched->get_dw(1), phi_label, patches, matls);
}
```

The iteration cycle is identical to Poisson1's timeAdvance algorithm using Jacobi iteration with a 7 point stencil. Refer to the discussion about the algorithm implementation in the Poisson1 description.

```
void Poisson2::iterate(const ProcessorGroup*, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw)
{
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);
    for(int m = 0;m<matls->size();m++){
      int matl = matls->get(m);
      constNCVariable<double> phi;
      old_dw->get(phi, phi_label, matl, patch, Ghost::AroundNodes, 1);
      NCVariable<double> newphi;
      new_dw->allocateAndPut(newphi, phi_label, matl, patch);
      newphi.copyPatch(phi, newphi.getLow(), newphi.getHigh());
      double residual=0;
      IntVector l = patch->getNodeLowIndex__New();
      IntVector h = patch->getNodeHighIndex__New();
      l += IntVector(patch->getBCType(Patch::xminus) == Patch::Neighbor?0:1,
      patch->getBCType(Patch::yminus) == Patch::Neighbor?0:1,
      patch->getBCType(Patch::zminus) == Patch::Neighbor?0:1);
      h -= IntVector(patch->getBCType(Patch::xplus) == Patch::Neighbor?0:1,
      patch->getBCType(Patch::yplus) == Patch::Neighbor?0:1,
      patch->getBCType(Patch::zplus) == Patch::Neighbor?0:1);
      for(NodeIterator iter(l, h);!iter.done(); iter++){
        newphi[*iter]=(1./6)*(
        phi[*iter+IntVector(1,0,0)]+phi[*iter+IntVector(-1,0,0)]+
        phi[*iter+IntVector(0,1,0)]+phi[*iter+IntVector(0,-1,0)]+
        phi[*iter+IntVector(0,0,1)]+phi[*iter+IntVector(0,0,-1)]);
        double diff = newphi[*iter]-phi[*iter];
        residual += diff*diff;
      }
      new_dw->put(sum_vartype(residual), residual_label);
    }
  }
}
```

The input file src/StandAlone/inputs/Examples/poisson2.ups is very similar to the poisson1.ups file shown above. The only additional tag that is used is the <maxresidual > specifying the tolerance within the iteration performed in the subscheduler.

To run the poisson2 input file execute the following in the dbg build StandAlone directory:

```
vaango inputs/Examples/poisson2.ups
```

3.3   Burger

In this example, the inviscid Burger's equation is solved in three dimensions:

$$\frac{du}{dt} = -u\frac{du}{dx} \tag{3.1}$$

with the initial conditions:

$$u = \sin(\pi x) + \sin(2\pi y) + \sin(3\pi z) \tag{3.2}$$

using Euler's method to advance in time. The majority of the code is very similar to the Poisson1 example code with the differences shown below.

The initialization of the grid values for the unknown variable, u, is done at every grid node using the NodeIterator construct. The x,y,z values for a given grid node is determined using the function, patch->getNodePosition(n), where n is the nodal index in i,j,k space.

```
void Burger::initialize(const ProcessorGroup*, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse*, DataWarehouse* new_dw)
{
```

```
    int matl = 0;
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);

    NCVariable<double> u;
    new_dw->allocateAndPut(u, u_label, matl, patch);

    //Initialize
    // u = sin( pi*x ) + sin( pi*2*y ) + sin(pi*3z )
    IntVector l = patch->getNodeLowIndex__New();
    IntVector h = patch->getNodeHighIndex__New();

    for( NodeIterator iter=patch->getNodeIterator__New(); !iter.done(); iter++ ){
      IntVector n = *iter;
      Point p = patch->nodePosition(n);
      u[n] = sin( p.x() * 3.14159265358 ) + sin( p.y() * 2*3.14159265358)  +  sin( p.z
          () * 3*3.14159265358);
    }
  }
}
```

The timeAdvance function is quite similar to the Poisson1's timeAdvance routine. The relevant differences are only shown.

```
void Burger::timeAdvance(const ProcessorGroup*, const PatchSubset* patches, const
    MaterialSubset* matls, DataWarehouse* old_dw, DataWarehouse* new_dw)
{
  int matl = 0;
  //Loop for all patches on this processor
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);
    . . . . .
```

The grid spacing and timestep values are stored.

```
    // dt, dx
    Vector dx = patch->getLevel()->dCell();
    delt_vartype dt;
    old_dw->get(dt, sharedState_->get_delt_label());
    . . . . .
```

Refer to the description in Poisson1 about the specification of the NodeIterator limits. The Euler algorithm is applied to solve the ordinary differential equation in time.

```
    //Iterate through all the nodes
    for(NodeIterator iter(l, h);!iter.done(); iter++){
      IntVector n = *iter;
      double dudx = (u[n+IntVector(1,0,0)] - u[n-IntVector(1,0,0)]) /(2.0 * dx.x());
      double dudy = (u[n+IntVector(0,1,0)] - u[n-IntVector(0,1,0)]) /(2.0 * dx.y());
      double dudz = (u[n+IntVector(0,0,1)] - u[n-IntVector(0,0,1)]) /(2.0 * dx.z());
      double du = - u[n] * dt * (dudx + dudy + dudz);
      new_u[n]= u[n] + du;
    }
```

Zero flux Neumann boundary conditions are applied to the node points on each of the grid faces.

```
    //--------------------------------
    // Boundary conditions: Neumann
    // Iterate over the faces encompassing the domain
    vector<Patch::FaceType>::const_iterator iter;
    vector<Patch::FaceType> bf;
    patch->getBoundaryFaces(bf);
    for (iter  = bf.begin(); iter != bf.end(); ++iter){
      Patch::FaceType face = *iter;

      IntVector axes = patch->faceAxes(face);
      int P_dir = axes[0]; // find the principal dir of that face

      IntVector offset(0,0,0);
      if (face == Patch::xminus || face == Patch::yminus || face == Patch::zminus){
```

```
        offset[P_dir] += 1;
      }
      if (face == Patch::xplus || face == Patch::yplus || face == Patch::zplus){
        offset[P_dir] -= 1;
      }

      Patch::FaceIteratorType FN = Patch::FaceNodes;
      for (CellIterator iter = patch->getFaceIterator__New(face,FN);!iter.done(); iter
          ++){
        IntVector n = *iter;
        new_u[n] = new_u[n + offset];
      }
    }
  }
}
```

The input file for the Burger (`burger.ups`) problem is very similar to the `poisson1.ups` with the addition, that the timestep increment used in the timeAdvance is quite small, 1.e-4 for stability reasons.

To run the Burger input file execute the following in the dbg build StandAlone directory:

```
vaango inputs/Examples/burger.ups
```
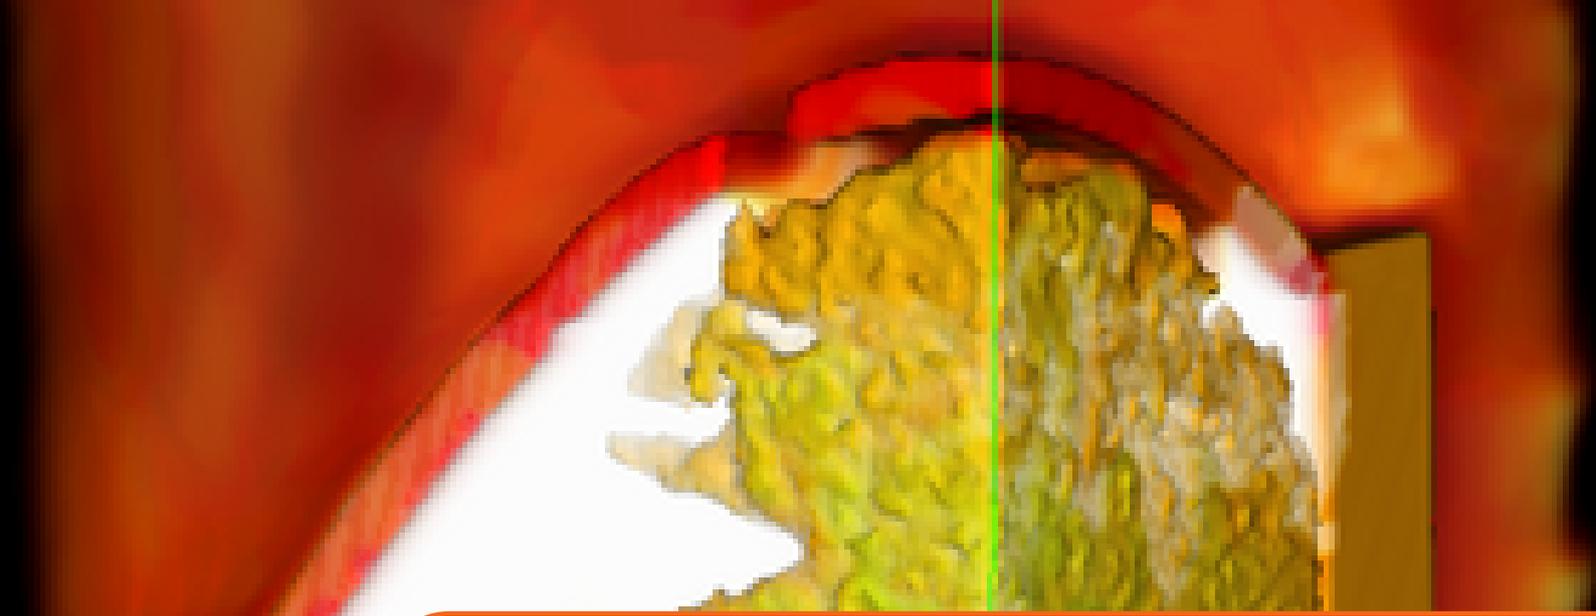
4 — The MPM component

The MPM component solves the momentum equations

$$\nabla \cdot \boldsymbol{\sigma} + \rho \mathbf{b} = \rho \dot{\mathbf{v}} \tag{4.1}$$

using an updated Lagrangian formulation. The momentum solve for solid materials is complicated by the fact that the equations need material constitutive models for closure. These material constitutive models vary significantly between materials and contribute a large fraction of the computational cost of a simulation.

In this chapter we discuss the algorithm used in Vaango to solve the momentum equations using MPM . The implementation follows the approach discussed in the previous chapter.

4.1 The MPM algorithm

The momentum equation is solved using the MPM algorithm while forward Euler time-stepping is use to integrate time derivatives. The pseudocode of the overall algorithm is given below. The main quantities of interest are:

- $t_{\max}$: The maximum time until which the simulation is to run.
- $t, \Delta t$: The current time ($t = t_n$) and the time step.
- $\boldsymbol{h}_g$: The grid spacing vector.
- m_p : The particle mass.
- V_p^n, V_p^{n+1} : The particle volume at $t = t_n$ and $t = t_{n+1}$.
- $\mathbf{x}_p^n, \mathbf{x}_p^{n+1}$: The particle position at $t = t_n$ and $t = t_{n+1}$.
- $\mathbf{u}_p^n, \mathbf{u}_p^{n+1}$: The particle displacement at $t = t_n$ and $t = t_{n+1}$.
- $\mathbf{v}_p^n, \mathbf{v}_p^{n+1}$: The particle velocity at $t = t_n$ and $t = t_{n+1}$.
- $\boldsymbol{\sigma}_p^n, \boldsymbol{\sigma}_p^{n+1}$: The particle Cauchy stress at time $t = t_n$ and $t = t_{n+1}$.
- $\boldsymbol{F}_p^n, \boldsymbol{F}_p^{n+1}$: The particle deformation gradient at time $t = t_n$ and $t = t_{n+1}$.

Algorithm 1 The algorithm

1: **procedure** RUN(inputUPSFile)
2:     $t_{\max}$, $\boldsymbol{h}_g$, `xmlProblemSpec`, `grid`, `globalState` $\leftarrow$ READINPUTUPSFILE(inputUPSFile)   ▷*Parse*
        ↪ *the input XML file (¡filename¿.ups), create the background grid, and*
        ↪ *set up a* SIMULATIONSTATE.

3:      mpmFlags, prescribedDefGrad, particleBC, contactModel, constitutiveModel,
        ↪ defGradComputer, damageModel ← PROBLEMSETUP(xmlProblemSpec, grid,
        ↪ globalState)   ▷*Set up flags, the constitutive model, and the deformation gradient*
        ↪ *algorithm based on data in input file.*

4:      $t \leftarrow 0, n \leftarrow 0$

5:      $\mathbf{x}_p^n, \mathbf{u}_p^n, m_p, V_p^n, \mathbf{v}_p^n, \boldsymbol{\sigma}_p^n, \boldsymbol{F}_p^n \leftarrow$ INITIALIZE(xmlProblemSpec)   ▷*Find the grid size and initialize*
        ↪ *particle variables based on geometry and other information in the input file.*

6:      isSuccess ← FALSE

7:      **repeat**

8:          $\Delta t \leftarrow$ COMPUTESTABLETIMESTEP($\boldsymbol{h}_g, \mathbf{v}_p$)   ▷*Find a stable time increment based on*
            ↪ *grid size and velocity*

9:          $t \leftarrow t + \Delta t, n \leftarrow n + 1$ ▷*Update the time*

10:         isSuccess, $\mathbf{x}_p^{n+1}, \mathbf{u}_p^{n+1}, v_p^{n+1}, \mathbf{v}_p^{n+1}, \boldsymbol{\sigma}_p^{n+1}, \boldsymbol{F}_p^{n+1} \leftarrow$ TIMEADVANCE($\boldsymbol{h}_g, \mathbf{x}_p^n, \mathbf{u}_p^n, m_p, V_p^n,$
            ↪ $\mathbf{v}_p^n, \boldsymbol{\sigma}_p^n, \boldsymbol{F}_p^n$)   ▷*Compute updated quantities*

11:         OUTPUTDATA($\mathbf{x}_p^{n+1}, \mathbf{u}_p^{n+1}, V_p^{n+1}, \mathbf{v}_p^{n+1}, \boldsymbol{\sigma}_p^{n+1}, \boldsymbol{F}_p^{n+1}$)   ▷*Save the solution*

12:         $n \leftarrow n + 1$

13:     **until** $t \geq t_{\max}$

14:     **return** isSuccess

15: **end procedure**

4.2 Reading the input file

The process used to read the input file is identical to that discussed earlier in this manual.

4.3 Problem setup

The overall structure of the problem setup code is given below. Details can be found in the code.

Algorithm 2 Problem setup

Require: xmlProblemSpec, grid, globalState

1: **procedure** PROBLEMSETUP(xmlProblemSpec, grid, globalState)

2:      flags ← READMPMFLAGS(xmlProblemSpec)   ▷*Read the option flags that determine*
        ↪ *the details of the MPM algorithm to be used in the simulation.*

3:      **if** flags.prescribeDeformation = TRUE **then**

4:          prescribedDefGrad ← READPRESCRIBEDDEFORMATIONS(flags.prescribedFileName)

5:      **end if**

6:      particleBC ← CREATEMPMPHYSICALBC(xmlProblemSpec, grid, flags)   ▷*Create the model*
        ↪ *used to apply pressures and forces directly to particles.*

7:      contactModel ← CREATECONTACTMODEL(xmlProblemSpec, grid, flags, globalState) ▷*Create*
        ↪ *the contact algorithm model used to compute interactions between objects.*

8:      constitutiveModel ← CREATECONSTITUTIVEMODELS(xmlProblemSpec, grid, flags,
        ↪ globalState)   ▷*Create the constitutive models that are needed*
        ↪ *for the simulation.*

9:      defGradComputer ← CREATEDEFORMATIONGRADIENTCOMPUTER(flags, globalState) ▷*Create*
        ↪ *the model that will be used to compute velocity and*
        ↪ *deformation gradients.*

10:     **if** flags.doBasicDamage = TRUE **then**

11:         damageModel ← CREATEBASICDAMAGEMODEL(flags, globalState)

12:     **end if**

13:     **return** flags, prescribedDefGrad, particleBC, contactModel, constitutiveModel,

       ↪ `damageModel`, `defGradComputer`

14: **end procedure**

4.4 Initialization

An outline of the initialization process is described below. Specific details have been discussed in earlier reports. The new quantities introduced in this section are

- n_p : The number of particles used to discretize a body.
- $\mathbf{b}_p^n, \mathbf{b}_p^{n+1}$: The particle body force acceleration at $t = t_n$ and $t = t_{n+1}$.
- D_p^n, D_p^{n+1} : The particle damage parameter at $t = t_n$ and $t = t_{n+1}$.
- $\mathbf{f}_p^{\text{ext},n}, \mathbf{f}_p^{\text{ext},n+1}$: The particle external force at $t = t_n$ and $t = t_{n+1}$.

Algorithm 3 Initialization

Require: `xmlProblemSpec`, `defGradComputer`, `constitutiveModel`, `damageModel`, `particleBC`,
       ↪ `mpmFlags` `materialList`,

1: **procedure** INITIALIZE
2:    **for** `matl` **in** `materialList` **do**
3:       $n_p[\texttt{matl}], \mathbf{x}_p^{\text{o}}[\texttt{matl}], \mathbf{u}_p^{\text{o}}[\texttt{matl}], m_p[\texttt{matl}], V_p^{\text{o}}[\texttt{matl}], \mathbf{v}_p^{\text{o}}[\texttt{matl}], \mathbf{b}_p^{\text{o}}[\texttt{matl}],$
         ↪ $\mathbf{f}_p^{\text{ext},\text{o}}[\texttt{matl}] \leftarrow$ `matl`.CREATEPARTICLES()
4:       $F_p^{\text{o}}[\texttt{matl}] \leftarrow$ `defGradComputer`.INITIALIZE(`matl`)
5:       $\boldsymbol{\sigma}_p^{\text{o}}[\texttt{matl}] \leftarrow$ `constitutiveModel`.INITIALIZE(`matl`)
6:       $D_p^{\text{o}}[\texttt{matl}] \leftarrow$ `damageModel`.INITIALIZE(`matl`)
7:    **end for**
8:    **if** `mpmFlags`.initializeStressWithBodyForce = TRUE **then**
9:       $\mathbf{b}_p^{\text{o}} \leftarrow$ INITIALIZEBODYFORCE()
10:      $\boldsymbol{\sigma}_p^{\text{o}}, F_p^{\text{o}} \leftarrow$ INITIALIZESTRESSANDDEFGRADFROMBODYFORCE()
11:    **end if**
12:    **if** `mpmFlags`.applyParticleBCs = TRUE **then**
13:       $\mathbf{f}_p^{\text{ext},\text{o}} \leftarrow$ `particleBC`.INITIALIZEPRESSUREBCS()
14:    **end if**
15:    **return** $n_p, \mathbf{x}_p^{\text{o}}, \mathbf{u}_p^{\text{o}}, m_p, V_p^{\text{o}}, \mathbf{v}_p^{\text{o}}, \mathbf{b}_p^{\text{o}}, \mathbf{f}_p^{\text{ext},\text{o}}, F_p^{\text{o}}, \boldsymbol{\sigma}_p^{\text{o}}, D_p^{\text{o}}$
16: **end procedure**

4.5 Time advance

The operations performed during a timestep are shown in the pseudocode below.

Algorithm 4 The MPM time advance algorithm

1: **procedure** TIMEADVANCE($\boldsymbol{h}_g, x_p^n, u_p^n, m_p, V_p^n, \mathbf{v}_p^n, \mathbf{f}_p^{\text{ext},n}, \boldsymbol{d}_p^n$)
2:    $\mathbf{b}_p^n \leftarrow$ COMPUTEPARTICLEBODYFORCE()   ▷*Compute the body force term*
3:    $\mathbf{f}_p^{\text{ext},n+1} \leftarrow$ APPLYEXTERNALLOADS()   ▷*Apply external loads to the particles*
4:    $m_g, V_g, \mathbf{v}_g, \mathbf{b}_g, \mathbf{f}_g^{\text{ext}} \leftarrow$ INTERPOLATEPARTICLESTOGRID() ▷*Interpolate particle data to the grid*
5:    EXCHANGEMOMENTUMINTERPOLATED()   ▷*Exchange momentum between bodies on grid.*
        ↪ *Not discussed in this report.*
6:    $\mathbf{f}_g^{\text{int}}, \boldsymbol{\sigma}_g, \mathbf{v}_g \leftarrow$ COMPUTEINTERNALFORCE()   ▷*Compute the internal force at the grid nodes*
7:    $\mathbf{v}_g^{\star}, \mathbf{a}_g \leftarrow$ COMPUTEANDINTEGRATEACCELERATION()   ▷*Compute the grid velocity*
        ↪ *and grid acceleration*
8:    EXCHANGEMOMENTUMINTEGRATED()   ▷*Exchange momentum between bodies on grid*
        ↪ *using integrated values. Not discussed in this report.*

9:    $\mathbf{v}_g^{\star}, \mathbf{a}_g \leftarrow$ SETGRIDBOUNDARYCONDITIONS()   ▷*Update the grid velocity and grid*
                          ↪ *acceleration using the BCs*

10:    $\mathbf{l}_p^n, \mathbf{F}_p^{n+1}, V_p^{n+1} \leftarrow$ COMPUTEDEFORMATIONGRADIENT()   ▷*Compute the velocity gradient*
                          ↪ *and the deformation gradient*

11:    $\boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\eta}_p^{n+1} \leftarrow$ COMPUTESTRESSTENSOR()   ▷*Compute the updated stress and*
                          ↪ *internal variables (if any)*

12:    $\boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\eta}_p^{n+1}, \chi_p^{n+1}, D_p^{n+1} \leftarrow$ COMPUTEBASICDAMAGE()   ▷*Compute the damage parameter*
                          ↪ *and update the stress and internal variables*

13:    $\chi_p^{n+1}, D_p^{n+1} \leftarrow$ UPDATEEROSIONPARAMETER()   ▷*Update the indicator variable that is used*
                          ↪ *to delete particles at the end of a time step*

14:    $V_p^{n+1}, \mathbf{u}_p^{n+1}, \mathbf{v}_p^{n+1}, \mathbf{x}_p^{n+1}, m_p, \boldsymbol{h}_p^{n+1} \leftarrow$ INTERPOLATETOPARTICLESANDUPDATE()   ▷*Update the*
                          ↪ *particle variables after interpolating grid quantities to particles*

15: **end procedure**

The algorithms used for the above operations are discussed next.

4.5.1 Computing the body force

The body force consists of a gravitational term and, optionally, centrifugal and coriolis terms that are needed for simulations inside a rotating frame such as a centrifuge.

Algorithm 5 Computing the body force on particles

Require: $\mathbf{x}_p^n, \mathbf{v}_p^n$, materialList, particleList, mpmFlags

1: **procedure** COMPUTEPARTICLEBODYFORCE

2:    **for** matl **in** materialList **do**

3:       **if** mpmFlags.rotatingCoordSystem = TRUE **then**

4:          $\mathbf{g} \leftarrow$ mpmFlags.gravityAcceleration

5:          $\mathbf{b}_p^n[\text{matl}] \leftarrow \mathbf{g}$

6:       **else**

7:          **for** part **in** particleList **do**

8:             $\mathbf{g} \leftarrow$ mpmFlags.gravityAcceleration

9:             $\mathbf{x}_{rc} \leftarrow$ mpmFlags.coordRotationCenter

10:            $\mathbf{z}_r \leftarrow$ mpmFlags.coordRotationAxis

11:            $w \leftarrow$ mpmFlags.coordRotationSpeed

12:            $\boldsymbol{\omega} \leftarrow w\mathbf{z}_r$ ▷*Compute angular velocity vector*

13:            $\mathbf{a}_{\text{corolis}} \leftarrow 2\boldsymbol{\omega} \times \mathbf{v}_p^n[\text{matl}, \text{part}]$ ▷*Compute Coriolis acceleration*

14:            $\mathbf{r} \leftarrow \mathbf{x}_p^n[\text{matl}, \text{part}] - \mathbf{x}_{rc}$

15:            $\mathbf{a}_{\text{centrifugal}} \leftarrow \boldsymbol{\omega} \times \boldsymbol{\omega} \times \mathbf{r}$ ▷*Compute the centrifugal body force acceleration*

16:            $\mathbf{b}_p^n[\text{matl}, \text{part}] \leftarrow \mathbf{g} - \mathbf{a}_{\text{centrifugal}} - \mathbf{a}_{\text{coriolis}}$ ▷*Compute the body force acceleration*

17:          **end for**

18:       **end if**

19:    **end for**

20:    **return** $\mathbf{b}_p^n$

21: **end procedure**

4.5.2 Applying external loads

Note that the updated deformation gradient has not been computed yet at this stage and the particle force is applied based on the deformation gradient at the beginning of the timestep. The new quantities introduced in this section are:

- $\boldsymbol{h}_p^n$: The particle size matrix at time $t = t_n$.

Algorithm 6 Applying external loads to particles

Require: $t_{n+1}, \mathbf{x}_p^n, h_p^n, \mathbf{u}_p^n, \mathbf{f}_p^{\text{ext},n}, F_p^n$, materialList, particleList, mpmFlags, particleBC

1: **procedure** APPLYEXTERNALLOADS
2:      $f_p \leftarrow 0$
3:      **if** mpmFlags.useLoadCurves = TRUE **then**
4:          $f_p \leftarrow$ particleBC.COMPUTEFORCEPERPARTICLE(t^{n+1})   ▷*Compute the force per particle*
            ↪ *due to the applied pressure*
5:      **end if**
6:      **for** matl **in** materialList **do**
7:          **if** mpmFlags.useLoadCurves = TRUE **then**
8:             **for** part **in** particleList **do**
9:                 $\mathbf{f}_p^{\text{ext},n+1}$[matl,part] $\leftarrow$ particleBC.GETFORCEVECTOR($t_{n+1}, \mathbf{x}_p^n, h_p^n, \mathbf{u}_p^n$,
                   ↪ f_p, F_p^n)   ▷*Compute the applied force vector at each particle*
10:             **end for**
11:          **else**
12:             $\mathbf{f}_p^{\text{ext},n+1}$[matl] $\leftarrow \mathbf{f}_p^{\text{ext},n}$[matl]
13:          **end if**
14:      **end for**
15:      **return** $\mathbf{f}_p^{\text{ext},n+1}$
16: **end procedure**

4.5.3 Interpolating particles to grid

The grid quantities computed during this procedure and not stored for the next timestep except for the purpose of visualization. The new quantities introduced in this section are

- m_g : The mass at a grid node.
- V_g : The volume at a grid node.
- $\mathbf{v}_g$: The velocity at a grid node.
- $\mathbf{f}_g^{\text{ext}}$: The external force at a grid node.
- $\mathbf{b}_g$: The body force at a grid node.

Algorithm 7 Interpolating particle data to background grid

Require: $m_p, V_p^n, \mathbf{x}_p^n, h_p^n, \mathbf{b}_p^n, \mathbf{f}_p^{\text{ext},n+1}, F_p^n$, materialList, particleList, gridNodeList mpmFlags, particleBC

1: **procedure** INTERPOLATEPARTICLESTOGRID
2:      interpolator $\leftarrow$ CREATEINTERPOLATOR(mpmFlags)   ▷*Create the interpolator*
         ↪ *and find number of grid nodes that can affect a particle*
3:      **for** matl **in** materialList **do**
4:          **for** part **in** particleList **do**
5:             $n_{gp}, S_{gp} \leftarrow$ interpolator.FINDCELLSANDWEIGHTS($\mathbf{x}_p^n, h_p^n, F_p^n$)   ▷*Find the node*
                ↪ *indices of the cells affecting the particle and the interpolation weights*
6:             $\mathbf{p}_p \leftarrow m_p$[matl][part] $\mathbf{v}_p^n$[matl][part]   ▷*Compute particle momentum*
7:             **for** node **in** n_{gp} **do**
8:                 m_g[matl][node] $\leftarrow m_g$[matl][node] $+ m_p$[matl][part] S_{gp}[node]
9:                 V_g[matl][node] $\leftarrow V_g$[matl][node] $+ V_p^n$[matl][part] S_{gp}[node]
10:                 $\mathbf{v}_g$[matl][node] $\leftarrow \mathbf{v}_g$[matl][node] $+ \mathbf{p}_p S_{gp}$[node]
11:                 $\mathbf{f}_g^{\text{ext}}$[matl][node] $\leftarrow \mathbf{f}_g^{\text{ext}}$[matl][node] $+ \mathbf{f}_p^{\text{ext},n+1}$[matl][part] S_{gp}[node]
12:                 $\mathbf{b}_g$[node] $\leftarrow \mathbf{b}_g$[node] $+ m_p$[matl][part] $\mathbf{b}_p^n$[matl][part] S_{gp}[node]
13:             **end for**
14:          **end for**

15:         **for** node **in** gridNodeList **do**
16:             $\mathbf{v}_g$[matl][node] ← $\mathbf{v}_g$[matl][node]/m_g[matl][node]
17:         **end for**
18:         $\mathbf{v}_g$[matl] ← APPLYSYMMETRYVELOCITYBC($\mathbf{v}_g$[matl])   ▷*Apply any symmetry*
                ↪ *velocity BCs that may be applicable*
19:     **end for**
20:     **return** m_g, V_g, $\mathbf{v}_g$, $\mathbf{b}_g$, $\mathbf{f}_g^{\text{ext}}$
21: **end procedure**

4.5.4 Exchanging momentum using interpolated grid values

The exchange of momentum is carried out using a contact model. Details can be found in the VAANGO Developers Manual.

4.5.5 Computing the internal force

This procedure computes the internal force at the grid nodes. The new quantities introduced in this section are

- n_{gp} : The number of grid nodes that are used to interpolate from particle to grid.
- S_{gp} : The nodal interpolation function evaluated at a particle
- $\boldsymbol{G}_{gp}$: The gradient of the nodal interpolation function evaluated at a particle
- $\boldsymbol{\sigma}_v$: A volume weighted grid node stress.
- $\mathbf{f}_g^{\text{int}}$: The internal force at a grid node.

Algorithm 8 Computing the internal force

Require: h_g, V_g, V_p^n, $\mathbf{x}_p^n$, h_p^n, $\boldsymbol{\sigma}_p^n$, F_p^n, materialList, particleList, gridNodeList mpmFlags
 1: **procedure** COMPUTEINTERNALFORCE
 2:     interpolator ← CREATEINTERPOLATOR(mpmFlags)   ▷*Create the interpolator and*
                ↪ *find number of grid nodes that can affect a particle*
 3:     **for** matl **in** materialList **do**
 4:         **for** part **in** particleList **do**
 5:             n_{gp}, S_{gp}, $\mathbf{G}_{gp}$ ←
                    ↪ interpolator.FINDCELLSANDWEIGHTSANDSHAPEDERVATIVES($\mathbf{x}_p^n$, h_p^n, F_p^n)
                    ↪   ▷*Find the node indices of the cells affecting the particle and*
                    ↪ *the interpolation weights and gradients*
 6:             $\boldsymbol{\sigma}_v$ ← V_p[matl][part] $\boldsymbol{\sigma}_p^n$[matl][part]
 7:             **for** node **in** n_{gp} **do**
 8:                 $\mathbf{f}_g^{\text{int}}$[matl][node] ← $\mathbf{f}_g^{\text{int}}$[matl][node] - ($\mathbf{G}_{gp}$[node]/$h_g$) ·$\boldsymbol{\sigma}_p^n$[matl][part] V_p^n[part]
 9:                 $\boldsymbol{\sigma}_g$[matl][node] ← $\boldsymbol{\sigma}_g$[matl][node] + $\boldsymbol{\sigma}_v$ S_{gp}[node]
10:             **end for**
11:         **end for**
12:         **for** node **in** gridNodeList **do**
13:             $\boldsymbol{\sigma}_g$[matl][node] ← $\boldsymbol{\sigma}_g$[matl][node]/V_g[matl][node]
14:         **end for**
15:         $\mathbf{v}_g$[matl] ← APPLYSYMMETRYTRACTIONBC()   ▷*Apply any symmetry tractions BCs*
                    ↪ *that may be applicable*
16:     **end for**
17:     **return** $\mathbf{f}_g^{\text{int}}$, $\boldsymbol{\sigma}_g$, $\mathbf{v}_g$
18: **end procedure**

4.5.6 Computing and integrating the acceleration

This procedure computes the accelerations at the grid nodes and integrates the grid accelerations using forward Euler to compute grid velocities. The new quantities introduced in this section are

- $\mathbf{a}_g$: The grid accelerations.
- $\mathbf{v}_g^\star$: The integrated grid velocities.

Algorithm 9 Computing and integrating the acceleration

Require: $\Delta t, m_g, \mathbf{f}_g^{\text{int}}, \mathbf{f}_g^{\text{ext}}, \mathbf{b}_g, \mathbf{v}_g$, materialList, gridNodeList, mpmFlags

1: **procedure** COMPUTEANDINTEGRATEACCELERATION
2:     **for** matl **in** materialList **do**
3:         **for** node **in** gridNodeList **do**
4:             $\mathbf{a}_g[\text{matl}][\text{node}] \leftarrow (\mathbf{f}_g^{\text{int}}[\text{matl}][\text{node}] + \mathbf{f}_g^{\text{ext}}[\text{matl}][\text{node}] + \mathbf{b}_g[\text{matl}][\text{node}])/m_g[\text{matl}][\text{node}]$
5:             $\mathbf{v}_g^\star \leftarrow \mathbf{v}_g[\text{matl}][\text{node}] + \mathbf{a}_g[\text{matl}][\text{node}] * \Delta t$
6:         **end for**
7:     **end for**
8:     **return** $\mathbf{v}_g^\star, \mathbf{a}_g$
9: **end procedure**

4.5.7 Exchanging momentum using integrated grid values

The exchange of momentum is carried out using a contact model. Details can be found in the VAANGO Developers Manual.

4.5.8 Setting grid boundary conditions

Algorithm 10 Setting grid boundary conditions

Require: $\Delta t, \mathbf{a}_g, \mathbf{v}_g^\star, \mathbf{v}_g$, materialList, gridNodeList, mpmFlags

1: **procedure** SETGRIDBOUNDARYCONDITIONS
2:     **for** matl **in** materialList **do**
3:         $\mathbf{v}_g^\star[\text{matl}] \leftarrow$ APPLYSYMMETRYVELOCITYBC$(\mathbf{v}_g^\star[\text{matl}])$
4:         **for** node **in** gridNodeList **do**
5:             $\mathbf{a}_g[\text{matl}][\text{node}] \leftarrow (\mathbf{v}_g^\star[\text{matl}][\text{node}] - \mathbf{v}_g[\text{matl}][\text{node}])/\Delta t$
6:         **end for**
7:     **end for**
8:     **return** $\mathbf{v}_g^\star, \mathbf{a}_g$
9: **end procedure**

4.5.9 Computing the deformation gradient

The velocity gradient is computed using the integrated grid velocities and then used to compute the deformation gradient. The new quantities introduced in this section are

- ΔF_p^n : The increment of the particle deformation gradient.
- l_p^{n+1} : The particle velocity gradient.
- ρ_{o} : The initial mass density of the material.

Algorithm 11 Computing the velocity gradient and deformation gradient

Require: $\Delta t, \mathbf{x}_p^n, m_p, V_p^n, \mathbf{h}_p^n, \mathbf{v}_p^n, l_p^n, F_p^n, \mathbf{h}_g, \mathbf{v}_g, \mathbf{v}_g^\star, \rho_{\text{o}}$ materialList, gridNodeList, mpmFlags, velGradComputer

1: **procedure** COMPUTEDEFORMATIONGRADIENT
2:     interpolator $\leftarrow$ CREATEINTERPOLATOR(mpmFlags)

3:    **for** `matl` **in** `materialList` **do**
4:        **for** `part` **in** `particleList` **do**
5:            l_p^{n+1}[`matl`,`part`] ← `velGradComputer`.COMPUTEVELGRAD(`interpolator`, h_g, $\mathbf{x}_p^n$[`matl`,`part`],
                ↪ h_p^n[`matl`,`part`], F_p^n[`matl`,`part`], $\mathbf{v}_g^\star$[`matl`])   ▷*Compute the velocity gradient*
6:            F_p^{n+1}[`matl`,`part`], ΔF_p^{n+1} ← COMPUTEDEFORMATIONGRADIENTFROMVELOCITY(l_p^n[`matl`,`part`],
                ↪ l_p^{n+1}[`matl`,`part`], F_p^n[`matl`,`part`])   ▷*Compute the deformation gradient*
7:            V_p^{n+1}[`matl`,`part`] ← m_p[`matl`,`part`]$/\rho_o *$ det(F_p^{n+1}[`matl`,`part`])
8:        **end for**
9:    **end for**
10:   **return** l_p^{n+1}, F_p^{n+1}, V_p^{n+1}
11: **end procedure**

Algorithm 12 Computing the deformation gradient using the velocity gradient

Require: Δt, l_p^{n+1}, F_p^n, `mpmFlags`
1: **procedure** COMPUTEDEFORMATIONGRADIENTFROMVELOCITY
2:      **if** `mpmFlags.defGradAlgorithm` = `"first_order"` **then**
3:          F_p^{n+1}, ΔF_p^{n+1} ← SERIESUPDATECONSTANTVELGRAD(numTerms = 1, Δt, l_p^{n+1}, F_p^n)
4:      **else if** `mpmFlags.defGradAlgorithm` = `"subcycle"` **then**
5:          F_p^{n+1}, ΔF_p^{n+1} ← SUBCYCLEUPDATECONSTANTVELGRAD(Δt, l_p^{n+1}, F_p^n)
6:      **else if** `mpmFlags.defGradAlgorithm` = `"taylor_series"` **then**
7:          F_p^{n+1}, ΔF_p^{n+1} ← SERIESUPDATECONSTANTVELGRAD(numTerms = `mpmFlags`.numTaylorSeriesTerms,
        Δt, l_p^{n+1}, F_p^n)
8:      **else**
9:          F_p^{n+1}, ΔF_p^{n+1} ← CAYLEYUPDATECONSTANTVELGRAD(Δt, l_p^{n+1}, F_p^n)
10:     **end if**
11:     **return** F_p^{n+1}, ΔF_p^{n+1}
12: **end procedure**

4.5.10 Computing the stress tensor

The stress tensor is compute by individual constitutive models. Details of the Arena partially saturated model are given later. The new quantities introduced in this section are

- η_p^n, η_p^{n+1} : The internal variables needed by the constitutive model.

Algorithm 13 Computing the stress tensor

Require: Δt, $\mathbf{x}_p^n$, m_p, V_p^{n+1}, h_p^n, l_p^{n+1}, F_p^{n+1}, σ_p^n, η_p^n, ρ_o, `materialList`, `mpmFlags`, `constitutiveModel`
1: **procedure** COMPUTESTRESSTENSOR
2:      **for** `matl` **in** `materialList` **do**
3:          σ^{n+1}, η_p^{n+1} ← `constitutiveModel`[`matl`].COMPUTESTRESSTENSOR(Δt, $\mathbf{x}_p^n$, m_p, V_p^{n+1}, h_p^n,
                ↪ l_p^{n+1}, F_p^{n+1}, σ_p^n, η_p^n, ρ_o, `mpmFlags`)   ▷*Update the stress and any*
                ↪ *internal variables needed by the constitutive model*
4:      **end for**
5:      **return** σ_p^{n+1}, η_p^{n+1}
6: **end procedure**

4.5.11 Computing the basic damage parameter

The damage parameter is updated and the particle stress is modified in this procedure. The new quantities introduced in this section are

- $\varepsilon_p^{f,n}, \varepsilon_p^{f,n+1}$: The particle strain to failure at $t = T_n$ and $t = T_{n+1}$.
- χ_p^n, χ_p^{n+1} : An indicator function that identifies whether a particle has failed completely.
- $t_p^{\chi,n}, t_p^{\chi,n+1}$: The time to failure of a particle.
- D_p^n, D_p^{n+1} : A particle damage parameter that can be used to modify the stress.

Algorithm 14 Computing the damage parameter

Require: $t^{n+1}, V_p^{n+1}, F_p^{n+1}, \sigma_p^{n+1}, D_p^n, \varepsilon_p^{f,n}, \chi_p^n, t_p^{\chi,n}$, materialList, mpmFlags

1: **procedure** COMPUTEDAMAGE
2:     **for** matl **in** materialList **do**
3:         **for** part **in** particleList **do**
4:             **if** brittleDamage = TRUE **then**
5:                 $\sigma_p^{n+1}, \varepsilon_p^{f,n+1}, \chi_p^{n+1}, t_p^{\chi,n+1}, D_p^{n+1} \leftarrow$ UPDATEDAMAGEANDMODIFYSTRESS($V_p^{n+1}, F_p^{n+1},$
                    $\hookrightarrow \sigma_p^{n+1}, D_p^n, \varepsilon_p^{f,n}, \chi_p^n, t_p^{\chi,n}$) $\triangleright$*Update the damage parameters and stress*
6:             **else**
7:                 $\sigma_p^{n+1}, \varepsilon_p^{f,n+1}, \chi_p^{n+1}, t_p^{\chi,n+1} \leftarrow$ UPDATEFAILEDPARTICLESANDMODIFYSTRESS($V_p^{n+1}, F_p^{n+1},$
                    $\hookrightarrow \sigma_p^{n+1}, \varepsilon_p^{f,n}, \chi_p^n, t_p^{\chi,n}, t^{n+1}$) $\triangleright$*Update the failed particles and stress*
8:             **end if**
9:         **end for**
10:     **end for**
11:     **return** $\sigma_p^{n+1}, \varepsilon_p^{f,n+1}, \chi_p^{n+1}, t_p^{\chi,n+1}, D_p^{n+1}$
12: **end procedure**

4.5.12 Updating the particle erosion parameter

The particle failure indicator function is updated in this procedure and used later for particle deletion if needed.

Algorithm 15 Updating the particle erosion parameter

Require: D_p^n, χ_p^n materialList, mpmFlags, constitutiveModel

1: **procedure** UPDATEEROSIONPARAMETER
2:     **for** matl **in** materialList **do**
3:         **for** part **in** particleList **do**
4:             **if** matl.doBasicDamage = TRUE **then**
5:                 $\chi_p^{n+1} \leftarrow$ damageModel.GETLOCALIZATIONPARAMETER() $\triangleright$*Just get the indicator*
                    $\hookrightarrow$ *parameter for particles that will be eroded.*
6:             **else**
7:                 $\chi_p^{n+1}, D_p^{n+1} \leftarrow$ constitutiveModel[matl].GETDAMAGEPARAMETER(χ_p^n, D_p^n)
                $\hookrightarrow$ $\triangleright$*Update the damage parameter in the constitutive model.*
8:             **end if**
9:         **end for**
10:     **end for**
11:     **return** χ_p^{n+1}, D_p^{n+1}
12: **end procedure**

4.5.13 Interpolating back to the particles and update

This is the final step at which the particle velocities and positions are updated and the grid is reset. Particle that are to be removed are dealt with in a subsequent relocation step.

Algorithm 16 Interpolating back to the particles and position update

---

Require: Δt, $\mathbf{a}_g$, $\mathbf{v}_g^\star$, $\mathbf{x}_p^n$, $\mathbf{v}_p^n$, $\mathbf{u}_p^n$, h_p^n, χ_p^{n+1}, F_p^{n+1}, V_p^{n+1}, materialList, particleList, gridNodeList, mpmFlags

1: **procedure** INTERPOLATETOPARTICLESANDUPDATE
2:     interpolator $\leftarrow$ CREATEINTERPOLATOR(mpmFlags)
3:     **for** matl **in** materialList **do**
4:         $h_p^{n+1} \leftarrow h_p^n$
5:         **for** part **in** particleList **do**
6:             $n_{gp}, S_{gp} \leftarrow$ interpolator.FINDCELLSANDWEIGHTS($\mathbf{x}_p^n$, h_p^{n+1}, F_p^{n+1})
7:             $\mathbf{v} \leftarrow \mathbf{o}, \quad \mathbf{a} \leftarrow \mathbf{o},$
8:             **for** node **in** gridNodeList **do**
9:                 $\mathbf{v} \leftarrow \mathbf{v} + \mathbf{v}_g^\star[\text{node}] * S_{gp}[\text{node}]$ ▷ *Update particle velocity*
10:                 $\mathbf{a} \leftarrow \mathbf{a} + \mathbf{a}_g[\text{node}] * S_{gp}[\text{node}]$ ▷ *Update particle acceleration*
11:             **end for**
12:             $\mathbf{x}_p^{n+1} \leftarrow \mathbf{x}_p^n + \mathbf{v} * \Delta t$ ▷ *Update position*
13:             $\mathbf{u}_p^{n+1} \leftarrow \mathbf{u}_p^n + \mathbf{v} * \Delta t$ ▷ *Update displacement*
14:             $\mathbf{v}_p^{n+1} \leftarrow \mathbf{v}_p^n + \mathbf{a} * \Delta t$ ▷ *Update velocity*
15:         **end for**
16:     **end for**
17:     DELETEROGUEPARTICLES()   ▷ *Delete particles that are to be eroded.*
18:     **return** V_p^{n+1}, $\mathbf{u}_p^{n+1}$, $\mathbf{v}_p^{n+1}$, $\mathbf{x}_p^{n+1}$, m_p, h_p^{n+1}
19: **end procedure**

4.6 Vaango Implementation

This section is intended to serve as a reference for users who find themselves needing to modify the source code, or those who are simply interested.

The goal of this section is to provide a mapping from the the algorithm described above to the software that carries it out. This won't be exhaustive, but will be a good starting point for the motivated reader.

The source code for the VAANGO -MPM implementation can be found in

```
src/CCA/Components/MPM
```

Within that directory are a number of files and subdirectories, these will be discussed as needed. For the moment, consider the various files that end in "MPM.cc":

```
AMRMPM.cc   FractureMPM.cc   ImpMPM.cc   RigidMPM.cc   SerialMPM.cc   ShellMPM.cc
```

AMRMPM.cc is the nascent beginnings of an AMR implementation of MPM. It is far from complete and should be ignored. FractureMPM.cc is an implementation of the work of Guo and Nairn [1], and while it is viable, it is undocumented and unsupported. ShellMPM.cc is a treatment of MPM particles as shell and membrane elements, developed by Biswajit Bannerjee. It is also viable, but also undocumented and unsupported. ImpMPM.cc is an implicit time integration form of MPM based on the work of Guilkey and Weiss [2]. It is also viable, and future releases of VAANGO will include documentation of its capabilities and uses. For now, interested readers should contact Jim Guilkey directly for more information. RigidMPM.cc contains a very reduced level of functionality, and is used solely in conjunction with the MPMArches component.

This leaves SerialMPM.cc This contains, despite its name, the parallel implementation of the algorithm described earlier and in the VAANGO Theory Manual. For now, we will skip over the initialization procedures such as:

```
SerialMPM::problemSetup
SerialMPM::scheduleInitialize
SerialMPM::actuallyInitialize
```

and focus mainly on the timestepping algorithm described above.

Each of the VAANGO components contains a function called scheduleTimeAdvance The algorithms implemented in these components are broken into a number of steps. The implementation of these steps in VAANGO take place in "tasks". Each task is responsible for performing the calculations needed to accomplish that step in the algorithm. Thus, each task requires some data upon which to operate, and it also creates some data, either as a final result, or as input to a subsequent task. Before individual tasks are executed, each is first "scheduled". The scheduling of tasks describes the dataflow and data dependencies for a given algorithm. By describing the data dependencies, both temporally and spatially, each task can be executed in the proper order, and communication tasks can automatically be generated by the VAANGO infrastructure to achieve parallelism. Thus, scheduleTimeAdvance calls a series of functions, each of which schedules the individual tasks. Let's begin by looking at the scheduleTimeAdvance for SerialMPM pasted below.

```cpp
void
SerialMPM::scheduleTimeAdvance(const LevelP & level,
                                 SchedulerP  & sched)
{
  MALLOC_TRACE_TAG_SCOPE("SerialMPM::scheduleTimeAdvance()");
  if (!flags->doMPMOnLevel(level->getIndex(), level->getGrid()->numLevels()))
    return;

  const PatchSet* patches = level->eachPatch();
  const MaterialSet* matls = d_sharedState->allMPMMaterials();

  scheduleApplyExternalLoads(           sched, patches, matls);
  scheduleInterpolateParticlesToGrid(   sched, patches, matls);
  scheduleExMomInterpolated(            sched, patches, matls);
  scheduleComputeContactArea(           sched, patches, matls);
  scheduleComputeInternalForce(         sched, patches, matls);

  scheduleComputeAndIntegrateAcceleration(sched, patches, matls);
  scheduleExMomIntegrated(              sched, patches, matls);
  scheduleSetGridBoundaryConditions(    sched, patches, matls);
  scheduleSetPrescribedMotion(          sched, patches, matls);
  scheduleComputeStressTensor(          sched, patches, matls);
  if(flags->d_doExplicitHeatConduction){
    scheduleComputeHeatExchange(        sched, patches, matls);
    scheduleComputeInternalHeatRate(    sched, patches, matls);
    scheduleComputeNodalHeatFlux(       sched, patches, matls);
    scheduleSolveHeatEquations(         sched, patches, matls);
    scheduleIntegrateTemperatureRate(   sched, patches, matls);
  }
  scheduleAddNewParticles(              sched, patches, matls);
  scheduleConvertLocalizedParticles(    sched, patches, matls);
  scheduleInterpolateToParticlesAndUpdate(sched, patches, matls);

  if(flags->d_canAddMPMMaterial){
    //  This checks to see if the model on THIS patch says that it's
    //  time to add a new material
    scheduleCheckNeedAddMPMMaterial(        sched, patches, matls);

    //  This one checks to see if the model on ANY patch says that it's
    //  time to add a new material
    scheduleSetNeedAddMaterialFlag(        sched, level,   matls);
  }

  sched->scheduleParticleRelocation(level, lb->pXLabel_preReloc,
                                 d_sharedState->d_particleState_preReloc,
                                 lb->pXLabel,
                                 d_sharedState->d_particleState,
                                 lb->pParticleIDLabel, matls);
  if(d_analysisModule){
    d_analysisModule->scheduleDoAnalysis( sched, level);
  }
}
```

The preceding includes scheduling for a number of rarely used features. For now, let's condense the pre-

ceding to the essential tasks:

```
void
SerialMPM::scheduleTimeAdvance(const LevelP & level,
                                    SchedulerP  & sched)
{
  if (!flags->doMPMOnLevel(level->getIndex(), level->getGrid()->numLevels()))
    return;

  const PatchSet* patches = level->eachPatch();
  const MaterialSet* matls = d_sharedState->allMPMMaterials();

  scheduleApplyExternalLoads(             sched, patches, matls);
  scheduleInterpolateParticlesToGrid(     sched, patches, matls);
  scheduleExMomInterpolated(              sched, patches, matls);
  scheduleComputeInternalForce(           sched, patches, matls);

  scheduleComputeAndIntegrateAcceleration(sched, patches, matls);
  scheduleExMomIntegrated(                sched, patches, matls);
  scheduleSetGridBoundaryConditions(      sched, patches, matls);
  scheduleComputeStressTensor(            sched, patches, matls);
  scheduleInterpolateToParticlesAndUpdate(sched, patches, matls);

  sched->scheduleParticleRelocation(level, lb->pXLabel_preReloc,
                                    d_sharedState->d_particleState_preReloc,
                                    lb->pXLabel,
                                    d_sharedState->d_particleState,
                                    lb->pParticleIDLabel, matls);
}
```

As described above, each of the "schedule" functions describes dataflow, and it also calls the function that actually executes the task. The naming convention is illustrated by an example, scheduleComputeAndIntegrateAcceleration calls computeAndIntegrateAcceleration Let's examine this particular task more carefully. First, the scheduling of the task:

```
void SerialMPM::scheduleComputeAndIntegrateAcceleration(SchedulerP& sched,
                                                  const PatchSet* patches,
                                                  const MaterialSet* matls)
{
  if (!flags->doMPMOnLevel(getLevel(patches)->getIndex(),
                           getLevel(patches)->getGrid()->numLevels()))
    return;

  printSchedule(patches,cout_doing,"MPM::scheduleComputeAndIntegrateAcceleration\t\t\t\t
      ");

  Task* t = scinew Task("MPM::computeAndIntegrateAcceleration",
                    this, &SerialMPM::computeAndIntegrateAcceleration);

  t->requires(Task::OldDW, d_sharedState->get_delt_label() );

  t->requires(Task::NewDW, lb->gMassLabel,          Ghost::None);
  t->requires(Task::NewDW, lb->gInternalForceLabel, Ghost::None);
  t->requires(Task::NewDW, lb->gExternalForceLabel, Ghost::None);
  t->requires(Task::NewDW, lb->gVelocityLabel,      Ghost::None);

  t->computes(lb->gVelocityStarLabel);
  t->computes(lb->gAccelerationLabel);

  sched->addTask(t, patches, matls);
}
```

The if statement basically directs the schedule to only do this task on the finest level (MPM can be used in AMR simulations, but only at the finest level.) The printSchedule command is in place for debugging purposes, this type of print statement can be turned on by setting an environmental variable. The real business of this task begins with the declaration of the Task. In the task declaration, the function associated with that task is identified. Subsequent to that is a description of the data dependencies. Namely, this task requires the mass, internal and external forces as well as velocity on the grid. No ghost data are required as this task is a node by node calculation. It also requires the timestep size. Note also that most of the

required data are needed from the NewDW where DW refers to DataWarehouse. This simply means that these data were calculated by an earlier task in the current timestep. The timestep size for this step was computed in the previous timestep, and thus is required from the OldDW Finally, this task computes the acceleration and time advanced velocity at each node.

The code to execute this task is as follows:

```
void SerialMPM::computeAndIntegrateAcceleration(const ProcessorGroup*,
                                                const PatchSubset* patches,
                                                const MaterialSubset*,
                                                DataWarehouse* old_dw,
                                                DataWarehouse* new_dw)
{
  for(int p=0;p<patches->size();p++){
    const Patch* patch = patches->get(p);
    printTask(patches, patch,cout_doing,"Doing computeAndIntegrateAcceleration\t\t\t\t")
        ;

    Ghost::GhostType  gnone = Ghost::None;
    Vector gravity = d_sharedState->getGravity();
    for(int m = 0; m < d_sharedState->getNumMPMMatls(); m++){
      MPMMaterial* mpm_matl = d_sharedState->getMPMMaterial( m );
      int dwi = mpm_matl->getDWIndex();

      // Get required variables for this patch
      constNCVariable<Vector> internalforce, externalforce, velocity;
      constNCVariable<double> mass;

      delt_vartype delT;
      old_dw->get(delT, d_sharedState->get_delt_label(), getLevel(patches) );

      new_dw->get(internalforce,lb->gInternalForceLabel, dwi, patch, gnone, 0);
      new_dw->get(externalforce,lb->gExternalForceLabel, dwi, patch, gnone, 0);
      new_dw->get(mass,         lb->gMassLabel,          dwi, patch, gnone, 0);
      new_dw->get(velocity,     lb->gVelocityLabel,      dwi, patch, gnone, 0);

      // Create variables for the results
      NCVariable<Vector> velocity_star,acceleration;
      new_dw->allocateAndPut(velocity_star, lb->gVelocityStarLabel, dwi, patch);
      new_dw->allocateAndPut(acceleration,  lb->gAccelerationLabel, dwi, patch);

      acceleration.initialize(Vector(0.,0.,0.));
      double damp_coef = flags->d_artificialDampCoeff;

      for(NodeIterator iter=patch->getExtraNodeIterator__New();
                       !iter.done();iter++){
        IntVector c = *iter;
        Vector acc(0.,0.,0.);
        if (mass[c] > flags->d_min_mass_for_acceleration){
          acc  = (internalforce[c] + externalforce[c])/mass[c];
          acc -= damp_coef*velocity[c];
        }
        acceleration[c] = acc +  gravity;
        velocity_star[c] = velocity[c] + acceleration[c] * delT;
      }
    }    // matls
  }
}
```

This task contains three nested for loops. First, is a loop over all of the "patches" that the processor executing this task is responsible for. Next is a loop over all materials (imagine a simulation involving the interaction between, say, tungsten and copper). Within this loop, the required data are retrieved from the new_dw (New DataWarehouse) and space for the data to be created is allocated. The final loop is over all of the nodes on the current patch, and the calculations described by Equations 4.2 and 4.3 are carried out. (This also includes a linear damping term not described above.)

$$\mathbf{a}_i = \frac{\mathbf{F}_i^{\text{ext}} - \mathbf{F}_i^{\text{int}}}{m_i} \tag{4.2}$$

An explicit forward Euler method is used for the time integration:

$$\mathbf{v}_i^L = \mathbf{v}_i + \mathbf{a}_i \Delta t \tag{4.3}$$

Let's consider each task in turn. The remaining tasks will be described in much less detail, but the preceding dissection of a fairly simple task, along with a description of what the remaining tasks are intended to accomplish, should allow interested individuals to follow the remainder of the Vaango -MPM implementation.

1. scheduleApplyExternalLoads This task is mainly responsible for applying traction boundary conditions described in the input file. This is done by assigning external force vectors to the particles. If the user wishes to apply a load that is not possible to acheive via the input file options, it is straightforward to modify the code here to do "one-off" tests.

2. scheduleInterpolateParticlesToGrid The name of this task was poorly chosen, but has persisted. This task carries out the operations given in Equation 4.4. It also sets boundary conditions on some of the variables, such as the grid temperature, and grid velocity (in the case of symmetry BCs).

$$m_i = \sum_p S_{ip} m_p, \quad \mathbf{v}_i = \frac{\sum_p S_{ip} m_p \mathbf{v}_p}{m_i}, \quad \mathbf{F}_i^{\text{ext}} = \sum_p S_{ip} \mathbf{F}_p^{\text{ext}} \tag{4.4}$$

where i refers to individual nodes of the grid, m_p is the particle mass, $\mathbf{v}_p$ is the particle velocity, and $\mathbf{F}_p^{\text{ext}}$ is the external force on the particle.

3. scheduleExMomInterpolated This task actually exists in one of the contact models which can be found in the Contact directory. Each of those models has two main tasks. This is the the first of those. It is responsible for modifying the grid velocity computed by interpolateParticlesToGrid according to the rules for the particular contact model chosen in the input file. These models are briefly described in the Vaango Theory Manual.

4. scheduleComputeInternalForce This task computes the volume integral of the divergence of stress. Specifically, it carries out the operation given in Equation 4.5. It also computes some diagnostic data, if requested in the input file, such as the reaction forces (tractions) on the boundaries of the computational domain.

$$\mathbf{F}_i^{\text{int}} = \sum_p \mathbf{G}_{ip} \boldsymbol{\sigma}_p v_p, \tag{4.5}$$

where $\mathbf{G}_{ip}$ is the gradient of the shape function of the ith node evaluated at $\mathbf{x}_p$, and $\boldsymbol{\sigma}_p$ and v_p are the time n values of particle stress and volume respectively.

5. scheduleComputeAndIntegrateAcceleration As described previously, this task carries out the operations described in Equations 4.2 and 4.3.

6. scheduleExMomIntegrated This is the second of the contact tasks (see above). It is responsible for modifiying the time advanced grid velocity computed in computeAndIntegrateAcceleration

7. scheduleSetGridBoundaryConditions This task sets boundary conditions on the time advanced grid velocity. It also sets an acceleration boundary condition as well. However, rather than just setting the acceleration to a given value, it is computed by solving Equation 4.3 for acceleration, and then recomputing the acceleration (on all nodes) as:

$$\mathbf{a}_i = \frac{\mathbf{v}_i^L - \mathbf{v}_i}{\Delta t} \tag{4.6}$$

Doing this operation on all nodes has several advantages. For most interior nodes, the value for acceleration will be unchanged, but for nodes on the where the velocity has been altered by enforcing boundary conditions, and for nodes at which the contact models have altered the velocity, the acceleration will be modified to reflect that alteration.

8. scheduleComputeStressTensor The task, computeStressTensor exists in each of the models in the ConstitutiveModel directory. Each model is responsible for carrying out the operations given

in Equation 4.7,

$$v_p^{n+1} = \text{Det}(\mathbf{dF}_p^{n+1})v_p^n, \quad \mathbf{F}_p^{n+1} = \mathbf{dF}_p^{n+1}\mathbf{F}_p^n \tag{4.7}$$

and of course, as the name implies, it also computes the material stress. This task has one additional important function, and that is computing the timestep size for the subsequent step. The CFL condition dictates that the timestep size be limited according to:

$$\Delta t < \frac{\Delta x}{c + |u|} \tag{4.8}$$

where Δx is the cell spacing, c is the wavespeed in the material, and $|u|$ is the magnitude of the local velocity. Because the wavespeed may depend on the state of stress that a material is in, this task provides a convenient time at which to make this calculation. A timestep size is computed for all particles, and the minimum for the particle set on a given patch is put into a "reduction variable". The VAANGO infrastructure then does a global reduction to select the smallest timestep from across the domain.

9. scheduleInterpolateToParticlesAndUpdate This task carries out the operations in Equations 4.9 and 4.10, namely updating the particle state based on the grid solution.

$$\mathbf{v}_p(t + \Delta t) = \mathbf{v}_p(t) + \sum_i S_{ip}\mathbf{a}_i \Delta t \tag{4.9}$$

$$\mathbf{x}_p(t + \Delta t) = \mathbf{x}_p(t) + \sum_i S_{ip}\mathbf{v}_i^L \Delta t \tag{4.10}$$

10. scheduleParticleRelocation This task is not actually located in the MPM code, but in the VAANGO infrastructure. The idea is that as particles move, some will cross patch boundaries, and their data will need to be sent to other processors. This task is responsible for identifying particles that have left the patch that they were on, finding where they went, and sending their data to the correct processor.
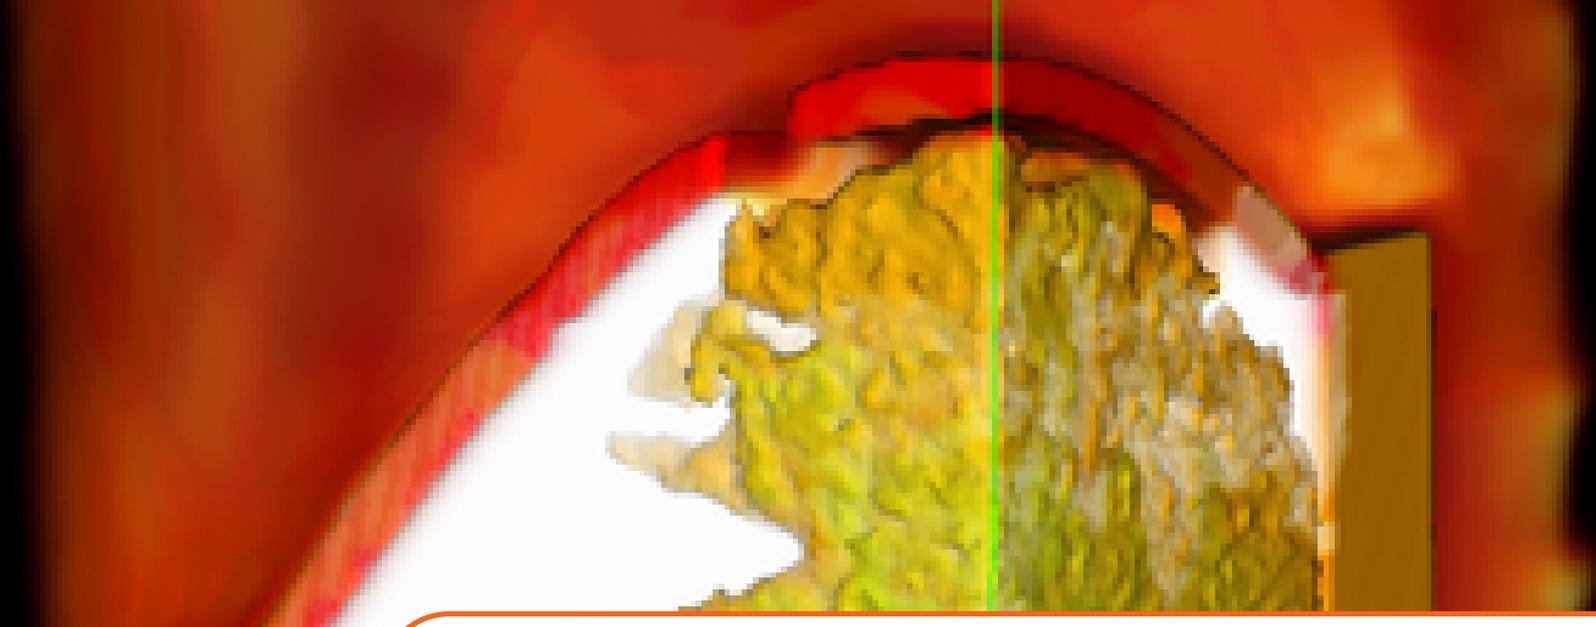
5 — Example MPM material model

In this chapter, we will examine the pseudocode for a reasonably complex material model that exercises most of the machinery for explicitly time integrated MPM material models in Vaango . The material model discussed in here is called Arena . A detailed description of the theory behind the model can be found in the Vaango Theory manual.

The main purpose of the material models is to compute the stress in a MPM particle given a state of deformation.

The stress tensor computation procedure calls the computeStressTensor routine that is specific to each constitutive model.

Algorithm 17 Computing the stress tensor

Require: $\Delta t, \mathbf{x}_p^n, m_p, V_p^{n+1}, \boldsymbol{h}_p^n, \boldsymbol{l}_p^{n+1}, \boldsymbol{F}_p^{n+1}, \boldsymbol{\sigma}_p^n, \boldsymbol{\eta}_p^n, \rho_\text{o}, \texttt{materialList}, \texttt{mpmFlags}, \texttt{constitutiveModel}$

1: **procedure** computeStressTensor
2:     **for** matl **in** materialList **do**
3:         $\boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\eta}_p^{n+1} \leftarrow \texttt{constitutiveModel}[\texttt{matl}].\text{computeStressTensor}(\Delta t, \mathbf{x}_p^n, m_p, V_p^{n+1}, \boldsymbol{h}_p^n,$
            $\hookrightarrow \boldsymbol{l}_p^{n+1}, \boldsymbol{F}_p^{n+1}, \boldsymbol{\sigma}_p^n, \boldsymbol{\eta}_p^n, \rho_\text{o}, \texttt{mpmFlags})$ ▷*Update the stress and any internal variables*
            $\hookrightarrow$ *needed by the constitutive model*
4:     **end for**
5:     **return** $\boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\eta}_p^{n+1}$
6: **end procedure**

The Arena model is a high strain-rate soil plasticity model that uses a "consistency bisection" algorithm to find the plastic strain direction and to update the internal state variables. A closest-point return algorithm in transformed stress space is used to project the trial stress state on to the yield surface. Because of the nonlinearities in the material models, it is easier to solve the problem by dividing the strain increment to substeps.

The Arena model treats the porosity (ϕ) and saturation (S_w) as internal variables in addition to the hydrostatic compressive strength (X), the isotropic backstress (ζ), and the plastic strain ($\boldsymbol{\varepsilon}^\text{p}$).

5.1 Initialization of the model

The model is initialized in two steps. In the first step, the constitutive model object is created followed by initialization of the stress (and the deformation gradient if needed). Here

- ϕ_o : The initial porosity.
- S_o : The initial porosity.
- $n_{\max}$: The maximum number of subcycles in the plastic return algorithm.
- $\varepsilon_{v,p}^{e,n}$: The elastic volumetric strain at a particle at $t = t_n$.
- σ_p^n : The dynamic Cauchy stress at a particle at $t = t_n$.
- $\sigma_{qs,p}^n$: The quasistatic Cauchy stress at a particle at $t = t_n$.
- σ^o : The initial Cauchy stress at a particle.

Algorithm 18 Creating the Arena constitutive model object

Require: `mpmFlags`, `xmlProblemSpec`

1: **procedure** CREATECONSTITUTIVEMODEL
2:     `elasticityModel` ← `ElasticModuliModelFactory`.CREATE(`xmlProblemSpec`)
3:     `yieldCondition` ← `YieldConditionFactory`.CREATE(`xmlProblemSpec`)
4:     $p_o, p_1, p_1^{\text{sat}}, p_2$ ← READCRUSHCURVEPARAMETERS(`xmlProblemSpec`)
5:     $\phi_o, S_o, \overline{p_o^w}$ ← READINITIALPOROSITYANDSATURATION(`xmlProblemSpec`)
6:     $n_{\max}$ ← READSUBCYCLINGCHARACTERISTICNUMBER(`xmlProblemSpec`)
7:     **return** `elasticityModel`, `yieldCondition`, $\phi_o, S_o, \overline{p_o^w}, n_{\max}, p_o, p_1, p_1^{\text{sat}}, p_2$
8: **end procedure**

Algorithm 19 Initializing the Arena particle variables

Require: $\phi_o, S_o, \overline{p_o^w}$, `particleList`, $V_p^o, m_p^o, \mathbf{v}_p^o$, `fluidParams`, `elasticityModel`, `yieldCondition`

1: **procedure** INITIALIZE
2:     `yieldCondition`.INITIALIZE(`particleList`, V_p^o)
3:     `yieldParams` ← `yieldCondition`.GETPARAMETERS()
4:     **for** `part` **in** `particleList` **do**
5:         $\phi_p^o[\text{part}] \leftarrow \phi_o$
6:         $S_{w,p}^o[\text{part}] \leftarrow S_o$
7:         $\chi_p^o[\text{part}] \leftarrow 0$
8:         $\varepsilon_{v,p}^{e,o}[\text{part}] \leftarrow 0$
9:         $\sigma_p^o[\text{part}] \leftarrow (\text{fluidParams}.\overline{p_o^w})\, \boldsymbol{I}$
10:         $\sigma_{qs,p}^o[\text{part}] \leftarrow \boldsymbol{\sigma}^o$
11:     **end for**
12:     Δt ← COMPUTESTABLETIMESTEP($V_p^o, m_p, \mathbf{v}_p^o$, `elasticityModel`)
13:     **return** $\phi_p^o, S_{w,p}^o, \chi_p^o, \varepsilon_{v,p}^{e,o}, \sigma_{qs,p}^o, \sigma_p^o, \Delta t$
14: **end procedure**

5.2 Computing the stress and internal variables

The COMPUTESTRESSTENSOR routine in the partially saturated Arena model assumes that the Biot coefficient is $B = 1$, and has the following form. Here we introduce the new variables

- ϕ_p^n, ϕ_p^{n+1} : The porosity at $t = t_n$ and $t = t_{n+1}$.
- $S_{w,p}^n, S_{w,p}^{n+1}$: The saturation at $t = t_n$ and $t = t_{n+1}$.
- $a_{1,p}, a_{2,p}, a_{3,p}, a_{4,p}$: Yield condition parameters at each particle.
- $p_{3,p}^n$: The particle crush curve parameter p_3 at $t = t_n$.
- X_p^n : The particle hydrostatic compressive strength at $t = t_n$.
- κ_p^n : The yield function branch point at $t = t_n$.
- $\boldsymbol{\epsilon}_p^{p,n}$: The particle plastic strain tensor at $t = t_n$.
- $\boldsymbol{\alpha}_p^n$: The particle backstress tensor at $t = t_n$.

- $\mathbf{d}^n$: The particle rate of deformation tensor at $t = t_n$.
- $\boldsymbol{R}^n, \boldsymbol{U}^n$: The particle rotation and stretch tensors at $t = t_n$.
- K^n, G^n : The particle bulk and shear modulus at $t = t_n$.

Algorithm 20 Computing the Arena stress tensor

Require: $\Delta t, \mathbf{x}_p^n, m_p, V_p^{n+1}, \boldsymbol{h}_p^n, \boldsymbol{l}_p^{n+1}, \boldsymbol{F}_p^n, \boldsymbol{F}_p^{n+1}, X_p^n, \kappa_p^n, \varepsilon_{v,p}^n, p_{3,p}^n, \boldsymbol{\epsilon}_p^{\mathrm{p},n}, \boldsymbol{\alpha}_p^n, \phi_p^n, S_{w,p}^n, \chi_p^n, \varepsilon_{v,p}^{\mathrm{e},n}, \boldsymbol{\sigma}_{qs,p}^n, \boldsymbol{\sigma}_p^n, \rho_\mathrm{o},$
    `particleList`, `mpmFlags`, `elasticityModel`, `yieldCondition`

1: **procedure** COMPUTESTRESSTENSOR
2:     `yieldParams` ← `yieldCondition`.GETPARAMETERS()
3:     $a_{1,p}, a_{2,p}, a_{3,p}, a_{4,p}, I_{1,p}^{\mathrm{peak}}, R_{c,p}$ ← `yieldCondition`.GETLOCALVARIABLES()
        ↪   ▷*Yield condition parameters vary at each particle.*
        ↪ *Get the per-particle values of these parameters.*
4:     **for** `part` **in** `particleList` **do**
5:         $\chi_p^{n+1}[\texttt{part}] \leftarrow \chi_p^n[\texttt{part}]$ ▷*Copy over failure indicator variable*
6:         $\boldsymbol{d}^{n+1} \leftarrow [\boldsymbol{l}_p^{n+1}[\texttt{part}] + (\boldsymbol{l}_p^{n+1}[\texttt{part}])^T]/2$ ▷*Compute rate of deformation*
7:         $\boldsymbol{R}^n, \boldsymbol{U}^n \leftarrow$ POLARDECOMPOSITION$(\boldsymbol{F}_p^n[\texttt{part}])$ ▷*Compute rotation and stretch tensors*
8:         $\boldsymbol{d}_{\mathrm{unrot}}^{n+1} \leftarrow (\boldsymbol{R}^n)^T \cdot \boldsymbol{d}^{n+1} \cdot \boldsymbol{R}^n$ ▷*Unrotate the rate of deformation tensor*
9:         $\boldsymbol{\sigma}_{qs,\mathrm{unrot}}^n \leftarrow (\boldsymbol{R}^n)^T \cdot \boldsymbol{\sigma}_{qs,p}^n[\texttt{part}] \cdot \boldsymbol{R}^n$ ▷*Unrotate the quasistatic stress*
10:        $\boldsymbol{\sigma}_{\mathrm{unrot}}^n \leftarrow (\boldsymbol{R}^n)^T \cdot \boldsymbol{\sigma}_p^n[\texttt{part}] \cdot \boldsymbol{R}^n$ ▷*Unrotate the total stress*
11:        $K^n, G^n, \boldsymbol{s}^n, (\overline{p^w})^n, I_1^{\mathrm{eff},n}, \sqrt{J_2^n}, r^n, z_{\mathrm{eff}}^n, \varepsilon_v^{\mathrm{p},n} \leftarrow$
        ↪ COMPUTEELASTICPROPERTIES$(\boldsymbol{\sigma}_{qs,\mathrm{unrot}}^n[\texttt{part}],$
        ↪ $\phi_p^n[\texttt{part}], S_{w,p}^n[\texttt{part}], \boldsymbol{\epsilon}_p^{\mathrm{p},n}[\texttt{part}], \boldsymbol{\alpha}_p^n[\texttt{part}], p_{3,p}^n[\texttt{part}])$
        ↪   ▷*Compute elastic properties and stress invariants*
12:        `isSuccess`, $\boldsymbol{\sigma}_{qs,\mathrm{unrot}}^{n+1}, \phi_{qs}^{n+1}, S_{w,qs}^{n+1}, X_p^{n+1}[\texttt{part}], \boldsymbol{\alpha}_p^{n+1}[\texttt{part}], \boldsymbol{\epsilon}_p^{\mathrm{p},n+1}[\texttt{part}] \leftarrow$
        ↪ RATEINDEPENDENTPLASTICUPDATE$(\Delta t, \boldsymbol{d}_{\mathrm{unrot}}^{n+1}, K^n, G^n, \boldsymbol{s}^n, (\overline{p^w})^n, I_1^{\mathrm{eff},n}, \sqrt{J_2^n}, r^n, z_{\mathrm{eff}}^n,$
        ↪ $\varepsilon_v^{\mathrm{p},n}, \boldsymbol{\sigma}_{qs,\mathrm{unrot}}^n, \phi_p^n[\texttt{part}], S_{w,p}^n[\texttt{part}], X_p^n[\texttt{part}], \boldsymbol{\alpha}_p^n[\texttt{part}], \boldsymbol{\epsilon}_p^{\mathrm{p},n}[\texttt{part}], p_{3,p}^n[\texttt{part}],$
        ↪ $a_{1,p}[\texttt{part}], a_{2,p}[\texttt{part}], a_{3,p}[\texttt{part}], a_{4,p}[\texttt{part}])$
        ↪   ▷*Compute updated quasistatic state using the consistency bisection algorithm*
13:        **if** `isSuccess` = FALSE **then**
14:           FLAGPARTICLEFORDELETION(`part`)
15:        **end if**
16:        $\boldsymbol{\sigma}_{\mathrm{unrot}}^{n+1} \leftarrow$ RATEDEPENDENTPLASTICUPDATE$(\Delta t, \boldsymbol{d}_{\mathrm{unrot}}^{n+1}, \boldsymbol{\sigma}_{qs,\mathrm{unrot}}^n, \boldsymbol{\sigma}_{qs,\mathrm{unrot}}^{n+1}, \phi_p^n[\texttt{part}], \phi_{qs}^{n+1},$
        ↪ $S_{w,p}^n[\texttt{part}], S_{w,qs}^{n+1}, X_p^n[\texttt{part}], X_p^{n+1}[\texttt{part}], \boldsymbol{\alpha}_p^n[\texttt{part}], \boldsymbol{\alpha}_p^{n+1}[\texttt{part}],$
        ↪ $\boldsymbol{\epsilon}_p^{\mathrm{p},n}[\texttt{part}], \boldsymbol{\epsilon}_p^{\mathrm{p},n+1}[\texttt{part}], p_{3,p}^n[\texttt{part}],$
        ↪ $a_{1,p}[\texttt{part}], a_{2,p}[\texttt{part}], a_{3,p}[\texttt{part}], a_{4,p}[\texttt{part}])$
17:        $\boldsymbol{R}^{n+1}, \boldsymbol{U}^{n+1} \leftarrow$ POLARDECOMPOSITION$(\boldsymbol{F}_p^{n+1}[\texttt{part}])$ ▷*Compute rotation and stretch tensors*
18:        $\boldsymbol{\sigma}_{p,qs}^{n+1}[\texttt{part}] \leftarrow \boldsymbol{R}^{n+1} \cdot \boldsymbol{\sigma}_{qs,\mathrm{unrot}}^{n+1} \cdot (\boldsymbol{R}^{n+1})^T$ ▷*Rotate the quasistatic stress*
19:        $\boldsymbol{\sigma}_p^{n+1}[\texttt{part}] \leftarrow \boldsymbol{R}^{n+1} \cdot \boldsymbol{\sigma}_{\mathrm{unrot}}^{n+1} \cdot (\boldsymbol{R}^{n+1})^T$ ▷*Rotate the dynamic stress*
20:     **end for**
21:     $\Delta t^{n+1} \leftarrow$ COMPUTESTABLETIMESTEP$(V_p^{n+1}, m_p, \boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\sigma}_p^n, \boldsymbol{l}_p^{n+1}, \texttt{elasticityModel})$
22:     **return** $\boldsymbol{\sigma}_p^{n+1}, \boldsymbol{\sigma}_{p,qs}^{n+1}, \phi_p^{n+1}, S_{w,p}^{n+1}, \boldsymbol{\alpha}_p^{n+1}, X_p^{n+1}, \boldsymbol{\epsilon}_p^{\mathrm{p},n+1}, \Delta t^{n+1}$.
23: **end procedure**

5.2.1   Compute elastic properties

The pseudocode for the generic COMPUTEELASTICPROPERTIES is listed below. The function has side effects beyond computing the elastic properties and should be used carefully. Note that the subscript p has been dropped for simplicity because all quantities are particle-based.

Algorithm 21 Computing the elastic properties

Require: $\boldsymbol{\sigma}, \phi, S_w, \boldsymbol{\epsilon}^{\mathrm{p}}, \boldsymbol{\alpha}, p_3,$ `elasticityModel`

1: **procedure** COMPUTEELASTICPROPERTIES
2:    $\boldsymbol{s}, \overline{p^w}, I_1^{\mathrm{eff}}, \sqrt{J_2}, r, z_{\mathrm{eff}} \leftarrow$ UPDATESTRESSINVARIANTS$(\boldsymbol{\sigma}, \boldsymbol{\alpha})$ ▷*Compute the deviatoric stress and the invariants of the input stress tensor.*
3:    $\varepsilon_v^{\mathrm{p}} \leftarrow$ UPDATEVOLUMETRICPLASTICSTRAIN$(\boldsymbol{\epsilon}^{\mathrm{p}})$ ▷*Compute the volumetric plastic strain from the input plastic strain tensor.*
4:    $K, G \leftarrow$ `elasticityModel`.GETCURRENTELASTICMODULI$(I_1^{\mathrm{eff}}, \overline{p^w}, \varepsilon_v^{\mathrm{p}}, \phi, S_w)$ ▷*Compute the elastic moduli corresponding to the input state.*
5:    **if** useDisaggregationAlgorithm = TRUE **then**
6:       scale = MAX$(\exp[-(p_3 + \varepsilon_v^{\mathrm{p}})], 10^{-5})$
7:       $K \leftarrow K^\star$scale , $G \leftarrow G^\star$scale
8:    **end if**
9:    **return** $K, G, \boldsymbol{s}, \overline{p^w}, I_1^{\mathrm{eff}}, \sqrt{J_2}, r, z_{\mathrm{eff}}, \varepsilon_v^{\mathrm{p}}$
10: **end procedure**

Algorithm 22 Updating the stress invariants

1: **procedure** UPDATESTRESSINVARIANTS$(\boldsymbol{\sigma}, \boldsymbol{\alpha})$
2:    $I_1 \leftarrow \mathrm{tr}(\boldsymbol{\sigma})$
3:    $\boldsymbol{s} \leftarrow \boldsymbol{\sigma} - (I_1/3)\boldsymbol{I}$ ▷*Compute deviatoric stress*
4:    $\overline{p^w} \leftarrow -\mathrm{tr}(\boldsymbol{\alpha})/3$ ▷*Compute pore pressure*
5:    $I_1^{\mathrm{eff}} \leftarrow I_1 + 3\overline{p^w}$, $\sqrt{J_2} \leftarrow \sqrt{(1/2)\boldsymbol{s}:\boldsymbol{s}}$ ▷*Compute invariants of the effective stress*
6:    $r \leftarrow \sqrt{2J_2}$, $z_{\mathrm{eff}} \leftarrow I_1^{\mathrm{eff}}/\sqrt{3}$ ▷*Compute Lode coordinates of the effective stress*
7:    **return** $\boldsymbol{s}, \overline{p^w}, I_1^{\mathrm{eff}}, \sqrt{J_2}, r, z_{\mathrm{eff}}$
8: **end procedure**

The elastic modulus computation procedures:

The functions used to compute the moduli are listed in the pseudocode below.

Algorithm 23 Computing the current elastic moduli

Require: $I_1^{\mathrm{eff}}, \overline{p^w}, \varepsilon_v^{\mathrm{p}}, \phi, S_w$

1: **procedure** GETCURRENTELASTICMODULI
2:    $\overline{I}_1^{\mathrm{eff}} \leftarrow -I_1^{\mathrm{eff}}, \overline{\varepsilon_v^{\mathrm{p}}} \leftarrow -\varepsilon_v^{\mathrm{p}},$
3:    $K \leftarrow 0, G \leftarrow 0$
4:    **if** $S_w > 0$ **then**
5:       $K, G \leftarrow$ COMPUTEPARTIALSATURATEDMODULI$(\overline{I}_1^{\mathrm{eff}}, \overline{p^w}, \overline{\varepsilon_v^{\mathrm{p}}}, \phi, S_w)$
6:    **else**
7:       $K, G \leftarrow$ COMPUTEDRAINEDMODULI$(\overline{I}_1^{\mathrm{eff}}, \overline{\varepsilon_v^{\mathrm{p}}})$
8:    **end if**
9:    **return** K, G
10: **end procedure**

Algorithm 24 Computing the partially saturated elastic moduli

---

Require: $K_{s0}, n_s, \overline{p}_{s0}, K_{w0}, n_w, \overline{p}_{w0}, \gamma, \overline{p}_r$

1: **procedure** COMPUTEPARTIALSATURATEDMODULI$(\overline{I}_1^{\mathrm{eff}}, \overline{p^w}, \overline{\varepsilon_v^{\mathrm{p}}}, \phi, S_w)$
2:    **if** $\overline{I}_1^{\mathrm{eff}} > 0$ **then**
3:       $\overline{p}^{\mathrm{eff}} \leftarrow \overline{I}_1^{\mathrm{eff}}/3$
4:       $K_s \leftarrow K_{s0} + n_s(\overline{p}^{\mathrm{eff}} - \overline{p}_{s0})$

5: $\quad K_w \leftarrow K_{wo} + n_w(\overline{p^w} - \overline{p}_{wo})$

6: $\quad K_a \leftarrow \gamma(\overline{p^w} + \overline{p}_r)$

7: $\quad K_d, G \leftarrow \text{computeDrainedModuli}(\overline{I}_1^{\text{eff}}, \overline{\varepsilon_v^p})$

8: $\quad K_f \leftarrow 1/\left[S_w/K_w + (1 - S_w)/K_a\right]$ $\qquad\qquad$ ▷*Bulk modulus of air + water mixture*

9: $\quad \text{numer} \leftarrow (1 - K_d/K_s)^2$

10: $\quad \text{denom} \leftarrow (1/K_s)(1 - K_d/K_s) + \phi(1/K_f - 1/K_s)$

11: $\quad K \leftarrow K_d + \text{numer}/\text{denom}$ $\qquad\qquad$ ▷*Bulk modulus of partially saturated material*

$\qquad\quad \hookrightarrow$ *(Biot-Gassman model)*

12: **else**

13: $\quad K, G \leftarrow \text{computeDrainedModuli}(\overline{I}_1, \overline{\varepsilon_v^p})$

14: **end if**

15: **return** K, G

16: **end procedure**

Algorithm 25 Computing the drained elastic moduli

Require: $K_{so}, n_s, \overline{p}_{so}, b_0, b_1, b_2, b_3, b_4, G_0, v_1, v_2$

1: **procedure** computeDrainedModuli$(\overline{I}_1^{\text{eff}}, \overline{\varepsilon_v^p})$

2: $\quad$ **if** $\overline{I}_1^{\text{eff}} > 0$ **then**

3: $\qquad \overline{p}^{\text{eff}} \leftarrow \overline{I}_1^{\text{eff}}/3$

4: $\qquad K_s \leftarrow K_{so} + n_s(\overline{p}^{\text{eff}} - \overline{p}_{so})$

5: $\qquad K_s^{\text{ratio}} \leftarrow K_s/(1 - n_s\overline{p}^{\text{eff}}/K_s)$

6: $\qquad \varepsilon_v^e \leftarrow \text{pow}((b_3\overline{p}^{\text{eff}})/(b_1 K_s - b_2\overline{p}^{\text{eff}}), (1/b_4));$

7: $\qquad y \leftarrow \text{pow}(\varepsilon_v^e, b_4)$

8: $\qquad z \leftarrow b_2 y + b_3$

9: $\qquad K \leftarrow K_s^{\text{ratio}}[b_0 + (1/\varepsilon_v^e)b_1 b_3 b_4 y/z^2];$ $\qquad$ ▷ *Compute compressive bulk modulus*

10: $\qquad v = v_1 + v_2 \exp(-K/K_s)$

11: $\qquad G \leftarrow G_0$

12: $\qquad$ **if** $v > 0$ **then**

13: $\qquad\quad G \leftarrow 1.5K(1 - 2v)/(1 + v)$ $\qquad\qquad$ ▷*Update the shear modulus (if $v_1, v_2 > 0$)*

14: $\qquad$ **end if**

15: $\quad$ **else**

16: $\qquad K \leftarrow b_0 K_{so}$ $\qquad\qquad$ ▷*Tensile bulk modulus = Bulk modulus at $p = 0$*

17: $\qquad G \leftarrow G_0$ $\qquad\qquad$ ▷*Tensile shear modulus*

18: $\quad$ **end if**

19: $\quad$ **return** K, G

20: **end procedure**

5.2.2 Rate-independent stress update

The partially saturated soil model uses a "consistency bisection" algorithm to find the plastic strain direction and to update the internal state variables. A closest-point return algorithm in transformed stress space is used to project the trial stress state on to the yield surface. Because of the nonlinearities in the material models, it is easier to solve the problem by dividing the strain increment into substeps.

The pseudocode for the algorithm is given below. All quantities are particle-based and the subscript p has been dropped for convenience.

Algorithm 26 The rate-independent stress and internal variable update algorithm

---

Require: $\Delta t, \boldsymbol{d}^{n+1}, K^n, G^n, \boldsymbol{s}^n, (\overline{p^w})^n, I_1^{\text{eff},n}, \sqrt{J_2^n}, r^n, z_{\text{eff}}^n, \varepsilon_v^{p,n}, a_1, a_2, a_3, a_4, p_3^n, \boldsymbol{\sigma}_{qs}^n, \phi^n, S_w^n, X^n, \boldsymbol{\alpha}^n, \boldsymbol{\varepsilon}^{p,n},$
$\qquad n_{\max}, \epsilon_{\text{sub}}, \chi_{\max}$

1: **procedure** rateIndependentPlasticUpdate

2:      $\boldsymbol{\sigma}_{\text{trial}} \leftarrow \textsc{ComputeTrialStress}(\boldsymbol{\sigma}_{qs}^n, K^n, G^n, \boldsymbol{d}^{n+1}, \Delta t)$ ▷*Compute trial stress*

3:      $\boldsymbol{\alpha}^{\text{trial}} \leftarrow \boldsymbol{\alpha}^n, p_3^{\text{trial}} \leftarrow p_3^n, \phi^{\text{trial}} \leftarrow \phi^n, S_w^{\text{trial}} \leftarrow S_w^n,$

        $\hookrightarrow X^{\text{trial}} \leftarrow X^n, \boldsymbol{\varepsilon}^{\text{p,trial}} \leftarrow \boldsymbol{\varepsilon}^{\text{p},n}$ ▷*Set all other trial quantities to the values*

        $\hookrightarrow$ *at the beginning of the timestep*

4:      $K^{\text{trial}}, G^{\text{trial}}, \boldsymbol{s}^{\text{trial}}, (\overline{p^w})^{\text{trial}}, I_1^{\text{eff,trial}}, \sqrt{J_2^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, \varepsilon_v^{\text{p,trial}} \leftarrow$

        $\hookrightarrow \textsc{ComputeElasticProperties}(\boldsymbol{\sigma}_{\text{trial}}, \phi^{\text{trial}}, S_w^{\text{trial}}, \boldsymbol{\varepsilon}^{\text{p,trial}}, \boldsymbol{\alpha}^{\text{trial}}, p_3^{\text{trial}})$ ▷*Update the trial*

        $\hookrightarrow$ *values of the moduli and compute the invariants of the trial stress*

5:      $n_{\text{sub}} \leftarrow \textsc{ComputeStepDivisions}(n_{\max}, \epsilon_{\text{sub}}, K^n, K^{\text{trial}}, I_1^{\text{peak}}, a_1, X^n, \boldsymbol{\sigma}_{qs}^n, \boldsymbol{\sigma}_{\text{trial}})$

        $\hookrightarrow$ ▷*Compute number of substeps used by the return algorithm*

6:     **if** $n_{\text{sub}} < 0$ **then**

7:        **return** isSuccess = FALSE

8:     **end if**

9:      $\delta t \leftarrow \dfrac{\Delta t}{n_{\text{sub}}}$ ▷*Substep timestep*

10:     $\chi \leftarrow 1, t_{\text{local}} \leftarrow 0$ ▷*Initialize substep multiplier and accumulated time increment*

11:     $\boldsymbol{\sigma}^k \leftarrow \boldsymbol{\sigma}_{qs}^n, \boldsymbol{\varepsilon}^{\text{p},k} \leftarrow \boldsymbol{\varepsilon}^{\text{p},n}, \phi^k \leftarrow \phi^n, S_w^k \leftarrow S_w^n, X^k \leftarrow X^n, \boldsymbol{\alpha}^k \leftarrow \boldsymbol{\alpha}^n, K^k \leftarrow K^n, G^k \leftarrow G^n, p_3^k \leftarrow p_3^n,$

        $\hookrightarrow \boldsymbol{s}^k \leftarrow \boldsymbol{s}^n, (\overline{p^w})^k \leftarrow (\overline{p^w})^n, I_1^{\text{eff},k} \leftarrow I_1^{\text{eff},n}, \sqrt{J_2^k} \leftarrow \sqrt{J_2^n}, r^k \leftarrow r^n, z_{\text{eff}}^k \leftarrow z_{\text{eff}}^n, \varepsilon_v^{\text{p},k} \leftarrow \varepsilon_v^{\text{p},n}$

12:     **repeat**

13:        isSuccess, $\boldsymbol{\sigma}^{k+1}, \boldsymbol{\varepsilon}^{\text{p},k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1}, \boldsymbol{\alpha}^{k+1}, K^{k+1}, G^{k+1}, p_3^{k+1} \leftarrow$

        $\hookrightarrow \textsc{ComputeSubstep}(\boldsymbol{\sigma}^k, \boldsymbol{\varepsilon}^{\text{p},k}, \phi^k, S_w^k, X^k, \boldsymbol{\alpha}^k, K^k, G^k, \boldsymbol{s}^k, (\overline{p^w})^k, I_1^{\text{eff},k}, \sqrt{J_2^k}, r^k, z_{\text{eff}}^k,$

        $\hookrightarrow \varepsilon_v^{\text{p},k}, p_3^k, \boldsymbol{d}^{n+1}, \delta t)$

        $\hookrightarrow$ ▷*Compute updated stress and internal variables for the current substep*

14:        **if** isSuccess = TRUE **then**

15:           $t_{\text{local}} \leftarrow t_{\text{local}} + \delta t$

16:           $\boldsymbol{\sigma}^k \leftarrow \boldsymbol{\sigma}^{k+1}, \boldsymbol{\varepsilon}^{\text{p},k} \leftarrow \boldsymbol{\varepsilon}^{\text{p},k+1}, \phi^k \leftarrow \phi^{k+1}, S_w^k \leftarrow S_w^{k+1}, X^k \leftarrow X^{k+1}, \boldsymbol{\alpha}^k \leftarrow \boldsymbol{\alpha}^{k+1}$

17:           $K^k \leftarrow K^{k+1}, G^k \leftarrow G^{k+1}, p_3^k \leftarrow p_3^{k+1}$

18:        **else**

19:           $\delta t \leftarrow \delta t / 2$ ▷*Halve the timestep*

20:           $\chi \leftarrow 2\chi$ ▷*Keep a count of how many times the timestep has been halved.*

21:           **if** $\chi > \chi_{\max}$ **then**

22:              **return** isSuccess = FALSE, $\boldsymbol{\sigma}^k, \phi^k, S_w^k, X^k, \boldsymbol{\alpha}^k, \boldsymbol{\varepsilon}^{\text{p},k}, K^k, G^k, p_3^k$

             $\hookrightarrow$ ▷*Algorithm has failed to converge*

23:           **end if**

24:        **end if**

25:     **until** $t_{\text{local}} \geq \Delta t$

26:     $\boldsymbol{\sigma}_{qs}^{n+1} \leftarrow \boldsymbol{\sigma}^{k+1}, \boldsymbol{\alpha}^{n+1} \leftarrow \boldsymbol{\alpha}^{k+1}, \boldsymbol{\varepsilon}^{\text{p},n+1} \leftarrow \boldsymbol{\varepsilon}^{\text{p},k+1}, \phi^{n+1} \leftarrow \phi^{k+1}, S_w^{n+1} \leftarrow S_w^{k+1}, X^{n+1} \leftarrow X^{k+1}$

27:     $K^{n+1} \leftarrow K^{k+1}, G^{n+1} \leftarrow G^{k+1}, p_3^{n+1} \leftarrow p_3^{k+1}$

28:     **return** isSuccess = TRUE, $\boldsymbol{\sigma}_{qs}^{n+1}, \phi^{n+1}, S_w^{n+1}, X^{n+1}, \boldsymbol{\alpha}^{n+1}, \boldsymbol{\varepsilon}^{\text{p},n+1}, K^{n+1}, G^{n+1}, p_3^{n+1}$

        $\hookrightarrow$ ▷*Algorithm has converged*

29: **end procedure**

Computing the trial stress

The pseudocode of the trial stress algorithm is given below.

Algorithm 27 Computing the trial stress

1:    **procedure** $\textsc{ComputeTrialStress}(\boldsymbol{\sigma}_{qs}^n, K^n, G^n, \boldsymbol{d}^{n+1}, \Delta t)$

2:      $\Delta \boldsymbol{\varepsilon} \leftarrow \boldsymbol{d}^{n+1} \Delta t$ ▷*Total strain increment*

3:      $\Delta \boldsymbol{\varepsilon}^{\text{iso}} \leftarrow \frac{1}{3} \text{tr}(\Delta \boldsymbol{\varepsilon}) \boldsymbol{I}$

4:      $\Delta \boldsymbol{\varepsilon}^{\text{dev}} \leftarrow \Delta \boldsymbol{\varepsilon} - \Delta \boldsymbol{\varepsilon}^{\text{iso}}$

5:      $\boldsymbol{\sigma}_{\text{trial}} \leftarrow \boldsymbol{\sigma}_{qs}^n + 3K^n \Delta \boldsymbol{\varepsilon}^{\text{iso}} + 2G^n \Delta \boldsymbol{\varepsilon}^{\text{dev}}$

6:       **return** $\boldsymbol{\sigma}_{\text{trial}}$

7: **end procedure**

Computing the number of subcycles in the return algorithm

To allow for nonlinear parameter variations, the algorithm breaks a trial loading step into subcycles. The algorithm below, determines the number of substeps based on the magnitude of the trial stress increment relative to the characteristic dimensions of the yield surface. Another comparison uses the value of the pressure dependent elastic properties at $\boldsymbol{\sigma}_{qs}^{n}$ and $\boldsymbol{\sigma}_{\text{trial}}$ and adjusts the number of substeps if there is a large change in elastic moduli. This ensures an accurate solution for nonlinear elasticity even with fully elastic loading.

Algorithm 28 Computing the number of subcycles

Require: $n_{\max}, \epsilon_{\text{sub}} \leftarrow 10^{-4}, K^{n}, K^{\text{trial}}, I_{1}^{\text{peak}}, a_{1}, X^{n}, \boldsymbol{\sigma}_{qs}^{n}, \boldsymbol{\sigma}_{\text{trial}}$

1: **procedure** COMPUTESTEPDIVISIONS

2:    $n_{\text{bulk}} \leftarrow \left\lceil \left| K^{n} - K^{\text{trial}} \right| / K^{n} \right\rceil$ ▷*Compute change in bulk modulus*

3:    $\Delta\boldsymbol{\sigma} \leftarrow \boldsymbol{\sigma}_{\text{trial}} - \boldsymbol{\sigma}_{qs}^{n}$

4:    $L \leftarrow \frac{1}{2}\left(I_{1}^{\text{peak}} - X^{n} \right)$

5:    **if** $a_{1} > 0$ **then**

6:        $L \leftarrow \text{MIN}(L, a_{1})$

7:    **end if**

8:    $n_{\text{yield}} \leftarrow \left\lceil \epsilon_{\text{sub}} \times \left\| \Delta\boldsymbol{\sigma} \right\| / L \right\rceil$ ▷*Compute trial stress increment relative to yield surface size*

9:    $n_{\text{sub}} \leftarrow \text{MAX}(n_{\text{bulk}}, n_{\text{yield}})$ ▷*n_{sub} is the maximum of the two values*

10:    **if** $n_{\text{sub}} > n_{\max}$ **then**

11:        $n_{\text{sub}} \leftarrow -1$

12:    **end if**

13:    **return** n_{sub}

14: **end procedure**

Updating the stress for a substep: consistency bisection

This procedure computes the updated stress state for a substep that may be either elastic, plastic, or partially elastic. It uses Homel's consistency bisection and non-hardening return concepts [3].

Algorithm 29 Computing the stress and internal variable update for a substep

Require: $\boldsymbol{d}^{n+1}, \delta t, \boldsymbol{\sigma}^{k}, \boldsymbol{\varepsilon}^{\text{p},k}, \phi^{k}, S_{w}^{k}, X^{k}, \boldsymbol{\alpha}^{k}, K^{k}, G^{k}, \boldsymbol{s}^{k}, (\overline{p^{w}})^{k}, I_{1}^{\text{eff},k}, \sqrt{J_{2}^{k}}, r^{k}, z_{\text{eff}}^{k}, \varepsilon_{v}^{\text{p},k}, p_{3}^{k}, a_{1}, a_{2}, a_{3}, a_{4},$
    $I_{1}^{\text{peak}}, R_{c}, \beta,$ `yieldCondition`

1: **procedure** COMPUTESUBSTEP

2:    $\delta\boldsymbol{\varepsilon} \leftarrow \boldsymbol{d}^{n+1}\delta t$ ▷*Compute strain increment*

3:    $\boldsymbol{\sigma}_{\text{trial}} \leftarrow \text{COMPUTETRIALSTRESS}(\boldsymbol{\sigma}^{k}, K^{k}, G^{k}, \boldsymbol{d}^{n+1}, \Delta t)$ ▷*Compute substep trial stress*

4:    $\boldsymbol{\alpha}^{\text{trial}} \leftarrow \boldsymbol{\alpha}^{k}, K^{\text{trial}} \leftarrow K^{k}, G^{\text{trial}} \leftarrow G^{k}, p_{3}^{\text{trial}} \leftarrow p_{3}^{k}, \phi^{\text{trial}} \leftarrow \phi^{k}, S_{w}^{\text{trial}} \leftarrow S_{w}^{k},$
        ↪ $X^{\text{trial}} \leftarrow X^{k}, \boldsymbol{\varepsilon}^{\text{p,trial}} \leftarrow \boldsymbol{\varepsilon}^{\text{p},k}$ ▷*Set all other trial quantities to the values*
      ↪ *at the beginning of the substep*

5:    $K^{\text{trial}}, G^{\text{trial}}, \boldsymbol{s}^{\text{trial}}, (\overline{p^{w}})^{\text{trial}}, I_{1}^{\text{eff,trial}}, \sqrt{J_{2}^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, \varepsilon_{v}^{\text{p,trial}} \leftarrow$
        ↪ $\text{COMPUTEELASTICPROPERTIES}(\boldsymbol{\sigma}_{\text{trial}}, \phi^{\text{trial}}, S_{w}^{\text{trial}}, \boldsymbol{\varepsilon}^{\text{p,trial}}, \boldsymbol{\alpha}^{\text{trial}}, p_{3}^{\text{trial}})$
        ↪   ▷*Compute elastic moduli and stress invariants for the trial state*

6:    $\text{isElastic} \leftarrow$ `yieldCondition`$.\text{EVALYIELDCONDITION}(I_{1}^{\text{eff,trial}}, \sqrt{J_{2}^{\text{trial}}}, X^{\text{trial}}, (\overline{p^{w}})^{\text{trial}}, \phi^{\text{trial}}, S_{w}^{\text{trial}},$
        ↪ $a_{1}, a_{2}, a_{3}, a_{4}, I_{1}^{\text{peak}}, R_{c}, \beta)$
        ↪   ▷*Determine whether the trial stress is elastic or not*

7:    **if** isElastic = TRUE **then**

8:     $\sigma^{k+1} \leftarrow \sigma_{\text{trial}}, \varepsilon^{\text{p},k+1} \leftarrow \varepsilon^{\text{p,trial}}, \phi^{k+1} \leftarrow \phi^{\text{trial}}, S_w^{k+1} \leftarrow S_w^{\text{trial}}, X^{k+1} \leftarrow X^{\text{trial}}, \alpha^{k+1} \leftarrow \alpha^{\text{trial}}$

9:     $K^{k+1} \leftarrow K^{\text{trial}}, G^{k+1} \leftarrow G^{\text{trial}}, p_3^{k+1} \leftarrow p_3^{\text{trial}}$

    ↪   *▷This is an elastic substep. Update the state to the trial value.*

10:     isSuccess = TRUE

11:     **return** isSuccess, $\sigma^{k+1}, \varepsilon^{\text{p},k+1} \phi^{k+1}, S_w^{k+1}, X^{k+1}, \alpha^{k+1}, K^{k+1}, G^{k+1}, p_3^{k+1}$

12:   **end if**

13:   $\sigma_{\text{fixed}}, \delta\varepsilon_{\text{fixed}}^{\text{p}} \leftarrow \text{NONHARDENINGRETURN}(\sigma^k, \delta\varepsilon, X^k, K^k, G^k, (\overline{p^w})^k,$

    ↪ $\mathbf{s}^{\text{trial}}, \sqrt{J_2^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta)$

    ↪   *▷Compute return to updated yield surface (no hardening)*

14:   isSuccess, $\sigma^{k+1}, \varepsilon^{\text{p},k+1}, \alpha^{k+1}, (\overline{p^w})^{k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1}, K^{k+1}, G^{k+1}, \mathbf{s}^{k+1}, (\overline{p^w})^{k+1}, I_1^{\text{eff},k+1}, \sqrt{J_2^{k+1}},$

    ↪ $r^{k+1}, z_{\text{eff}}^{k+1}, \varepsilon_v^{\text{p},k+1} \leftarrow \text{CONSISTENCYBISECTION}(\delta\varepsilon, \varepsilon^{\text{p},k}, \sigma^k, K^k, G^k, (\overline{p^w})^k, \phi^k, S_w^k, X^k,$

    ↪ $\mathbf{s}^{\text{trial}}, I_1^{\text{eff,trial}}, \sqrt{J_2^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, \varepsilon_v^{\text{p,trial}}, p_3^{\text{trial}}, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta, i_{\max}, j_{\max},$

    ↪ $\sigma_{\text{fixed}}, \delta\varepsilon_{\text{fixed}}^{\text{p}})$ *▷The bisection return algorithm to take care of yield surface hardening.*

15:   **if** iSuccess = FALSE **then**

16:     **return** isSuccess, $\sigma^k, \varepsilon^{\text{p},k}, \phi^k, S_w^k, X^k, \alpha^k, K^k, G^k, p_3^k$

17:   **end if**

18:   **return** isSuccess, $\sigma^{k+1}, \varepsilon^{\text{p},k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1}, \alpha^{k+1}, K^{k+1}, G^{k+1}, p_3^{k+1}$

19: **end procedure**

The nonhardening return algorithm

The nonhardening return algorithm uses a transformed space (see [3]) where the computation is carried out in special Lode coordinates (z_{eff}, r') where

$$z_{\text{eff}} := \frac{\text{tr}(\sigma - \alpha)}{\sqrt{3}} \quad \text{and} \quad r' = \beta r \sqrt{\frac{3K}{2G}}, \quad r := \sqrt{2J_2}. \tag{5.1}$$

If the flow rule is non-associative, the yield surface parameter $\beta \neq 1$.

The nonhardening return algorithm pseudocode is listed below:

Algorithm 30 Non-hardening return algorithm

Require: $\sigma^k, \delta\varepsilon, X^k, K^k, G^k, (\overline{p^w})^k, \mathbf{s}^{\text{trial}}, \sqrt{J_2^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta,$ `yieldCondition`

1: **procedure** NONHARDENINGRETURN

2:   $r'_{\text{trial}} \leftarrow \beta \, r^{\text{trial}} \sqrt{\frac{3K^k}{2G^k}}$ *▷Transform the trial r coordinate*

3:   $X_{\text{eff}}^k \leftarrow X^k + 3(\overline{p^w})^k$

4:   $z_{\text{eff}}^{\text{close}}, r'_{\text{close}} \leftarrow$ `yieldCondition`.GETCLOSESTPOINT$(K^k, G^k, X_{\text{eff}}^k, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta,$
    ↪ $z_{\text{eff}}^{\text{trial}}, r'_{\text{trial}})$

5:   $I_1^{\text{close}} \leftarrow \sqrt{3} z_{\text{eff}}^{\text{close}} - 3(\overline{p^w})^k, \sqrt{J_2^{\text{close}}} \leftarrow \frac{1}{\beta}\sqrt{\frac{G^k}{3K^k}} r'_{\text{close}}$

6:   **if** $\sqrt{J_2^{\text{trial}}} > 0$ **then**

7:     $\sigma^{\text{fixed}} = \frac{1}{3} I_1^{\text{close}} \mathbf{I} + \frac{\sqrt{J_2^{\text{close}}}}{\sqrt{J_2^{\text{trial}}}} \mathbf{s}^{\text{trial}}$ *▷Compute updated total stress*

8:   **else**

9:     $\sigma^{\text{fixed}} = \frac{1}{3} I_1^{\text{close}} \mathbf{I} + \mathbf{s}^{\text{trial}}$ *▷Compute updated total stress when the trial stress is hydrostatic*

10:   **end if**

11:   $\delta\sigma_{\text{fixed}} \leftarrow \sigma^{\text{fixed}} - \sigma^k$ *▷Compute stress increment*

12:   $\delta\sigma_{\text{fixed}}^{\text{iso}} \leftarrow \frac{1}{3}\text{tr}(\delta\sigma_{\text{fixed}})\mathbf{I}, \quad \delta\sigma_{\text{fixed}}^{\text{dev}} \leftarrow \delta\sigma_{\text{fixed}} - \delta\sigma_{\text{fixed}}^{\text{iso}}$

13:   $\delta\varepsilon^{\text{p,fixed}} = \delta\varepsilon - \frac{1}{3K^k}\delta\sigma_{\text{fixed}}^{\text{iso}} - \frac{1}{2G^k}\delta\sigma_{\text{fixed}}^{\text{dev}}$ *▷Compute plastic strain increment*

14:     **return** $\boldsymbol{\sigma}^{\text{fixed}}, \delta\boldsymbol{\varepsilon}^{\text{p,fixed}}$
15: **end procedure**

Finding the closest point in transformed space

Algorithm 31 Compute the closest point from the trial state to transformed non-hardening yield surface

---

Require: $K^k, G^k, X_{\text{eff}}^k, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta, z_{\text{eff}}^{\text{trial}}, r_{\text{trial}}'$

1: **procedure** GETCLOSESTPOINT
2:     $n_{\text{poly}} \leftarrow 5$
3:     Set up $I_1^{\text{max}}, I_1^{\text{min}}, I_1^{\text{mid}} = 1/2(I_1^{\text{max}} + I_1^{\text{min}})$ limits of the yield surface.
4:     Set up bisection: $\eta_{\text{low}} \leftarrow 0, \eta_{\text{high}} \leftarrow 1, \eta_{\text{mid}} \leftarrow 1/2(\eta_{\text{low}} + \eta_{\text{high}})$
5:     **while** ABS($\eta^{\text{high}} - \eta^{\text{low}}$) > TOLERANCE **do**
6:         $\mathbf{x}_{\text{poly}} \leftarrow$ GETYIELDSURFACEPOINTSALL_RPRIMEZ($n_{\text{poly}}, K^k, G^k, X_{\text{eff}}^k,$
              $\hookrightarrow a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta$)
              $\hookrightarrow$ ▷*Get the polygon that represents the yield surface in z_{eff}-r' space.*
7:         $\mathbf{x}_{\text{close}} \leftarrow$ FINDCLOSESTPOINT($z_{\text{eff}}^{\text{trial}}, r_{\text{trial}}', \mathbf{x}_{\text{poly}}$)
              $\hookrightarrow$ ▷*Find the closest point in the discretized segments to the trial stress state.*
8:         Compute I_1^{close} from $\mathbf{x}_{\text{close}}$
9:         **if** $I_1^{\text{close}} < I_1^{\text{mid}}$ **then**
10:             $I_1^{\text{max}} \leftarrow I_1^{\text{mid}}, \eta_{\text{high}} \leftarrow \eta_{\text{mid}}$ ▷*Update mid point*
11:         **else**
12:             $I_1^{\text{min}} \leftarrow I_1^{\text{mid}}, \eta_{\text{low}} \leftarrow \eta_{\text{mid}}$ ▷*Update mid point*
13:         **end if**
14:         Recompute $I_1^{\text{mid}}, \eta_{\text{mid}}$ and update old closest point.
15:     **end while**
16:     **return** isSuccess = TRUE, $\mathbf{x}_{\text{close}}.z_{\text{eff}}$, $\mathbf{x}_{\text{close}}.r'$
17: **end procedure**

Finding the yield surface polygon in z_{eff}-r' space

Algorithm 32 Find points in a closed polygon that describes the yield surface in z_{eff}-r' space

---

Require: $n_{\text{poly}}, K^k, G^k, X_{\text{eff}}^k, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta$

1: **procedure** GETYIELDSURFACEPOINTSALL_RPRIMEZ
2:     $\kappa \leftarrow I_1^{\text{peak}} - R_c(I_1^{\text{peak}} - X_{\text{eff}}^k)$ ▷*Compute κ.*
3:     $I_1^{\text{eff}} \leftarrow$ LINSPACE($\texttt{from} = X_{\text{eff}}^k, \texttt{to} = I_1^{\text{peak}}, \texttt{points} = n_{\text{poly}}$)
              $\hookrightarrow$ ▷*Create an equally spaced set of I_1^{eff} values.*
4:     **for** I_1 in I_1^{eff} **do**
5:         $F_f = a_1 - a_3 \exp(a_2 I_1) - a_4 I_1;$ ▷*Compute F_f.*
6:         $F_c^2 \leftarrow 1$
7:         **if** $I_1 < \kappa$ and $X_{\text{eff}}^k < I_1$ **then**
8:             $F_c^2 = 1 - \left[\frac{\kappa - I_1}{\kappa - X_{\text{eff}}^k} \right]^2;$ ▷*Compute F_c.*
9:         **end if**
10:         $J_2 = F_f^2 F_c^2$ ▷*Compute J_2 and push into a vector*
11:     **end for**
12:     $z_{\text{eff}} \leftarrow I_1^{\text{eff}}/\sqrt{3}$, $r' \leftarrow \beta\sqrt{\frac{3K^k}{2G^k}}\sqrt{2J_2}$
13:     $\mathbf{x}_{\text{poly}}.z_{\text{eff}} \leftarrow z_{\text{eff}} \cup$ REVERSE(z_{eff}) , $\mathbf{x}_{\text{poly}}.r' \leftarrow r' \cup$ REVERSE($-r'$)
              $\hookrightarrow$ ▷*Add the points on the negative r' side of the polygon*
14:     $\mathbf{x}_{\text{poly}}[2n_{\text{poly}} + 1] \leftarrow \mathbf{x}_{\text{poly}}[1]$ ▷*Add the first point to close the polygon*
15:     **return** $\mathbf{x}_{\text{poly}}$
16: **end procedure**

Finding the closest point on yield surface in z_{eff}-r' space

Algorithm 33 Find the closest point from the trial stress state on the polyline describing the yield surface

Require: $\mathbf{x}_{\text{trial}}.z_{\text{eff}}$, $\mathbf{x}_{\text{trial}}.r'$, $\mathbf{x}_{\text{poly}}$
1: **procedure** FINDCLOSESTPOINT
2:　　　$i \leftarrow 0$
3:　　　**for** $\{\mathbf{x}_{\text{start}}, \mathbf{x}_{\text{end}}\}$ **in** $\mathbf{x}_{\text{poly}}$ **do**
4:　　　　　$\mathbf{x}_{\text{seg}} \leftarrow \mathbf{x}_{\text{end}} - \mathbf{x}_{\text{start}}$
5:　　　　　$\mathbf{x}_{\text{proj}} \leftarrow \mathbf{x}_{\text{trial}} - \mathbf{x}_{\text{start}}$
6:　　　　　$t \leftarrow \frac{\mathbf{x}_{\text{proj}} \cdot \mathbf{x}_{\text{seg}}}{\|\mathbf{x}_{\text{seg}}\|}$
7:　　　　　$i \leftarrow i + 1$
8:　　　　　**if** $t < 0$ **then**
9:　　　　　　　$\mathbf{x}[i] \leftarrow \mathbf{x}_{\text{start}}$
10:　　　　　**else if** $t > 1$ **then**
11:　　　　　　　$\mathbf{x}[i] \leftarrow \mathbf{x}_{\text{end}}$
12:　　　　　**else**
13:　　　　　　　$\mathbf{x}[i] \leftarrow \mathbf{x}_{\text{start}} + t\,\mathbf{x}_{\text{seg}}$
14:　　　　　**end if**
15:　　　**end for**
16:　　　$d^2_{\min} \leftarrow \text{DOUBLE\_MAX}$
17:　　　$\mathbf{x}_{\text{close}} \leftarrow \mathbf{0}$
18:　　　**for** $\mathbf{x}_i$ **in** $\mathbf{x}$ **do**
19:　　　　　$d^2 \leftarrow \text{DISTANCESQ}(\mathbf{x}_i, \mathbf{x}_{\text{trial}})$
20:　　　　　**if** $d^2 < d^2_{\min}$ **then**
21:　　　　　　　$d^2_{\min} \leftarrow d^2$
22:　　　　　　　$\mathbf{x}_{\text{close}} \leftarrow \mathbf{x}_i$
23:　　　　　**end if**
24:　　　**end for**
25:　　　**return** $\mathbf{x}_{\text{close}}$
26: **end procedure**

Consistency bisection algorithm

Algorithm 34 The consistency bisection algorithm for partially saturated materials

Require: $\delta\boldsymbol{\varepsilon}$, $\boldsymbol{\varepsilon}^{\text{p},k}$, $\boldsymbol{\sigma}^k$, K^k, G^k, $(\overline{p^w})^k$, ϕ^k, S_w^k, X^k, $\boldsymbol{s}^{\text{trial}}$, $I_1^{\text{eff,trial}}$, $\sqrt{J_2^{\text{trial}}}$, r^{trial}, $z_{\text{eff}}^{\text{trial}}$, $\varepsilon_v^{\text{p,trial}}$, p_3^{trial}, a_1, a_2, a_3,
　　　a_4, I_1^{peak}, R_c, β, $i_{\max}$, $j_{\max}$, $\boldsymbol{\sigma}_{\text{fixed}}$, $\delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p}}$, `yieldCondition`
1: **procedure** CONSISTENCYBISECTION
2:　　　$\delta\varepsilon_v^{\text{p,fixed}} \leftarrow \text{tr}(\delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p}})$
3:　　　$i \leftarrow 1$
4:　　　$\eta^{\text{in}} \leftarrow 0, \eta^{\text{out}} \leftarrow 1$
5:　　　**while** ABS$(\eta^{\text{out}} - \eta^{\text{in}}) > \text{TOLERANCE}$ **do**
6:　　　　　$j \leftarrow 1$
7:　　　　　isElastic $\leftarrow$ TRUE
8:　　　　　**while** isElastic = TRUE **do**
9:　　　　　　　$\eta^{\text{mid}} \leftarrow \frac{1}{2}(\eta^{\text{in}} + \eta^{\text{out}})$
10:　　　　　　　$\delta\varepsilon_v^{\text{p,mid}} \leftarrow \eta^{\text{mid}}\delta\varepsilon_v^{\text{p,fixed}}$
11:　　　　　　　$(\overline{p^w})^{\text{mid}}, \phi^{\text{mid}}, S_w^{\text{mid}}, X^{\text{mid}} \leftarrow \text{COMPUTEINTERNALVARIABLES}(K^k, G^k, (\overline{p^w})^k, \phi^k, S_w^k,$
　　　　　　　　　$\hookrightarrow X^k, \delta\varepsilon_v^{\text{p,mid}})$　　　$\triangleright$*Update the internal variables using the bisected increment*
　　　　　　　$\hookrightarrow$ *of the volumetric plastic strain*

12:        isElastic $\leftarrow$ `yieldCondition`.EVALYIELDCONDITION($I_1^{\text{eff,trial}}, \sqrt{J_2^{\text{trial}}}, X^{\text{mid}}, (\overline{p^w})^{\text{mid}},$
             $\hookrightarrow \ \phi^{\text{mid}}, S_w^{\text{mid}}, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta)$
             $\hookrightarrow$ *▷Determine whether the trial stress is elastic or not*

13:       **if** isElastic = TRUE **then**

14:          $\eta^{\text{out}} \leftarrow \eta^{\text{mid}}$ *▷If the local trial state is inside the updated yield surface, the yield*
                      $\hookrightarrow$ condition evaluates to "elastic". We need to reduce the size of the
                      $\hookrightarrow$ yield surface by decreasing the plastic strain increment.

15:          $j \leftarrow j + 1$

16:          **if** $j \geq j_{\max}$ **then**

17:              **return** isSuccess $\leftarrow$ FALSE   *▷The bisection algorithm failed because of*
                      $\hookrightarrow$ too many iterations.

18:          **end if**

19:       **end if**

20:     **end while**

21:     $\boldsymbol{\sigma}_{\text{fixed}}^{\text{new}}, \delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p,new}} \leftarrow$ NONHARDENINGRETURN($\boldsymbol{\sigma}^k, \delta\boldsymbol{\varepsilon}, X^{\text{mid}}, K^k, G^k, (\overline{p^w})^{\text{mid}},$
         $\hookrightarrow \ \boldsymbol{s}^{\text{trial}}, \sqrt{J_2^{\text{trial}}}, r^{\text{trial}}, z_{\text{eff}}^{\text{trial}}, a_1, a_2, a_3, a_4, I_1^{\text{peak}}, R_c, \beta)$
         $\hookrightarrow$ *▷Compute return to updated yield surface (no hardening)*

22:     **if** SIGN$(\text{tr}(\boldsymbol{\sigma}_{\text{trial}} - \boldsymbol{\sigma}_{\text{fixed}}^{\text{new}})) \neq$ SIGN$(\text{tr}(\boldsymbol{\sigma}_{\text{trial}} - \boldsymbol{\sigma}_{\text{fixed}}))$ **then**

23:       $\eta^{\text{out}} \leftarrow \eta^{\text{mid}}$ *▷Too much plastic strain*

24:       **continue**

25:     **end if**

26:     **if** $\left\| \delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p,new}} \right\| > \eta^{\text{mid}} \left\| \delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p}} \right\|$ **then**

27:       $\eta^{\text{in}} \leftarrow \eta^{\text{mid}}$ *▷Too little plastic strain*

28:     **else**

29:       $\eta^{\text{out}} \leftarrow \eta^{\text{mid}}$ *▷Too much plastic strain*

30:     **end if**

31:     $i \leftarrow i + 1$

32:     **if** $i \geq i_{\max}$ **then**

33:       **return** isSuccess $\leftarrow$ FALSE   *▷Too many iterations*

34:     **end if**

35:   **end while**

36:   $\delta\varepsilon_{v,\text{fixed}}^{\text{p},k+1} \leftarrow \text{tr}(\delta\boldsymbol{\varepsilon}_{\text{fixed}}^{\text{p},k+1})$

37:   $(\overline{p^w})^{k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1} \leftarrow$ COMPUTEINTERNALVARIABLES($K^k, G^k, (\overline{p^w})^k, \phi^k, S_w^k,$
     $\hookrightarrow X^k, \delta\varepsilon_{v,\text{fixed}}^{\text{p},k+1})$ *▷Update the internal variables using the bisected increment*
   $\hookrightarrow$ of the volumetric plastic strain

38:   $\boldsymbol{\sigma}^{k+1} \leftarrow \boldsymbol{\sigma}_{\text{fixed}}^{\text{new}}, \ \boldsymbol{\alpha}^{k+1} \leftarrow -(\overline{p^w})^{k+1}\boldsymbol{I}, \ p_3^{k+1} \leftarrow p_3^{\text{trial}}$

39:   $\boldsymbol{\varepsilon}^{\text{p},k+1} = \boldsymbol{\varepsilon}^{\text{p},k} + \delta\boldsymbol{\varepsilon}^{\text{p},k+1}$ *▷Update the plastic strain*

40:   $K^{k+1}, G^{k+1}, \boldsymbol{s}^{k+1}, (\overline{p^w})^{k+1}, I_1^{\text{eff},k+1}, \sqrt{J_2^{k+1}}, r^{k+1}, z_{\text{eff}}^{k+1}, \varepsilon_v^{\text{p},k+1} \leftarrow$
     $\hookrightarrow$ COMPUTEELASTICPROPERTIES($\boldsymbol{\sigma}^{k+1}, \phi^{k+1}, S_w^{k+1}, \boldsymbol{\varepsilon}^{\text{p},k+1}, \boldsymbol{\alpha}^{k+1}, p_3^{k+1}$)
     $\hookrightarrow$ *▷Compute elastic moduli and stress invariants for the new state*

41:   **return** isSuccess $\leftarrow$ TRUE, $\boldsymbol{\sigma}^{k+1}, \boldsymbol{\varepsilon}^{\text{p},k+1}, \boldsymbol{\alpha}^{k+1}, (\overline{p^w})^{k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1}, K^{k+1}, G^{k+1},$
     $\hookrightarrow \boldsymbol{s}^{k+1}, (\overline{p^w})^{k+1}, I_1^{\text{eff},k+1}, \sqrt{J_2^{k+1}}, r^{k+1}, z_{\text{eff}}^{k+1}, \varepsilon_v^{\text{p},k+1}, p_3^{k+1}$

42: **end procedure**

Updating the internal variables

Algorithm 35 Updating the internal variables for partially saturated materials

Require: $\boldsymbol{\sigma}^k, \varepsilon_v^{\text{p},k}, K^k, G^k, (\overline{p^w})^k, \phi^k, S_w^k, X^k, \delta\varepsilon_v^{\text{p}}$, `fluidParams`, `crushParams`, `airModel`, `waterModel`

1:   **procedure** COMPUTEINTERNALVARIABLES

2:     $\overline{\varepsilon_v^{\text{p},k}} \leftarrow -\varepsilon_v^{\text{p},k}, \delta\overline{\varepsilon_v^{\text{p}}} \leftarrow -\delta\varepsilon_v^{\text{p}}$

3: $\quad \overline{p_o^w} \leftarrow \texttt{fluidParams}.\overline{p_o^w}, S_o \leftarrow \texttt{fluidParams}.S_o, \phi_o \leftarrow \texttt{fluidParams}.\phi_o, p_1^{\text{sat}} \leftarrow \texttt{crushParams}.p_1^{\text{sat}}$

4: $\quad K_a \leftarrow \texttt{airModel}.\textsc{computeBulkModulus}((\overline{p^w})^k)$

5: $\quad K_w \leftarrow \texttt{waterModel}.\textsc{computeBulkModulus}((\overline{p^w})^k)$

6: $\quad \varepsilon_v^{a,o} \leftarrow \texttt{airModel}.\textsc{computeElasticVolumetricStrain}(\overline{p_o^w})$

7: $\quad \varepsilon_v^{a} \leftarrow \texttt{airModel}.\textsc{computeElasticVolumetricStrain}((\overline{p^w})^k)$

8: $\quad \varepsilon_v^{w} \leftarrow \texttt{waterModel}.\textsc{computeElasticVolumetricStrain}((\overline{p^w})^k, \overline{p_o^w})$

9: $\quad \overline{\varepsilon_v^{a}} \leftarrow -(\varepsilon_v^a - \varepsilon_v^{a,o}), \overline{\varepsilon_v^{w}} \leftarrow -\varepsilon_v^w$

10: $\quad \mathcal{C}_p \leftarrow S_o \exp(\overline{\varepsilon_v^a} - \overline{\varepsilon_v^w})$

11: $\quad \mathcal{D}_p \leftarrow \frac{1 - S_o}{(1 - S_o + \mathcal{C}_p)^2}$

12: $\quad \dfrac{d\mathcal{C}_p}{d\overline{p^w}} \leftarrow \mathcal{C}_p \left[\dfrac{1}{K_a} - \dfrac{1}{K_w} \right]$

13: $\quad \mathcal{G}_a \leftarrow \exp(\overline{\varepsilon_v^{p,k}} - \overline{\varepsilon_v^a}), \ \mathcal{G}_w \leftarrow \exp(\overline{\varepsilon_v^{p,k}} - \overline{\varepsilon_v^w})$

14: $\quad \mathcal{B}_p \leftarrow \dfrac{1}{(1 - S_o)\mathcal{G}_a + S_o \mathcal{G}_w} \left[-\dfrac{(1-\phi^k)\phi^k}{\phi_o} \left(\dfrac{S_w^k}{K_w} + \dfrac{1 - S_w^k}{K_a} \right) + \dfrac{1 - S_o}{K_a} \mathcal{G}_a + \dfrac{S_o}{K_w} \mathcal{G}_w \right]$

15: $\quad (\overline{p^w})^{k+1} \leftarrow \textsc{max} \left[(\overline{p^w})^k + \dfrac{1}{\mathcal{B}_p} \delta\varepsilon_v^p, o \right] \quad$ ▷*Update the pore pressure making sure that pressure does*
$\quad\quad \hookrightarrow$ *not become negative during dilatative plastic deformations.*

16: $\quad \overline{X}_d, \dfrac{d\overline{X}_d}{d\varepsilon_v^p} \leftarrow \textsc{computeDrainedHydrostaticStrengthAndDeriv}(\overline{\varepsilon_v^{p,k}})$
$\quad\quad \hookrightarrow$ ▷*Compute the drained hydrostatic compressive strength and its derivative*

17: $\quad \overline{X}^{k+1} = -X^k + \left[\left(1 - S_w^k + p_1^{\text{sat}} S_w^k \right) \dfrac{d\overline{X}_d}{d\varepsilon_v^p} + \overline{X}_d \left(p_1^{\text{sat}} - 1\right) \dfrac{\mathcal{D}_p}{\mathcal{B}_p} \dfrac{d\mathcal{C}_p}{d\overline{p^w}} + \dfrac{3}{\mathcal{B}_p} \right] \delta\varepsilon_v^p$
$\quad\quad \hookrightarrow$ ▷*Update the hydrostatic compressive strength*

18: $\quad X^{k+1} \leftarrow -\overline{X}^{k+1}$

19: $\quad \varepsilon_v^{p,k+1} \leftarrow \varepsilon_v^{p,k} + \delta\varepsilon_v^p \quad\quad\quad\quad\quad\quad\quad\quad\quad$ ▷*Compute the updated volumetric plastic strain.*

20: $\quad \varepsilon_v^{a,k+1} \leftarrow \texttt{airModel}.\textsc{computeElasticVolumetricStrain}((\overline{p^w})^{k+1})$

21: $\quad \varepsilon_v^{w,k+1} \leftarrow \texttt{waterModel}.\textsc{computeElasticVolumetricStrain}((\overline{p^w})^{k+1}, \overline{p_o^w})$

22: $\quad \overline{\varepsilon_v^{a,k+1}} \leftarrow -(\varepsilon_v^{a,k+1} - \varepsilon_v^{a,o}), \overline{\varepsilon_v^{w,k+1}} \leftarrow -\varepsilon_v^{w,k+1} \quad$ ▷*The updated strains in the fluid phases.*

23: $\quad \mathcal{C}_p^{k+1} \leftarrow S_o \exp(\overline{\varepsilon_v^{a,k+1}} - \overline{\varepsilon_v^{w,k+1}})$

24: $\quad S_w^{k+1} \leftarrow \dfrac{\mathcal{C}_p^{k+1}}{1 - S_o + \mathcal{C}_p^{k+1}} \quad\quad\quad\quad\quad\quad\quad\quad\quad\quad$ ▷*Update the saturation*

25: $\quad \mathcal{G}_a^{k+1} \leftarrow \exp(\overline{\varepsilon_v^{p,k+1}} - \overline{\varepsilon_v^{a,k+1}}), \ \mathcal{G}_w^{k+1} \leftarrow \exp(\overline{\varepsilon_v^{p,k+1}} - \overline{\varepsilon_v^{w,k+1}})$

26: $\quad \phi^{k+1} \leftarrow (1 - S_o)\phi_o \mathcal{G}_a^{k+1} + S_o \phi_o \mathcal{G}_w^{k+1} \quad\quad\quad\quad\quad$ ▷*Update the porosity*

27: $\quad$ **return** $(\overline{p^w})^{k+1}, \phi^{k+1}, S_w^{k+1}, X^{k+1}$

28: **end procedure**

Algorithm 36 Computing the drained hydrostatic strength and its derivative

---

Require: $\overline{\varepsilon_v^{p,k}}$, $\texttt{fluidParams}$, $\texttt{crushParams}$

1: **procedure** computeDrainedHydrostaticStrengthAndDeriv

2: $\quad \phi_o \leftarrow \texttt{fluidParams}.\phi_o$

3: $\quad p_o \leftarrow \texttt{crushParams}.p_o, p_1 \leftarrow \texttt{crushParams}.p_1, p_1^{\text{sat}} \leftarrow \texttt{crushParams}.p_1^{\text{sat}}, p_2 \leftarrow \texttt{crushParams}.p_2$

4: $\quad p_3 \leftarrow -\log(1 - \phi_o)$

5: $\quad \overline{X}_d \leftarrow \textsc{max}(p_o, 1000); \quad\quad\quad\quad\quad\quad\quad$ ▷$\overline{X}_d$ *has a minimum value of 1000 pressure units*

6: $\quad \dfrac{d\overline{X}_d}{d\varepsilon_v^p} \leftarrow o$

7: $\quad$ **if** $\overline{\varepsilon_v^{p,k}} > o$ **then**

8: $\quad\quad \phi_{\text{temp}} \leftarrow \exp(-p_3 + \overline{\varepsilon_v^{p,k}})$

9: $\quad\quad \phi \leftarrow 1 - \phi_{\text{temp}}$

10: $\quad\bar{\xi} \leftarrow p_1 \, \text{POW}\!\left(\frac{\phi_o}{\phi} - 1, \frac{1}{p_2}\right)$

11: $\quad\overline{X}_d \leftarrow \overline{X}_d + \bar{\xi}$

12: $\quad\dfrac{d\overline{X}_d}{d\varepsilon_v^{\text{p}}} \leftarrow \frac{1}{p_2} \, \frac{\phi_o}{\phi} \, \phi_{\text{temp}} \, \dfrac{\bar{\xi}}{\phi\left(\frac{\phi_o}{\phi} - 1\right)}$

13: **end if**

14: **return** $\overline{X}_d, \dfrac{d\overline{X}_d}{d\varepsilon_v^{\text{p}}}$

15: **end procedure**

5.2.3 Rate-dependent plastic update

Our implementation does not consider rate-dependent updates of the internal variables. We approximate the trial stress using the average of the elastic moduli at the start and end of the step.

Algorithm 37 Computing the correction to the stress due to rate-dependent plasticity

Require: $\Delta t, \, \boldsymbol{d}^{n+1}, \, \boldsymbol{\sigma}^n, \, K^n, \, G^n, \, \phi^n, \, S_w^n, \, X^n, \, \boldsymbol{\alpha}^n, \, \boldsymbol{\epsilon}^{\text{p},n}, \, p_3^n, \, \boldsymbol{\sigma}_{qs}^n, \, \boldsymbol{\sigma}_{qs}^{n+1}, \, K^{n+1}, \, G^{n+1}, \, \phi^{n+1}, \, S_w^{n+1}, \, X^{n+1}, \, \boldsymbol{\alpha}^{n+1},$
$\quad\boldsymbol{\epsilon}^{\text{p},n+1}, \, p_3^{n+1}, \, a_1, \, a_2, \, a_3, \, a_4, \, I_1^{\text{peak}}, \, R_c, \, \texttt{yieldParams}$

1: **procedure** RATEDEPENDENTPLASTICUPDATE

2: $\quad T_1 \leftarrow \texttt{yieldParams}.T_1, \, T_2 \leftarrow \texttt{yieldParams}.T_2$

3: $\quad$**if** $T_1 = 0$ **or** $T_2 = 0$ **then** $\qquad\qquad\qquad$ ▷*Check if rate-dependent plasticity has been turned on*

4: $\qquad$**return** isRateDependent $\leftarrow$ FALSE, $\boldsymbol{\sigma}_{qs}^{n+1}$

5: $\quad$**end if**

6: $\quad K_{\text{dyn}} \leftarrow \frac{1}{2}\left(K^n + K^{n+1}\right), \, G_{\text{dyn}} \leftarrow \frac{1}{2}\left(G^n + G^{n+1}\right)$ $\qquad$ ▷*Compute mid-step bulk and shear modulus*

7: $\quad\Delta\boldsymbol{\varepsilon} \leftarrow \Delta t \, \boldsymbol{d}^{n+1}$

8: $\quad\boldsymbol{\sigma}_{\text{trial,dyn}} \leftarrow \text{COMPUTETRIALSTRESS}(\boldsymbol{\sigma}^n, K_{\text{dyn}}, G_{\text{dyn}}, \boldsymbol{d}^{n+1}, \Delta t)$ $\qquad$ ▷*Compute substep trial stress*

9: $\quad\dot{\varepsilon} \leftarrow \text{MAX}(\|\boldsymbol{d}^{n+1}\|, \text{ABS\_DOUBLE\_MIN})$

10: $\quad\tau \leftarrow T_1 \, \text{POW}(\dot{\varepsilon}, T_2)$ $\qquad\qquad\qquad$ ▷*The characteristic time is defined from the rate-dependence*
$\qquad\qquad\qquad\qquad\qquad$ ↪ *input parameters and the magnitude of the strain rate*

11: $\quad r_h \leftarrow \exp\!\left(-\frac{\Delta t}{\tau}\right)$

12: $\quad R_H \leftarrow \frac{1 - r_h}{\frac{\Delta t}{\tau}}$

13: $\quad\boldsymbol{\sigma}^{n+1} \leftarrow \boldsymbol{\sigma}_{qs}^{n+1} + \left[\left(\boldsymbol{\sigma}_{\text{trial,dyn}} - \boldsymbol{\sigma}^n\right) - \left(\boldsymbol{\sigma}_{qs}^{n+1} - \boldsymbol{\sigma}_{qs}^n\right)\right] R_H + \left(\boldsymbol{\sigma}^n - \boldsymbol{\sigma}_{qs}^n\right) r_h$ $\qquad$ ▷*Stress update*

14: $\quad$**return** $\boldsymbol{\sigma}^{n+1}$, isRateDependent $\leftarrow$ TRUE

15: **end procedure**
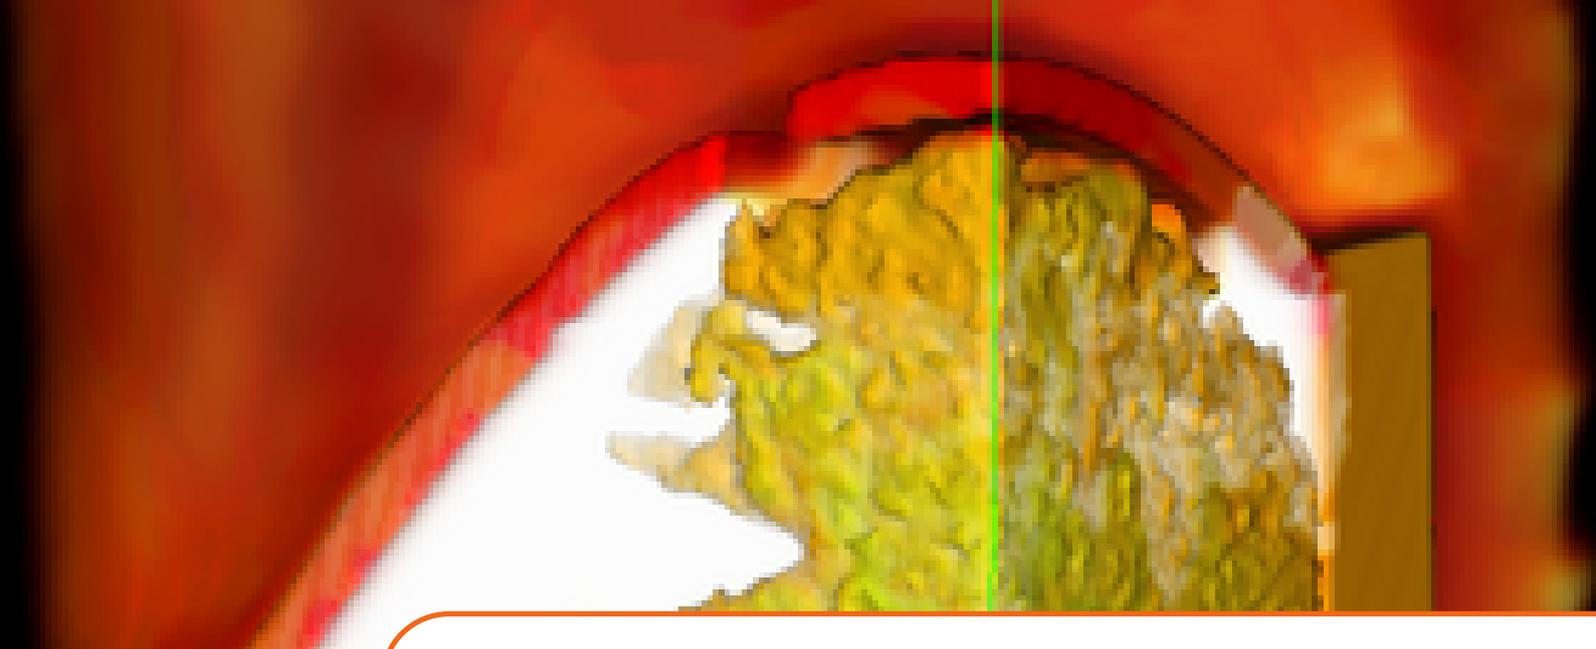

6 — Submodels of MPM material models

The previous chapter on the ARENA model showed that MPM constitutive models can rapidly become unmanageable without sufficient object granularity. To allow for simpler model development and testing processes to be used in VAANGO , constitutive models are sometime broken up into submodels.

In VAANGO , we have the main constitutive models are located in the directory src/CCA/Components/MPM/ConstitutiveModel . Some of the models in this directory depend of one of two sets of submodels. The first set of submodels were designed circa 2004-2005 and are located in the directory ...ConstitutiveModel/PlasticityModels . These models have been retained largely for backward compatibility with older plasticity models. A more recent set of models (designed circa 2012) can be found in ...ConstitutiveModel/Models .
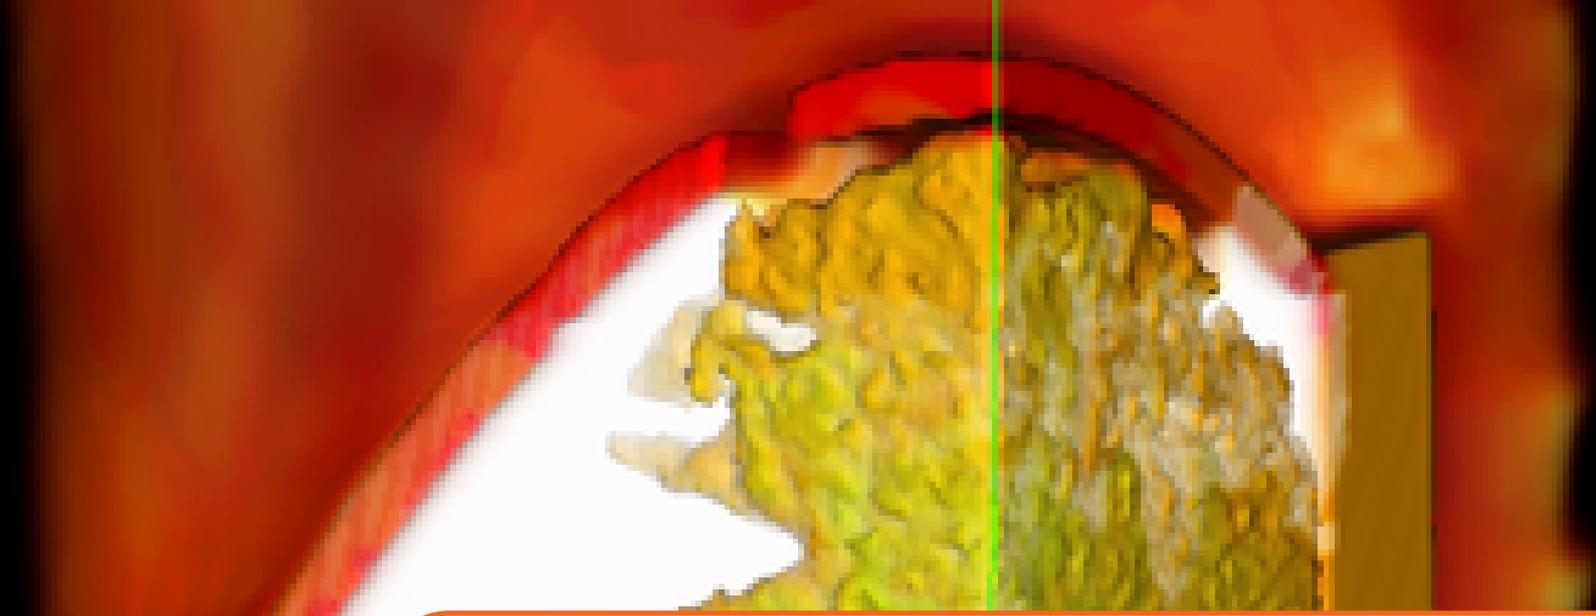
7 — Metal plasticity models in MPM

The metal plasticity models in MPM have been designed with the intention of making the replacement of one set of sub-models with another via the input file mechanism. In this chapter we discuss the basic ideas behind a typical metal plasticity implementation and how new models can be added.

7.1 Adding new models

In the parallel implementation of the stress update algorithm, sockets have been added to allow for the incorporation of a variety of plasticity, damage, yield, and bifurcation models without requiring any change in the stress update code. The algorithm is shown in Algorithm 7.1. The equation of state, plasticity model, yield condition, damage model, and the stability criterion are all polymorphic objects created using a factory idiom in C++ ([4]). The algorithm below is just an example and may not represent any particular algorithm currently implemented in VAANGO in all its details.

Addition of a new model requires the following steps (the example below is only for the flow stress model but the same idea applies to other models) :

1. Creation of a new class that encapsulates the plasticity model. The template for this class can be copied from the existing plasticity models. The data that is unique to the new model are specified in the form of
   - A structure containing the constants for the plasticity model.
   - Particle variables that specify the variables that evolve in the plasticity model.
2. The implementation of the plasticity model involves the following steps.
   - Reading the input file for the model constants in the constructor.
   - Adding the variables that evolve in the plasticity model appropriately to the task graph.
   - Adding the appropriate flow stress calculation method.
3. The `PlasticityModelFactory` is then modified so that it recognizes the added plasticity model.

7.2 Implementation

The elastic response is assumed to be isotropic. The material constants that are taken as input for the elastic response are the bulk and shear modulus. The flow rule is determined from the input and the appropriate plasticity model is created using the `PlasticityModelFactory` class. The damage evolution rule is determined from the input and a damage model is created using the `DamageModelFactory` class. The

Table 7.1: Stress Update Algorithm

Persistent:Initial moduli, temperature, porosity,
        scalar damage, equation of state, plasticity model,
        yield condition, stability criterion, damage model
Temporary:Particle state at time t
Output: Particle state at time $t + \Delta t$

For *all the patches in the domain*
    Read the particle data and initialize updated data storage
    For *all the particles in the patch*
        Compute the velocity gradient and the rate of deformation tensor
        Compute the deformation gradient and the rotation tensor
        Rotate the Cauchy stress and the rate of deformation tensor
            to the material configuration
        Compute the current shear modulus and melting temperature
        Compute the pressure using the equation of state,
            update the hydrostatic stress, and
            compute the trial deviatoric stress
        Compute the flow stress using the plasticity model
        Evaluate the yield function
        If *particle is elastic*
            Update the elastic deviatoric stress from the trial stress
            Rotate the stress back to laboratory coordinates
            Update the particle state
        Else
            Compute the elastic-plastic deviatoric stress
            Compute updated porosity, scalar damage, and
                temperature increase due to plastic work
            Compute elastic-plastic tangent modulus and evaluate stability condition
            Rotate the stress back to laboratory coordinates
            Update the particle state
        End If
        If *Temperature > Melt Temperature* or *Porosity > Critical Porosity* or *Unstable*
            Tag particle as failed
        End If
        Convert failed particles into a material with a different velocity field
    End For
End For

equation of state that is used to determine the pressure is also determined from the input. The equation of state model is created using the MPMEquationOfStateFactory class.

In addition, a damage evolution variable (D) is stored at each time step (this need not be the case and will be transfered to the damage models in the future). The left stretch and rotation are updated incrementally at each time step (instead of performing a polar decomposition) and the rotation tensor is used to rotate the Cauchy stress and rate of deformation to the material coordinates at each time step (instead of using a objective stress rate formulation).

Any evolution variables for the plasticity model, damage model or the equation of state are specified in the class that encapsulates the particular model.

The flow stress is calculated from the plasticity model using a function call of the form

```
double flowStress = d_plasticity->computeFlowStress(tensorEta, tensorS,
                                                pTemperature[idx],
                                                delT, d_tol, matl, idx);
```

A number of plasticity models can be evaluated using the inputs in the computeFlowStress call. The variable d_plasticity is polymorphic and can represent any of the plasticity models that can be created by the plasticity model factory. The plastic evolution variables are updated using a polymorphic function along the lines of computeFlowStress.

The equation of state is used to calculate the hydrostatic stress using a function call of the form

```
Matrix3 tensorHy = d_eos->computePressure(matl, bulk, shear,
                                            tensorF_new, tensorD,
                                            tensorP, pTemperature[idx],
                                            rho_cur, delT);
```

Similarly, the damage model is called using a function of the type

```
double damage = d_damage->computeScalarDamage(tensorEta, tensorS,
                                                pTemperature[idx],
                                                delT, matl, d_tol,
                                                pDamage[idx]);
```

Therefore, the plasticity, damage and equation of state models are easily be inserted into any other type of stress update algorithm without any change being needed in them.

7.3 Models in the "PlasticityModels" directory

PlasticityModels contains several submodels that are primarily applicable to high strain-rate metal plasticity. The models in this folder are derived from the following base classes:

```
PlasticityModels/
|-- DamageModel.h
|-- DevStressModel.h
|-- FlowModel.h
|-- KinematicHardeningModel.h
|-- MeltingTempModel.h
|-- MPMEquationOfState.h
|-- ShearModulusModel.h
|-- SpecificHeatModel.h
|-- StabilityCheck.h
|-- ViscoPlasticityModel.h
|-- YieldCondition.h
```

These models are created inside a ConstitutiveModel using the factory idiom. The following factories are available:

```
PlasticityModels/
|-- DamageModelFactory.cc
|-- DevStressModelFactory.cc
|-- FlowStressModelFactory.cc
|-- KinematicHardeningModelFactory.cc
```

```
|-- MeltingTempModelFactory.cc
|-- MPMEquationOfStateFactory.cc
|-- ShearModulusModelFactory.cc
|-- SpecificHeatModelFactory.cc
|-- StabilityCheckFactory.cc
|-- ViscoPlasticityModelFactory.cc
|-- YieldConditionFactory.cc
```

The material state is communicated to the submodels through the two structs:

```
PlasticityModels/
|-- DeformationState.h
|-- PlasticityState.h
```

7.3.1   Implemented "PlasticityModels" models

Each model factory can produce several types of submodel objects. The VAANGO implementation contains the following specialized submodels.

Damage models

```
PlasticityModels/
        HancockMacKenzieDamage.cc
        JohnsonCookDamage.cc
        NullDamage.cc
```

Deviatoric stress models

```
PlasticityModels/
        HypoElasticDevStress.cc
        HypoViscoElasticDevStress.cc
```

Flow stress models

```
PlasticityModels/
        IsoHardeningFlow.cc
        JohnsonCookFlow.cc
        MTSFlow.cc
        PTWFlow.cc
        SCGFlow.cc
        ZAFlow.cc
        ZAPolymerFlow.cc
```

Kinematic hardening models

```
PlasticityModels/
        ArmstrongFrederickKinematicHardening.cc
        NoKinematicHardening.cc
        PragerKinematicHardening.cc
```

Melting temperature models

```
PlasticityModels/
        BPSMeltTemp.cc
        ConstantMeltTemp.cc
        LinearMeltTemp.cc
        SCGMeltTemp.cc
```

Equation of state models

```
PlasticityModels/
        DefaultHypoElasticEOS.cc
        HyperElasticEOS.cc
        MieGruneisenEOS.cc
        MieGruneisenEOSEnergy.cc
```

Shear modulus models

```
PlasticityModels/
        ConstantShear.cc
        MTSShear.cc
        NPShear.cc
        PTWShear.cc
        SCGShear.cc
```

Specific heat models

```
PlasticityModels/
        ConstantCp.cc
        CopperCp.cc
        CubicCp.cc
        SteelCp.cc
```

Stability check models

```
PlasticityModels/
        BeckerCheck.cc
        DruckerBeckerCheck.cc
        DruckerCheck.cc
        NoneCheck.cc
```

Viscoplasticity models

```
PlasticityModels/
        SuvicI.cc
```

Yield condition models

```
PlasticityModels/
        GursonYield.cc
        VonMisesYield.cc
```

7.3.2 Using the models in "PlasticityModels"

Suppose that you want to design a new constitutive model MyModel but want to reuse some of the flow stress models in the PlasticityModels directory. The following steps are needed to integrate the submodels into your new model.

1. In your header file, MyModel.h , create a private pointer to the model:

```
#ifndef __MPM_CM_MyModel_H__
#define __MPM_CM_MyModel_H__
#include <CCA/Components/MPM/ConstitutiveModel/ConstitutiveModel.h>
#include <CCA/Components/MPM/ConstitutiveModel/PlasticityModels/FlowModel.h>
namespace Vaango {
  class MyModel : public ConstitutiveModel {
    public:
        //......
    private:
        Uintah::FlowModel* d_flow;
  };
}
#endif
```

2. In your implementation file, MyModel.cc , create a copy of the model in the constructor and delete the copy in the destructor:

```
#include <CCA/Components/MPM/ConstitutiveModel/MyModel.h>
#include <CCA/Components/MPM/ConstitutiveModel/PlasticityModels/
    FlowStressModelFactory.h>
// Constructor
MyModel::MyModel(Uintah::ProblemSpecP& ps, Uintah::MPMFlags* flags) :
    ConstitutiveModel(flags)
{
  d_flow = FlowStressModelFactory::create(ps);
  if (!d_flow) {
    std::ostringstream err;
    err << "An error occured in the FlowModelFactory that has \n"
        << " slipped through the existing bullet proofing. \n";
```

```
      throw Uintah::ProblemSetupExcepton(err.str(), __FILE__, __LINE__);
   }
}
// Copy constructor
MyModel::MyModel(const MyModel* model) : ConstitutiveModel(model)
{
   d_flow = FlowStressModelFactory::createCopy(model->d_flow);
}
// Destructor
MyModel::~MyModel()
{
   delete d_flow;
}
```

3. To make sure that the details of the flow stress model are added to the output for restarting the simulation from a checkpoint, you will have to add the following to MyModel.cc :

```
void
MyModel::outputProblemSpec(Uintah::ProblemSpecP& ps, bool output_cm_tag)
{
   Uintah::ProblemSpecP model_ps = ps;
   if (output_cm_tag) {
      model_ps = ps->appendChild("constitutive_model");
      model_ps->setAttribute("type", "my_model_tag");
   }
   d_flow->outputProblemSpec(model_ps);
}
```

4. Some flow stress models have their own associated internal variables. You will have to make sure that these are initialized, even if you don't plan to use a model with submodel internal variables.

```
// Set up particle state
void
MyModel::addParticleState(std::vector<const Uintah::VarLabel*>& from, std::vector<
      const Uintah::VarLabel*>& to)
{
   d_flow->addParticleState(from, to);
}
// Set up initialization task
void
MyModel::addInitialComputesAndRequires(Uintah::Task* task, const Uintah::
      MPMMaterial* matl, const Uintah::PatchSet* patch) const
{
   const Uintah::MaterialSubset* matlset = matl->thisMaterial();
   d_flow->addInitialComputesAndRequires(task, matl, patch);
}
// Do the actual initialization
void
MyModel::initializeCMData(const Uintah::Patch* patch, const Uintah::MPMMaterial*
      matl, Uintah::DataWarehouse* new_dw)
{
   Uintah::ParticleSubset* pset = new_dw->getParticleSubset(matl->getDWIndex(),
         patch);
   d_flow->initializeInternalVars(pset, new_dw);
}
```

5. Now you are almost ready to use the flow stress model in the stress computation logic. To complete the process, you will have to add a task that makes sure any submodel internal variables are updated correctly during the stress computation process:

```
void
MyModel::addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial* matl
      , const Uintah::PatchSet* patches) const
{
   d_flow->addComputesAndRequires(task, matl, patches);
}
```

6. Finally, you can use the flow stress model object for stress computation:

```
void
```

```
MyModel::computeStressTensor(const Uintah::PatchSubset* patches, const Uintah::
    MPMMaterial* matl, Uintah::DataWarehouse* old_dw, Uintah::DataWarehouse* new_dw
    )
{
  // Set up initial state
  Uintah::PlasticityState state;
  state.initialTemperature = ...;
  state.initialDensity = ...;
  state.initialVolume = ...;
  state.initialBulkModulus = ...;
  state.initialShearModulus = ...;
  state.initialMeltTemp = ...;
  state.energy = ...;
  // Loop through patches
  for (int patchIndex = 0; patchIndex < patches->size(); patchIndex++) {
    const Patch* patch = patches->get(patchIndex);
    int dwi = matl->getDWIndex();
    // Get the particle set and do gets and allocations
    ParticleSubset* pset = old_dw->getParticleSubset(dwi, patch);
    d_flow->getInternalVars(pset, old_dw);
    d_flow->allocateAndPutInternalVars(pset, new_dw);
    // Loop through particles in a patch
    for (auto idx : *pset) {
      // Update the state
      state.strainRate = ...;
      state.plasticStrainRate = ...;
      state.plasticStrain = ...;
      state.pressure = ...;
      state.temperature = ...;
      state.density = ...;
      state.volume = ...;
      state.bulkModulus = ...;
      state.shearModulus = ...;
      state->meltingTemp = ...;
      state->specificHeat = ...;
      // Calculate the flow stress using the flow stress model
      state->yieldStress = d_flow->computeFlowStress(state, delT, d_tol, matl, idx)
          ;
      // .......
    } // End particle loop
  } // End patch loop
}
```

7.3.3 Creating a new model in "PlasticityModels"

Let us now suppose that you want to add a new flow stress model called MyFlow to Vaango. You will have to use the following process to create the model so that it can be used by the existing constitutive models as well as the new model that you created in the previous section.

1. First create a header file MyFlow.h in the PlasticityModels directory:

```
#ifndef __MPM_CM_PM_MyFlow_H__
#define __MPM_CM_PM_MyFlow_H__
#include <CCA/Components/MPM/ConstitutiveModels/PlasticityModels/FlowModel.h>
....
namespace Vaango {
  class MyFlow : public FlowModel
  {
  public:
    // Model parameters
    struct MyFlowParameters {
      double parameter1;
      double parameter2;
      ....
    };

    // Internal variables
    constParticleVariable<double> pInternalVar1, pInternalVar2;
    ParticleVariable<double> pInternalVar1_new, pInternalVar2_new;
```

```cpp
// Internal variable labels
const Uintah::VarLabel* pInternalVar1Label, pInternalVar2Label;
const Uintah::VarLabel* pInternalVar1Label_preReloc,
    pInternalVar2Label_preReloc;

// Constructors and destructor
MyFlow(Uintah::ProblemSpecP& ps);
MyFlow(const MyFlow* flow);
~MyFlow() override;

// Delete assignment operator
MyFlow& operator=(const MyFlow& flow) = delete;

// Compute the flow stress
double computeFlowStress(const Uintah::PlasticityState* state, const double&
    delT, const double& tolerance, const Uintah::MPMMaterial* matl, const
    Uintah::particleIndex idx) override;

// Calculate the plastic strain rate [epdot(tau,ep,T)]
double computeEpdot(const Uintah::PlasticityState* state, const double& delT,
    const double& tolerance, const Uintah::MPMMaterial* matl, const Uintah::
    particleIndex idx) override;

// Evaluate derivative of flow stress with respect to scalar and internal
    variables.
void evalDerivativeWRTScalarVars(const Uintah::PlasticityState* state, const
    Uintah::particleIndex idx, Uintah::Vector& derivs) override;

// Evaluate derivative of flow stress with respect to plastic strain
double evalDerivativeWRTPlasticStrain(const Uintah::PlasticityState* state,
    const Uintah::particleIndex idx) override;

// Evaluate derivative of flow stress with respect to strain rate.
double evalDerivativeWRTStrainRate(const Uintah::PlasticityState* state, const
    Uintah::particleIndex idx) override;

//Compute the elastic-plastic tangent modulus
void computeTangentModulus(const Uintah::Matrix3& stress, const Uintah::
    PlasticityState* state, const double& delT, const Uintah::MPMMaterial* matl
    , const Uintah::particleIndex idx, Uintah::TangentModulusTensor& Ce, Uintah
    ::TangentModulusTensor& Cep) override;

// Flow stress update methods
void updateElastic(const Uintah::particleIndex idx) override;
void updatePlastic(const Uintah::particleIndex idx, const double& delGamma)
    override;

// Shear modulus and melting temperature computations
double computeShearModulus(const Uintah::PlasticityState* state) override;
double computeMeltingTemp(const Uintah::PlasticityState* state) override;

// Boilerplate methods used by the computational framework
void addInitialComputesAndRequires(Uintah::Task* task, const Uintah::
    MPMMaterial* matl, const Uintah::PatchSet* patches) override;
void addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial* matl
    , const Uintah::PatchSet* patches) override;
void addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial* matl
    , const Uintah::PatchSet* patches, bool recurse, bool schedParent) override
    ;
void allocateCMDataAddRequires(Uintah::Task* task, const Uintah::MPMMaterial*
    matl, const Uintah::PatchSet* patch, Uintah::MPMLabel* lb) override;
void allocateCMDataAdd(Uintah::DataWarehouse* new_dw, Uintah::ParticleSubset*
    addset, Uintah::ParticleLabelVariableMap* newState, Uintah::ParticleSubset*
     delset, Uintah::DataWarehouse* old_dw) override;
void allocateAndPutRigid(Uintah::ParticleSubset* pset, Uintah::DataWarehouse*
    new_dw) override;

// Methods needed if the model has its own internal variables
void addParticleState(std::vector<const VarLabel*>& from, std::vector<const
    VarLabel*>& to) override;
void initializeInternalVars(Uintah::ParticleSubset* pset, Uintah::DataWarehouse
    * new_dw) override;
void getInternalVars(Uintah::ParticleSubset* pset, Uintah::DataWarehouse*
```

```
        old_dw) override;
    void allocateAndPutInternalVars(Uintah::ParticleSubset* pset, Uintah::
        DataWarehouse* new_dw) override;

  private:
    MyFlowParameters d_flow;
    void createInternalVarLabels();
  }
}
#endif
```

2. In the implementation file MyFlow.cc , create the constructors and make sure that the model input parameters are copied to output files:

```
#include <CCA/Components/MPM/ConstitutiveModels/PlasticityModels/FlowModel.h>
MyFlow::MyFlow(Uintah::ProblemSpecP& ps)
{
  ps->require("my_flow_parameter1", d_flow.parameter1);
  ps->require("my_flow_parameter2", d_flow.parameter2);
  createInternalVarLabels();
}

MyFlow::MyFlow(const MyFlow* flow)
{
  d_flow.parameter1 = flow->d_flow.parameter1;
  d_flow.parameter2 = flow->d_flow.parameter2;
  createInternalVarLabels();
}

void
MyFlow::createInternalVarLabels()
{
  pInternalVar1Label = Uintah::VarLabel::create("p.my_flow_var1", Uintah::
      ParticleVariable<double>::getTypeDescription());
  pInternalVar2Label = Uintah::VarLabel::create("p.my_flow_var2", Uintah::
      ParticleVariable<double>::getTypeDescription());
  pInternalVar1Label_preReloc = Uintah::VarLabel::create("p.my_flow_var1+", Uintah
      ::ParticleVariable<double>::getTypeDescription());
  pInternalVar2Label_preReloc = Uintah::VarLabel::create("p.my_flow_var2+", Uintah
      ::ParticleVariable<double>::getTypeDescription());
}

void
MyFlow::outputProblemSpec(Uintah:;ProblemSpecP& ps)
{
  ProblemSpecP flow_ps = ps->appendChild("flow_model");
  flow_ps->setAttribute("type", "my_flow_model");
  flow_ps->appendElement("my_flow_parameter1", d_flow.parameter1);
  flow_ps->appendElement("my_flow_parameter2", d_flow.parameter2);
}
```

3. We can now implement the flow stress computation code:

```
double
MyFlow::computeFlowStress(const Uintah::PlasticityState* state, const double&,
    const double&, const Uintah::MPMMaterial*, const Uintah::particleIndex idx)
{
  // Get the internal variables
  double internalVar1 = pInternalVar1[idx];
  double internalVar2 = pInternalVar2[idx];

  // Compute the flow stress
  double flowStress = doSomeComputation(internalVar1, internalVar2);

  // Update the internal variables if needed
  pInternalVar1_new[idx] = internalVar1;
  pInternalVar2_new[idx] = internalVar2;

  return flowStress;
}
```

4. Compute the plastic strain rate if possible:

```
double
MyFlow::computeEpdot(const Uintah::PlasticityState* state, const double& delT,
    const double&, const Uintah::MPMMaterial*, const Uintah::particleIndex pidx)
{
  double epdot = doSomComputation();
  return epdot;
}
```

5. Compute the derivative of the flow stress function with respect to the scalar plastic strain and strain
   rate:

```
double
MyFlow::evalDerivativeWRTPlasticStrain(const Uintah::PlasticityState* state, const
    Uintah::particleIndex idx)
{
  // Do the calculation
  return derivative;
}
double
MyFlow::evalDerivativeWRTStrainRate(const Uintah::PlasticityState* state, const
    Uintah::particleIndex idx)
{
  // Do the calculation
  return derivative;
}
```

6. Compute the derivative of the flow stress function with respect to some scalar quantities if needed:

```
void
MyFlow::evalDerivativeWRTScalarVars(const Uintah::PlasticityState* state, const
    Uintah::particleIndex idx, Uintah::Vector& derivs)
{
  derivs[0] = evalDerivativeWRTStrainRate(state, idx);
  derivs[1] = evalDerivativeWRTTemperature(state, idx);
  derivs[2] = evalDerivativeWRTPlasticStrain(state, idx);
}
```

7. Compute the tangent modulus for implicit calculations. This is typically quite involved and usually
   avoided by the plasticity models in VAANGO .
8. Compute the shear modulus and melting temperature if needed. We usually use the separate shear
   modulus and melting temperature models to accomplish this task.
9. At this stage we should set up the boilerplate code for communicating the internal variables:

```
void
MyFlow::addInitialComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial
    * matl, const PatchSet*)
{
  const Uintah::MaterialSubset* matlset = matl->thisMaterial();
  task->computes(pInternalVar1Label, matlset);
  task->computes(pInternalVar2Label, matlset);
}

void
MyFlow::addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial* matl,
    const Uintah::PatchSet*)
{
  const Uintah::MaterialSubset* matlset = matl->thisMaterial();
  task->requires(Task::OldDW, pInternalVar1Label, matlset, Ghost::None);
  task->requires(Task::OldDW, pInternalVar2Label, matlset, Ghost::None);
  task->computes(pInternalVar1Label_preReloc, matlset);
  task->computes(pInternalVar2Label_preReloc, matlset);
}

void
MyFlow::addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial*,
    const Uintah::PatchSet*, bool, bool)
{
  const Uintah::MaterialSubset* matlset = matl->thisMaterial();
  task->requires(Task::ParentOldDW, pInternalVar1Label, matlset, Ghost::None);
  task->requires(Task::ParentOldDW, pInternalVar2Label, matlset, Ghost::None);
```

```cpp
}

void
MyFlow::addParticleState(std::vector<const Uintah::VarLabel*>&, std::vector<const
    Uintah::VarLabel*>&)
{
  from.push_back(pInternalVar1Label);
  from.push_back(pInternalVar2Label);
  to.push_back(pInternalVar1Label_preReloc);
  to.push_back(pInternalVar2Label_preReloc);
}

void
MyFlow::initializeInternalVars(Uintah::ParticleSubset* pset, Uintah::DataWarehouse*
    new_dw)
{
  new_dw->allocateAndPut(pInternalVar1_new, pInternalVar1Label, pset);
  new_dw->allocateAndPut(pInternalVar2_new, pInternalVar2Label, pset);
  for (auto pidx : *pset) {
    pInternalVar1_new[pidx] = 0.0;
    pInternalVar2_new[pidx] = 0.0;
  }
}

void
MyFlow::getInternalVars(Uintah::ParticleSubset* pset, Uintah::DataWarehouse* old_dw
    )
{
  old_dw->get(pInternalVar1, pInternalVar1Label, pset);
  old_dw->get(pInternalVar2, pInternalVar2Label, pset);
}

void
MyFlow::allocateAndPutInternalVars(Uintah::ParticleSubset* pset, Uintah::
    DataWarehouse* new_dw)
{
  new_dw->allocateAndPut(pInternalVar1_new, pInternalVar1Label_preReloc, pset);
  new_dw->allocateAndPut(pInternalVar2_new, pInternalVar2Label_preReloc, pset);
}

void
MyFlow::allocateAndPutRigid(Uintah::ParticleSubset* pset, Uintah::DataWarehouse*
    new_dw)
{
  allocateAndPutInternalVars(pset, new_dw);
  for (auto pidx : *pset) {
    pInternalVar1_new[pidx] = 0.0;
    pInternalVar2_new[pidx] = 0.0;
  }
}

void
MyFlow::allocateCMDataAddRequires(Uintah::Task* task, const Uintah::MPMMaterial*
    matl, const Uintah::PatchSet*, Uintah::MPMLabel*)
{
  const Uintah::MaterialSubset* matlset = matl->thisMaterial();
  task->requires(Task::NewDW, pInternalVar1Label_preReloc, matlset, Ghost::None);
  task->requires(Task::NewDW, pInternalVar2Label_preReloc, matlset, Ghost::None);
}

void
MyFlow::allocateCMDataAdd(Uintah::DataWarehouse* new_dw, Uintah::ParticleSubset*
    addset, Uintah::ParticleLabelVariableMap* newState, Uintah::ParticleSubset*
    delset, Uintah::DataWarehouse* old_dw)
{
  Uintah::ParticleVariable<double> n_internalVar1, n_internalVar2;
  Uintah::constParticleVariable<double> o_internalVar1, o_internalVar2;

  new_dw->allocateTemporary(n_internalVar1, addset);
  new_dw->allocateTemporary(n_internalVar2, addset);

  new_dw->get(o_internalVar1 , pInternalVar1Label_preReloc, delset);
  new_dw->get(o_internalVar2 , pInternalVar2Label_preReloc, delset);
```

```
    ParticleSubset::iterator o,n = addset->begin();
    for (o = delset->begin(); o != delset->end(); o++, n++) {
      n_internalVar1[*n] = o_internalVar1[*o];
      n_internalVar2[*n] = o_internalVar2[*o];
    }

    (*newState)[pInternalVar1Label] = n_internalVar1.clone();
    (*newState)[pInternalVar2Label] = n_internalVar2.clone();
}

void
MyFlow::updateElastic(const Uintah::particleIndex pidx)
{
    pInternalVar1_new[idx] = pInternalVar1[idx];
    pInternalVar2_new[idx] = pInternalVar2[idx];
}

void
MyFlow::updatePlastic(const particleIndex pidx, const double& someValue)
{
    pInternalVar1_new[idx] = pInternalVar1_new[idx];
    pInternalVar2_new[idx] = pInternalVar2_new[idx] + someValue;
}
```

10. We now add this model to the FlowStressModelFactory as follows:

```
#include <CCA/Components/MPM/ConstitutiveModel/PlasticityModels/MyFlow.h>
using namespace Uintah;
FlowModel*
FlowStressModelFactory::create(ProblemSpecP& ps)
{
    ProblemSpecP child = ps->findBlock("flow_model");
    if (!child)
        throw ProblemSetupException("Cannot find flow_model tag", __FILE__,
                                    __LINE__);
    string mat_type;
    if (!child->getAttribute("type", mat_type))
        throw ProblemSetupException("No type for flow_model", __FILE__, __LINE__);
    .....
    else if (mat_type == "my_flow_model") {
        return (scinew Vaango::MyFlow(child));
    }
    .....
}

FlowModel*
FlowStressModelFactory::createCopy(const FlowModel* pm)
{
    .....
    else if (dynamic_cast<const Vaango::MyFlow*>(pm))
        return (scinew Vaango::MyFlow(dynamic_cast<const Vaango::MyFlow*>(pm)));
    .....
}
```

11. Next we add the new file to the compilation list in CMakeLists.txt in the PlasticityModels directory:

```
SET(MPM_ConstitutiveModel_PlasticityModels_SRCS
    ................
    MyFlow.cc
)
```

12. Finally, we add the new model tags in the constitutive_models.xml file in the directory src/StandAlone/inputs/UPS_SPEC:

```
    <flow_model spec="OPTIONAL NO_DATA" attribute1="type REQUIRED STRING '
        isotropic_hardening, johnson_cook, mts_model, .... ,
        zerilli_armstrong_polymer, my_flow_model'" >
      ........
      <my_flow_parameter1 spec="REQUIRED DOUBLE" need_applies_to "type my_flow_model
          "/>
```

```
        <my_flow_parameter2 spec="REQUIRED DOUBLE" need_applies_to "type my_flow_model
            "/>
    </flow_model>
```

7.4 Models in the "Models" directory

Models contains submodels that include some of the older models from the PlasticityModels directory but mostly models that are applicable to soil and rock plasticity. The approach taken is similar to that in PlasticityModels in that the models are derived from base classes and generated using factories.

The models in the Models folder are derived from the following base classes:

```
Models/
|-- ElasticModuliModel.h
|-- InternalVariableModel.h
|-- KinematicHardeningModel.h
|-- PressureModel.h
|-- ShearModulusModel.h
|-- YieldCondition.h
```

The factory approach used for these models is similar to that used for the PlasticityModels. The following factories are available:

```
Models/
|-- ElasticModuliModelFactory.cc
|-- InternalVariableModelFactory.cc
|-- KinematicHardeningModelFactory.cc
|-- PressureModelFactory.cc
|-- ShearModulusModelFactory.cc
|-- YieldConditionFactory.cc
```

Two reduce inappropriate mixing of models, the material state is communicated to the submodels through a set of structures that are derived from a ModelStateBase object. The available state structures are:

```
Models/
|-- ModelState_Arena.h
|-- ModelState_Arenisca3.h
|-- ModelState_CamClay.h
|-- ModelState_Default.h
|-- ModelState_SoilModelBrannon.h
|-- ModelState_Tabular.h
```

Tabular models are created using the TabularData class. Details are given in Section 7.4.2.

In addition there are a few utility functions that are implemented in

```
Models/
|-- TableContainers.h
|-- TableUtils.cc
|-- YieldCondUtils.cc
```

7.4.1 Implemented "Model" models

Each model factory can produce several types of submodel objects. The Vaango implementation contains the following specialized submodels.

Elastic moduli models

```
        ElasticModuli_Arena.cc
        ElasticModuli_ArenaMixture.cc
        ElasticModuli_Arenisca.cc
        ElasticModuli_Constant.cc
        ElasticModuli_Tabular.cc
```

Internal variable models

```
Models/
        InternalVar_Arena.cc
        InternalVar_BorjaPressure.cc
        InternalVar_SoilModelBrannonKappa.cc
```

Kinematic hardening models

```
Models/
        KinematicHardening_Arena.cc
        KinematicHardening_Armstrong.cc
        KinematicHardening_None.cc
        KinematicHardening_Prager.cc
```

Pressure models

```
Models/
        Pressure_Air.cc
        Pressure_Borja.cc
        Pressure_Granite.cc
        Pressure_Hyperelastic.cc
        Pressure_Hypoelastic.cc
        Pressure_MieGruneisen.cc
        Pressure_Water.cc
```

Shear modulus models

```
Models/
        ShearModulus_Borja.cc
        ShearModulus_Constant.cc
        ShearModulus_Nadal.cc
```

Yield condition models

```
Models/
        YieldCond_Arena.cc
        YieldCond_ArenaMixture.cc
        YieldCond_Arenisca3.cc
        YieldCond_CamClay.cc
        YieldCond_Gurson.cc
        YieldCond_Tabular.cc
        YieldCond_vonMises.cc
```

7.4.2   Using the submodels in "Models"

The procedure for using the submodels in the Models directory in a new constitutive model is almost, but not exactly, identical to that for the PlasticityModels . As an example, let us see how the TabuarPlasticity model has been implemented.

The following steps are were used to integrate the submodels needed for TabularPlasticity into the model.

1. The header file, TabularPlasticity.h , contains:

    ```cpp
    #ifndef __MPM_CONSTITUTIVEMODEL_TABULAR_PLASTICITY_H__
    #define __MPM_CONSTITUTIVEMODEL_TABULAR_PLASTICITY_H__

    #include <CCA/Components/MPM/ConstitutiveModel/ConstitutiveModel.h>
    #include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuliModel.h>
    #include <CCA/Components/MPM/ConstitutiveModel/Models/ModelState_Tabular.h>
    #include <CCA/Components/MPM/ConstitutiveModel/Models/YieldCondition.h>

    namespace Vaango {
      class TabularPlasticity : public ConstitutiveModel {
        public:
          // Model parameters
          struct CMData
          {
            double yield_scale_fac;  // A scaling factor for number of substeps
            double subcycling_characteristic_number; // A max substeps value
          };

          // Local variables that can be communicated across processes
          const Uintah::VarLabel* pElasticStrainLabel; // Elastic Strain
    ```

```cpp
const Uintah::VarLabel* pElasticStrainLabel_preReloc;
const Uintah::VarLabel* pPlasticStrainLabel; // Plastic Strain
const Uintah::VarLabel* pPlasticStrainLabel_preReloc;
const Uintah::VarLabel* pPlasticCumEqStrainLabel; // Equivalent plastic
    strain
const Uintah::VarLabel* pPlasticCumEqStrainLabel_preReloc;
const Uintah::VarLabel* pElasticVolStrainLabel; // Elastic Volumetric Strain
const Uintah::VarLabel* pElasticVolStrainLabel_preReloc;
const Uintah::VarLabel* pPlasticVolStrainLabel; // Plastic Volumetric Strain
const Uintah::VarLabel* pPlasticVolStrainLabel_preReloc;
const Uintah::VarLabel* pRemoveLabel; // Flag for removal
const Uintah::VarLabel* pRemoveLabel_preReloc;

// Constructors/assignment/destructors
TabularPlasticity(Uintah::ProblemSpecP& ps, Uintah::MPMFlags* flag);
TabularPlasticity(const TabularPlasticity* cm);
TabularPlasticity(const TabularPlasticity& cm);
~TabularPlasticity() override;
TabularPlasticity& operator=(const TabularPlasticity& cm) = delete;
TabularPlasticity* clone() override;

// For restart purposes
void outputProblemSpec(Uintah::ProblemSpecP& ps, bool output_cm_tag = true)
    override;

// Model parameter dictionary get method
ParameterDict getParameters() const
{
  ParameterDict params;
  return params;
}

// Initialization methods
void addInitialComputesAndRequires( Uintah::Task* task, const Uintah::
    MPMMaterial* matl, const Uintah::PatchSet* patches) const override;
void initializeCMData(const Uintah::Patch* patch, const Uintah::MPMMaterial*
    matl, Uintah::DataWarehouse* new_dw) override;

// Main stress computation methods
void addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial*
    matl, const Uintah::PatchSet* patches) const override;
void computeStableTimestep(const Uintah::Patch* patch, const Uintah::
    MPMMaterial* matl, Uintah::DataWarehouse* new_dw) override;
void computeStressTensor(const Uintah::PatchSubset* patches, const Uintah::
    MPMMaterial* matl, Uintah::DataWarehouse* old_dw, Uintah::DataWarehouse*
    new_dw) override;

// Boilerplate methods for data communication
void addRequiresDamageParameter( Uintah::Task* task, const Uintah::
    MPMMaterial* matl, const Uintah::PatchSet* patches) const override;
void getDamageParameter(const Uintah::Patch* patch, Uintah::ParticleVariable<
    int>& damage, int dwi, Uintah::DataWarehouse* old_dw, Uintah::
    DataWarehouse* new_dw) override;
void carryForward(const Uintah::PatchSubset* patches, const Uintah::
    MPMMaterial* matl, Uintah::DataWarehouse* old_dw, Uintah::DataWarehouse*
    new_dw) override;
void addComputesAndRequires(Uintah::Task* task, const Uintah::MPMMaterial*
    matl, const Uintah::PatchSet* patches, const bool recursion, const bool
    dummy) const override;
void addParticleState(std::vector<const Uintah::VarLabel*>& from, std::vector
    <const Uintah::VarLabel*>& to) override;
void allocateCMDataAdd(Uintah::DataWarehouse* new_dw, Uintah::ParticleSubset*
     addset, Uintah::ParticleLabelVariableMap* newState, Uintah::
    ParticleSubset* delset, Uintah::DataWarehouse* old_dw) override;

// For communication with MPMICE
double computeRhoMicroCM(double pressure, const double p_ref, const Uintah::
    MPMMaterial* matl, double temperature, double rho_guess) override;
void computePressEOSCM(double rho_m, double& press_eos, double p_ref, double&
     dp_drho, double& ss_new, const Uintah::MPMMaterial* matl, double
    temperature) override;
double getCompressibility() override;
```

```
    private:
      ElasticModuliModel* d_elastic;
      YieldCondition* d_yield;
      CMData d_cm;

      void initializeLocalMPMLabels();
      // Other private methods
  };
}
#endif
```

2. In the implementation file, TabularPlasticity.cc , create copies of the submodel in the constructor, initialize MPMLabel s, and delete the copies in the destructor:

```cpp
#include <CCA/Components/MPM/ConstitutiveModel/TabularPlasticity.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuliModelFactory.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/YieldConditionFactory.h>
....

// Constructor
TabularPlasticity::TabularPlasticity(Uintah::ProblemSpecP& ps, Uintah::MPMFlags*
    mpmFlags) : Uintah::ConstitutiveModel(mpmFlags)
{
  // Bulk and shear modulus models
  d_elastic = Vaango::ElasticModuliModelFactory::create(ps);
  if (!d_elastic) {
    std::ostringstream desc;
    desc << "**ERROR** Internal error while creating ElasticModuliModel."
         << std::endl;
    throw InternalError(desc.str(), __FILE__, __LINE__);
  }

  // Yield condition model
  d_yield = Vaango::YieldConditionFactory::create(ps);
  if (!d_yield) {
    std::ostringstream desc;
    desc << "**ERROR** Internal error while creating YieldConditionModel."
         << std::endl;
    throw InternalError(desc.str(), __FILE__, __LINE__);
  }

  // Algorithmic parameters
  ps->getWithDefault("yield_surface_radius_scaling_factor",
                     d_cm.yield_scale_fac, 1.0);
  ps->getWithDefault("subcycling_characteristic_number",
                     d_cm.subcycling_characteristic_number,
                     256); // allowable subcycles

  initializeLocalMPMLabels();
}

// Copy constructors (don't allow defaults)
TabularPlasticity::TabularPlasticity(const TabularPlasticity& cm)
  : ConstitutiveModel(cm)
{
  d_elastic = Vaango::ElasticModuliModelFactory::createCopy(cm.d_elastic);
  d_yield = Vaango::YieldConditionFactory::createCopy(cm.d_yield);

  // Yield surface scaling
  d_cm.yield_scale_fac = cm.d_cm.yield_scale_fac;

  // Subcycling
  d_cm.subcycling_characteristic_number =
    cm.d_cm.subcycling_characteristic_number;

  initializeLocalMPMLabels();
}
// *NOTE* Delegation is a C++11 feature
TabularPlasticity::TabularPlasticity(const TabularPlasticity* cm)
  : TabularPlasticity(*cm)
{
}
TabularPlasticity*
```

```cpp
TabularPlasticity::clone()
{
  return scinew TabularPlasticity(*this);
}

// Initialize all labels of the particle variables associated with
// TabularPlasticity.
void
TabularPlasticity::initializeLocalMPMLabels()
{
  pElasticStrainLabel = Uintah::VarLabel::create(
    "p.elasticStrain",
    Uintah::ParticleVariable<Uintah::Matrix3>::getTypeDescription());
  pElasticStrainLabel_preReloc = Uintah::VarLabel::create(
    "p.elasticStrain+",
    Uintah::ParticleVariable<Uintah::Matrix3>::getTypeDescription());

  pElasticVolStrainLabel = VarLabel::create(
    "p.elasticVolStrain", ParticleVariable<double>::getTypeDescription());
  pElasticVolStrainLabel_preReloc = VarLabel::create(
    "p.elasticVolStrain+", ParticleVariable<double>::getTypeDescription());

  pPlasticStrainLabel = Uintah::VarLabel::create(
    "p.plasticStrain",
    Uintah::ParticleVariable<Uintah::Matrix3>::getTypeDescription());
  pPlasticStrainLabel_preReloc = Uintah::VarLabel::create(
    "p.plasticStrain+",
    Uintah::ParticleVariable<Uintah::Matrix3>::getTypeDescription());

  pPlasticCumEqStrainLabel = Uintah::VarLabel::create(
    "p.plasticCumEqStrain",
    Uintah::ParticleVariable<double>::getTypeDescription());
  pPlasticCumEqStrainLabel_preReloc = Uintah::VarLabel::create(
    "p.plasticCumEqStrain+",
    Uintah::ParticleVariable<double>::getTypeDescription());

  pPlasticVolStrainLabel = Uintah::VarLabel::create(
    "p.plasticVolStrain",
    Uintah::ParticleVariable<double>::getTypeDescription());
  pPlasticVolStrainLabel_preReloc = Uintah::VarLabel::create(
    "p.plasticVolStrain+",
    Uintah::ParticleVariable<double>::getTypeDescription());

  pRemoveLabel = Uintah::VarLabel::create(
    "p.remove",
    Uintah::ParticleVariable<int>::getTypeDescription());
  pRemoveLabel_preReloc = Uintah::VarLabel::create(
    "p.remove+",
    Uintah::ParticleVariable<int>::getTypeDescription());
}

// Destructor
TabularPlasticity::~TabularPlasticity()
{
  VarLabel::destroy(pElasticStrainLabel);
  VarLabel::destroy(pElasticStrainLabel_preReloc);
  VarLabel::destroy(pElasticVolStrainLabel); // Elastic Volumetric Strain
  VarLabel::destroy(pElasticVolStrainLabel_preReloc);
  VarLabel::destroy(pPlasticStrainLabel);
  VarLabel::destroy(pPlasticStrainLabel_preReloc);
  VarLabel::destroy(pPlasticCumEqStrainLabel);
  VarLabel::destroy(pPlasticCumEqStrainLabel_preReloc);
  VarLabel::destroy(pPlasticVolStrainLabel);
  VarLabel::destroy(pPlasticVolStrainLabel_preReloc);
  VarLabel::destroy(pRemoveLabel);
  VarLabel::destroy(pRemoveLabel_preReloc);

  delete d_yield;
  delete d_elastic;
}
```

3. To make sure that the details of the submodels are added to the output for restarting the simulation

from a checkpoint, add the following to TabularPlasticity.cc :

```
void
TabularPlasticity::outputProblemSpec(ProblemSpecP& ps, bool output_cm_tag)
{
  ProblemSpecP cm_ps = ps;
  if (output_cm_tag) {
    cm_ps = ps->appendChild("constitutive_model");
    cm_ps->setAttribute("type", "tabular_plasticity");
  }

  d_elastic->outputProblemSpec(cm_ps);
  d_yield->outputProblemSpec(cm_ps);

  cm_ps->appendElement("yield_surface_radius_scaling_factor",
                       d_cm.yield_scale_fac);
  cm_ps->appendElement("subcycling_characteristic_number",
                       d_cm.subcycling_characteristic_number);
}
```

4. Some submodels have their own associated internal variables, in this case the YieldCondition . You will have to make sure that these are initialized, even if you don't plan to use a model with submodel internal variables.

```
// Add the labels to the label list
void
TabularPlasticity::addParticleState(std::vector<const VarLabel*>& from, std::vector
    <const VarLabel*>& to)
{
  from.push_back(pElasticStrainLabel);
  to.push_back(pElasticStrainLabel_preReloc);

  from.push_back(pElasticVolStrainLabel);
  to.push_back(pElasticVolStrainLabel_preReloc);

  from.push_back(pPlasticStrainLabel);
  to.push_back(pPlasticStrainLabel_preReloc);

  from.push_back(pPlasticCumEqStrainLabel);
  to.push_back(pPlasticCumEqStrainLabel_preReloc);

  from.push_back(pPlasticVolStrainLabel);
  to.push_back(pPlasticVolStrainLabel_preReloc);

  from.push_back(pRemoveLabel);
  to.push_back(pRemoveLabel_preReloc);

  // Add the particle state for the yield condition model
  d_yield->addParticleState(from, to);
}
// Set up initialization task
void
TabularPlasticity::addInitialComputesAndRequires(Task* task, const MPMMaterial*
    matl, const PatchSet* patch) const
{
  const MaterialSubset* matlset = matl->thisMaterial();

  task->computes(pElasticStrainLabel, matlset);
  task->computes(pElasticVolStrainLabel, matlset);
  task->computes(pPlasticStrainLabel, matlset);
  task->computes(pPlasticCumEqStrainLabel, matlset);
  task->computes(pPlasticVolStrainLabel, matlset);
  task->computes(pRemoveLabel, matlset);

  d_yield->addInitialComputesAndRequires(task, matl, patch);
}
// Do the actual initialization
void
TabularPlasticity::initializeCMData(const Patch* patch, const MPMMaterial* matl,
    DataWarehouse* new_dw)
{
  // Get the particles in the current patch
  ParticleSubset* pset = new_dw->getParticleSubset(matl->getDWIndex(), patch);
```

```cpp
  // Get the particle volume and mass
  constParticleVariable<double> pVolume, pMass;
  new_dw->get(pVolume, lb->pVolumeLabel, pset);
  new_dw->get(pMass, lb->pMassLabel, pset);

  // Initialize variables for yield function parameter variability
  d_yield->initializeLocalVariables(patch, pset, new_dw, pVolume);

  // Now initialize the other variables
  ParticleVariable<int> pRemove;
  ParticleVariable<double> pdTdt;
  ParticleVariable<double> pElasticVolStrain;
  ParticleVariable<double> pPlasticCumEqStrain, pPlasticVolStrain;
  ParticleVariable<Matrix3> pStress;
  ParticleVariable<Matrix3> pElasticStrain;
  ParticleVariable<Matrix3> pPlasticStrain;

  new_dw->allocateAndPut(pdTdt, lb->pdTdtLabel, pset);
  new_dw->allocateAndPut(pStress, lb->pStressLabel, pset);

  new_dw->allocateAndPut(pRemove, pRemoveLabel, pset);
  new_dw->allocateAndPut(pElasticStrain, pElasticStrainLabel, pset);
  new_dw->allocateAndPut(pElasticVolStrain, pElasticVolStrainLabel, pset);
  new_dw->allocateAndPut(pPlasticStrain, pPlasticStrainLabel, pset);
  new_dw->allocateAndPut(pPlasticCumEqStrain, pPlasticCumEqStrainLabel, pset);
  new_dw->allocateAndPut(pPlasticVolStrain, pPlasticVolStrainLabel, pset);

  for (const particleIndex& pidx : *pset) {
    pRemove[pidx] = 0;
    pdTdt[pidx] = 0.0;
    pStress[pidx] = Identity;
    pElasticStrain[pidx].set(0.0);
    pElasticVolStrain[pidx] = 0.0;
    pPlasticStrain[pidx].set(0.0);
    pPlasticCumEqStrain[pidx] = 0.0;
    pPlasticVolStrain[pidx] = 0.0;
  }

  // Compute timestep
  computeStableTimestep(patch, matl, new_dw);
}
// Compute a stable timestep
void
TabularPlasticity::computeStableTimestep(const Patch* patch, const MPMMaterial*
    matl, DataWarehouse* new_dw)
{
  int matID = matl->getDWIndex();

  // Compute initial elastic moduli
  ElasticModuli moduli = d_elastic->getInitialElasticModuli();
  double bulk = moduli.bulkModulus;
  double shear = moduli.shearModulus;

  // Initialize wave speed
  double c_dil = std::numeric_limits<double>::min();
  Vector dx = patch->dCell();
  Vector WaveSpeed(c_dil, c_dil, c_dil);

  // Get the particles in the current patch
  ParticleSubset* pset = new_dw->getParticleSubset(matID, patch);

  // Get particles mass, volume, and velocity
  constParticleVariable<double> pMass, pVolume;
  constParticleVariable<long64> pParticleID;
  constParticleVariable<Vector> pVelocity;

  new_dw->get(pMass, lb->pMassLabel, pset);
  new_dw->get(pVolume, lb->pVolumeLabel, pset);
  new_dw->get(pParticleID, lb->pParticleIDLabel, pset);
  new_dw->get(pVelocity, lb->pVelocityLabel, pset);

  // loop over the particles in the patch
```

```
for (const particleIndex& idx : *pset) {

  // Compute wave speed + particle velocity at each particle,
  // store the maximum
  c_dil = std::sqrt((bulk + four_third * shear) * (pVolume[idx] / pMass[idx]));

  WaveSpeed = Vector(Max(c_dil + std::abs(pVelocity[idx].x()), WaveSpeed.x()),
      Max(c_dil + std::abs(pVelocity[idx].y()), WaveSpeed.y()), Max(c_dil + std::
      abs(pVelocity[idx].z()), WaveSpeed.z()));
}
WaveSpeed = dx / WaveSpeed;
double delT_new = WaveSpeed.minComponent();
new_dw->put(delt_vartype(delT_new), lb->delTLabel, patch->getLevel());
}
```

5. Now you are almost ready to use the submodels in the stress computation logic. To complete the
   process, you will have to add a task that makes sure any submodel internal variables are updated
   correctly during the stress computation process:

```
void
TabularPlasticity::addComputesAndRequires(Task* task, const MPMMaterial* matl,
    const PatchSet* patches) const
{
  // Add the computes and requires that are common to all explicit
  // constitutive models.  The method is defined in the ConstitutiveModel
  // base class.
  const MaterialSubset* matlset = matl->thisMaterial();
  addSharedCRForHypoExplicit(task, matlset, patches);
  task->requires(Task::OldDW, lb->pParticleIDLabel, matlset, Ghost::None);

  // Add yield Function computes and requires
  d_yield->addComputesAndRequires(task, matl, patches);

  // Add internal variable computes and requires
  task->requires(Task::OldDW, pElasticStrainLabel, matlset, Ghost::None);
  task->requires(Task::OldDW, pElasticVolStrainLabel, matlset, Ghost::None);
  task->requires(Task::OldDW, pPlasticStrainLabel, matlset, Ghost::None);
  task->requires(Task::OldDW, pPlasticCumEqStrainLabel, matlset, Ghost::None);
  task->requires(Task::OldDW, pPlasticVolStrainLabel, matlset, Ghost::None);
  task->requires(Task::OldDW, pRemoveLabel, matlset, Ghost::None);

  task->computes(pElasticStrainLabel_preReloc, matlset);
  task->computes(pElasticVolStrainLabel_preReloc, matlset);
  task->computes(pPlasticStrainLabel_preReloc, matlset);
  task->computes(pPlasticCumEqStrainLabel_preReloc, matlset);
  task->computes(pPlasticVolStrainLabel_preReloc, matlset);
  task->computes(pRemoveLabel_preReloc, matlset);
}
```

6. Finally, you can use the submodel objects for stress computation:

```
void
TabularPlasticity::computeStressTensor(const PatchSubset* patches, const
    MPMMaterial* matl, DataWarehouse* old_dw, DataWarehouse* new_dw)
{
  // Global loop over each patch
  for (int p = 0; p < patches->size(); p++) {

    // Declare and initial value assignment for some variables
    const Patch* patch = patches->get(p);

    // Initialize wave speed
    double c_dil = std::numeric_limits<double>::min();
    Vector WaveSpeed(c_dil, c_dil, c_dil);
    Vector dx = patch->dCell();

    // Initialize strain energy
    double se = 0.0;

    // Get particle subset for the current patch
    int matID = matl->getDWIndex();
    ParticleSubset* pset = old_dw->getParticleSubset(matID, patch);
```

```cpp
// Set up local particle variables to be read and written
constParticleVariable<int> pRemove;
constParticleVariable<double> pEev, pEpv, pEpeq_old;
constParticleVariable<Matrix3> pEe, pEp;
old_dw->get(pEe,        pElasticStrainLabel, pset);
old_dw->get(pEev,       pElasticVolStrainLabel, pset);
old_dw->get(pEp,        pPlasticStrainLabel, pset);
old_dw->get(pEpv,       pPlasticVolStrainLabel, pset);
old_dw->get(pEpeq_old, pPlasticCumEqStrainLabel, pset);
old_dw->get(pRemove,    pRemoveLabel, pset);

ParticleVariable<int>    pRemove_new;
ParticleVariable<double> pEev_new, pEpv_new, pEpeq_new;
ParticleVariable<Matrix3> pEe_new, pEp_new;
new_dw->allocateAndPut(pEe_new,     pElasticStrainLabel_preReloc, pset);
new_dw->allocateAndPut(pEev_new,    pElasticVolStrainLabel_preReloc, pset);
new_dw->allocateAndPut(pEp_new,     pPlasticStrainLabel_preReloc, pset);
new_dw->allocateAndPut(pEpv_new,    pPlasticVolStrainLabel_preReloc, pset);
new_dw->allocateAndPut(pEpeq_new,   pPlasticCumEqStrainLabel_preReloc, pset);
new_dw->allocateAndPut(pRemove_new, pRemoveLabel_preReloc, pset);

// Set up global particle variables to be read and written
delt_vartype delT;
constParticleVariable<long64> pParticleID;
constParticleVariable<double> pMass;
constParticleVariable<Vector> pVelocity;
constParticleVariable<Matrix3> pDefGrad, pStress_old;

old_dw->get(delT,         lb->delTLabel, getLevel(patches));
old_dw->get(pMass,        lb->pMassLabel, pset);
old_dw->get(pParticleID,  lb->pParticleIDLabel, pset);
old_dw->get(pVelocity,    lb->pVelocityLabel, pset);
old_dw->get(pDefGrad,     lb->pDefGradLabel, pset);
old_dw->get(pStress_old,  lb->pStressLabel, pset);

// Get the particle variables computed in interpolateToParticlesAndUpdate()
constParticleVariable<double> pVolume;
constParticleVariable<Matrix3> pVelGrad_new, pDefGrad_new;
new_dw->get(pVolume,      lb->pVolumeLabel_preReloc, pset);
new_dw->get(pVelGrad_new, lb->pVelGradLabel_preReloc, pset);
new_dw->get(pDefGrad_new, lb->pDefGradLabel_preReloc, pset);

ParticleVariable<double> p_q, pdTdt;
ParticleVariable<Matrix3> pStress_new;
new_dw->allocateAndPut(p_q,         lb->p_qLabel_preReloc, pset);
new_dw->allocateAndPut(pdTdt,       lb->pdTdtLabel_preReloc, pset);
new_dw->allocateAndPut(pStress_new, lb->pStressLabel_preReloc, pset);

// Loop over particles
for (particleIndex& idx : *pset) {

  // No thermal effects
  pdTdt[idx] = 0.0;

  // Compute the symmetric part of the velocity gradient
  Matrix3 DD = (pVelGrad_new[idx] + pVelGrad_new[idx].Transpose()) * .5;

  // Use polar decomposition to compute the rotation and stretch tensors
  Matrix3 FF = pDefGrad[idx];
  Matrix3 RR, UU;
  FF.polarDecompositionRMB(UU, RR);

  // Compute the unrotated symmetric part of the velocity gradient
  DD = (RR.Transpose()) * (DD * RR);

  // Compute the unrotated stress at the start of the current timestep
  Matrix3 sigma_old = (RR.Transpose()) * (pStress_old[idx] * RR);

  // Set up model state
  ModelState_Tabular state_old;
  state_old.particleID = pParticleID[idx];
  state_old.stressTensor = sigma_old;
```

```cpp
        state_old.elasticStrainTensor = pEe[idx];
        state_old.plasticStrainTensor = pEp[idx];
        state_old.ep_cum_eq = pEpeq_old[idx];

        // Compute the elastic moduli at t = t_n
        computeElasticProperties(state_old);

        // Rate-independent plastic step
        ModelState_Tabular state_new;
        bool isSuccess = rateIndependentPlasticUpdate( DD, delT, idx, pParticleID[idx
            ], state_old, state_new);

        if (isSuccess) {
          pStress_new[idx] = state_new.stressTensor; // unrotated stress at end of
              step
          pEe_new[idx] = state_new.elasticStrainTensor; // elastic strain at end of
              step
          pEp_new[idx] = state_new.plasticStrainTensor; // plastic strain at end of
              step
          pEpv_new[idx] = pEp_new[idx].Trace(); // Plastic volumetric strain at end
              of step
          pEpeq_new[idx] = state_new.ep_cum_eq; // Equivalent plastic strain at end
              of step

          // Elastic volumetric strain at end of step, compute from updated
              deformation gradient.
          // H = ln(U) => tr(H) = tr(ln(U)) = ln(det(U)) = ln(sqrt(det(FT) det(F))) =
              ln J
          pEev_new[idx] = log(pDefGrad_new[idx].Determinant()) - pEpv_new[idx];

        } else {
          // If the updateStressAndInternalVars function can't converge it will
              return false.
          // This indicates substepping has failed, and the particle will be  deleted
              .
          pRemove_new[idx] = -999;
          std::cout << "** WARNING ** Bad step, deleting particle" << " idx = " <<
              idx << " particleID = " << pParticleID[idx] << ":" << __FILE__ << ":"
              << __LINE__ << std::endl;

          pStress_new[idx] = pStress_old[idx];
          pEe_new[idx] = state_old.elasticStrainTensor; // elastic strain at start of
              step
          pEp_new[idx] = state_old.plasticStrainTensor; // plastic strain at start of
              step
          pEpv_new[idx] = pEp_new[idx].Trace();
          pEpeq_new[idx] = pEpeq_old[idx];
          pEev_new[idx] = pEe_new[idx].Trace();
        }

        // Use polar decomposition to compute the rotation and stretch tensors.
        Matrix3 FF_new = pDefGrad_new[idx];
        double Fmax_new = FF_new.MaxAbsElem();
        double JJ_new = FF_new.Determinant();
        // These checks prevent failure of the polar decomposition algorithm if [
            F_new] has some extreme values.
        if ((Fmax_new > 1.0e16) || (JJ_new < 1.0e-16) || (JJ_new > 1.0e16)) {
          pRemove_new[idx] = -999;
          proc0cout << "Deformation gradient component unphysical: [F] = " << FF <<
              std::endl;
          proc0cout << "Resetting [F]=[I] for this step and deleting particle" << "
              idx = " << idx << " particleID = " << pParticleID[idx] << std::endl;
          Identity.polarDecompositionRMB(UU, RR);
        } else {
          FF_new.polarDecompositionRMB(UU, RR);
        }

        // Compute the rotated dynamic and quasistatic stress at the end of the
            current timestep
        pStress_new[idx] = (RR * pStress_new[idx]) * (RR.Transpose());

        // Compute wave speed + particle velocity at each particle, store the
            maximum
```

```
        computeElasticProperties(state_new);
        double bulk = state_new.bulkModulus;
        double shear = state_new.shearModulus;
        double rho_cur = pMass[idx] / pVolume[idx];
        c_dil = sqrt((bulk + four_third * shear) / rho_cur);
        WaveSpeed = Vector(Max(c_dil + std::abs(pVelocity[idx].x()), WaveSpeed.x()),
            Max(c_dil + std::abs(pVelocity[idx].y()), WaveSpeed.y()), Max(c_dil + std
            ::abs(pVelocity[idx].z()), WaveSpeed.z())));

        // Compute artificial viscosity term
        if (flag->d_artificial_viscosity) {
          double dx_ave = (dx.x() + dx.y() + dx.z()) * one_third;
          double c_bulk = sqrt(bulk / rho_cur);
          p_q[idx] = artificialBulkViscosity(DD.Trace(), c_bulk, rho_cur, dx_ave);
        } else {
          p_q[idx] = 0.;
        }

        // Compute the averaged stress
        Matrix3 AvgStress = (pStress_new[idx] + pStress_old[idx]) * 0.5;

        // Compute the strain energy increment associated with the particle
        double e = (DD(0, 0) * AvgStress(0, 0) + DD(1, 1) * AvgStress(1, 1) + DD(2,
            2) * AvgStress(2, 2) + 2.0 * (DD(0, 1) * AvgStress(0, 1) + DD(0, 2) *
            AvgStress(0, 2) + DD(1, 2) * AvgStress(1, 2))) *
          pVolume[idx] * delT;

        // Accumulate the total strain energy
        se += e;

    } // End particle set loop

    // Compute the stable timestep based on maximum value of "wave speed + particle
        velocity"
    WaveSpeed = dx / WaveSpeed; // Variable now holds critical timestep (not speed)

    double delT_new = WaveSpeed.minComponent();

    // Put the stable timestep and total strain enrgy
    new_dw->put(delt_vartype(delT_new), lb->delTLabel, patch->getLevel());
    if (flag->d_reductionVars->accStrainEnergy || flag->d_reductionVars->
        strainEnergy) {
      new_dw->put(sum_vartype(se), lb->StrainEnergyLabel);
    }
  }
}
```

7. Note that the submodels are actually called inside the private methods computeElasticProperties and rateIndependentPlasticUpdate . The details are not relevant for the purposes of this manual and will be discussed in the VAANGO User manual and the VAANGO Theory Manual. It will suffice to examine the computeElasticProperties method here.

```
void
TabularPlasticity::computeElasticProperties(ModelState_Tabular& state)
{
  state.updateStressInvariants();
  state.updatePlasticStrainInvariants();
  ElasticModuli moduli = d_elastic->getCurrentElasticModuli(&state);
  state.bulkModulus = moduli.bulkModulus;
  state.shearModulus = moduli.shearModulus;
}
```

7.4.3 Creating a new model in "Models"

The additions of a model in the Models directory is similar to that discussed earlier for the Plasticity-Models . Let us examine the implementation of ElasticModuli_Tabular to get a feel for the process.

1. The header file ElasticModuli_Tabular.h in the Models directory contains the following code. Notice that the private structure BulkModulusParameters contains a TabularData object that is con-

structed from an input file upon initialization.

```cpp
#ifndef ___ELASTIC_MODULI_TABULAR_MODEL_H__
#define ___ELASTIC_MODULI_TABULAR_MODEL_H__

#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuliModel.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/TabularData.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/ModelStateBase.h>
#include <Core/ProblemSpec/ProblemSpecP.h>
#include <limits>

namespace Vaango {

class ElasticModuli_Tabular : public ElasticModuliModel
{

public:

  ElasticModuli_Tabular() = delete;
  ElasticModuli_Tabular(const ElasticModuli_Tabular& model) = delete;
  ~ElasticModuli_Tabular() = default;

  ElasticModuli_Tabular(Uintah::ProblemSpecP& ps);
  ElasticModuli_Tabular(const ElasticModuli_Tabular* model);
  ElasticModuli_Tabular& operator=(const ElasticModuli_Tabular& model) = delete;

  void outputProblemSpec(Uintah::ProblemSpecP& ps) override;

  // Get parameters
  std::map<std::string, double> getParameters() const override
  {
    std::map<std::string, double> params;
    params["G0"] = d_shear.G0;
    params["nu"] = d_shear.nu;
    return params;
  }

  // Compute the moduli
  ElasticModuli getInitialElasticModuli() const override;
  ElasticModuli getCurrentElasticModuli(const ModelStateBase* state) override;

  ElasticModuli getElasticModuliLowerBound() const override
  {
    return getInitialElasticModuli();
  }
  ElasticModuli getElasticModuliUpperBound() const override
  {
    return ElasticModuli(std::numeric_limits<double>::max(),
                         std::numeric_limits<double>::max());
  }
  /* Tangent bulk modulus parameters */
  struct BulkModulusParameters
  {
    TabularData table;
    BulkModulusParameters() = default;
    BulkModulusParameters(Uintah::ProblemSpecP& ps) : table(ps) {
      table.setup();
      table.translateIndepVar1ByIndepVar0<2>();
    }
    BulkModulusParameters(const BulkModulusParameters& bulk) {
      table = bulk.table;
    }
    BulkModulusParameters& operator=(const BulkModulusParameters& bulk) {
      if (this != &bulk) {
        table = bulk.table;
      }
      return *this;
    }
  };

  /* Tangent shear modulus parameters */
  struct ShearModulusParameters
  {
```

```
    double G0;
    double nu;
  };

  BulkModulusParameters d_bulk;
  ShearModulusParameters d_shear;

  void checkInputParameters();

  double computeBulkModulus(const double& elasticStrain, const double&
      plasticStrain) const;
  double computeShearModulus(const double& bulkModulus) const;

};
}
#endif
```

2. In the implementation file ElasticModuli_Tabular.cc , we create constructors and make sure that the model input parameters are copied to output files. Note that there are no internal variables in this model.

```
#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuli_Tabular.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/ModelState_Tabular.h>
#include <Core/Exceptions/InternalError.h>
#include <Core/Exceptions/InvalidValue.h>
#include <Core/Exceptions/ProblemSetupException.h>

using namespace Vaango;

// Construct a default elasticity model.
ElasticModuli_Tabular::ElasticModuli_Tabular(Uintah::ProblemSpecP& ps)
  : d_bulk(ps)
{
  ps->require("G0", d_shear.G0);
  ps->require("nu", d_shear.nu);

  checkInputParameters();
}

ElasticModuli_Tabular::ElasticModuli_Tabular(const ElasticModuli_Tabular* model)
{
  d_bulk = model->d_bulk;
  d_shear = model->d_shear;
}

void
ElasticModuli_Tabular::outputProblemSpec(Uintah::ProblemSpecP& ps)
{
  Uintah::ProblemSpecP elasticModuli_ps = ps->appendChild("elastic_moduli_model");
  elasticModuli_ps->setAttribute("type", "tabular");

  d_bulk.table.outputProblemSpec(elasticModuli_ps);

  elasticModuli_ps->appendElement("G0", d_shear.G0);
  elasticModuli_ps->appendElement("nu", d_shear.nu);
}
```

3. We can now implement the elastic modulus computation code. The main point of difference with the models in PlasticityModels is that we now enforce strict adherence to a particular ModelState , in this case ModelState_Tabular , to prevent the use of submodels that are not designed for a particular model. This check, of course, comes at a cost.

```
ElasticModuli
ElasticModuli_Tabular::getInitialElasticModuli() const
{
  double K = computeBulkModulus(1.0e-6, 0);
  double G = computeShearModulus(K);
  return ElasticModuli(K, G);
}

ElasticModuli
```

```
ElasticModuli_Tabular::getCurrentElasticModuli(const ModelStateBase* state_input)
{
  const ModelState_Tabular* state = dynamic_cast<const ModelState_Tabular*>(
      state_input);
  if (!state) {
    std::ostringstream out;
    out << "**ERROR** The correct ModelState object has not been passed." << " Need
        ModelState_Tabular.";
    throw Uintah::InternalError(out.str(), __FILE__, __LINE__);
  }

  // Make sure the quantities are positive in compression
  double ev_e_bar = -(state->elasticStrainTensor).Trace();
  double ev_p_bar = -(state->plasticStrainTensor).Trace();

  // Compute the elastic moduli
  double K = computeBulkModulus(ev_e_bar, ev_p_bar);
  double G = computeShearModulus(K);

  return ElasticModuli(K, G);
}
```

4. Now we implement the method for the actual bulk and shear modulus computation.

```
double
ElasticModuli_Tabular::computeBulkModulus(const double& elasticVolStrain,
                                          const double& plasticVolStrain) const
{
  double epsilon = 1.0e-6;
  DoubleVec1D pressure_lo;
  DoubleVec1D pressure_hi;
  try {
    pressure_lo = d_bulk.table.interpolate<2>({{plasticVolStrain, elasticVolStrain-
        epsilon}});
    pressure_hi = d_bulk.table.interpolate<2>({{plasticVolStrain, elasticVolStrain+
        epsilon}});
  } catch (Uintah::InvalidValue& e) {
    std::ostringstream out;
    out << "**ERROR** In computeBulkModulus:" << " elasticVolStrain = " <<
        elasticVolStrain
        << " plasticVolStrain = " << plasticVolStrain << "\n" << e.message();
    throw Uintah::InvalidValue(out.str(), __FILE__, __LINE__);
  }
  double K = (pressure_hi[0] - pressure_lo[0])/(2*epsilon);
  return K;
}

double
ElasticModuli_Tabular::computeShearModulus(const double& K) const
{
  double nu = d_shear.nu;
  double G = (nu > -1.0 && nu < 0.5) ? 1.5*K*(1.0 - 2.0*nu)/(1.0 + nu) : d_shear.G0
      ;
  return G;
}
```

5. We now add this model to the ElasticModuliModelFactory as follows:

```
#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuliModelFactory.h>
#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuli_Arena.h>
......
#include <CCA/Components/MPM/ConstitutiveModel/Models/ElasticModuli_Tabular.h>

using namespace Vaango;
ElasticModuliModel*
ElasticModuliModelFactory::create(Uintah::ProblemSpecP& ps)
{
  Uintah::ProblemSpecP child = ps->findBlock("elastic_moduli_model");
  if (!child) {
    std::ostringstream out;
    out << "**Error** No Elastic modulus model provided." << " Default (constant
        elasticity) model needs at least two input parameters." << std::endl;
    throw Uintah::ProblemSetupException(out.str(), __FILE__, __LINE__);
```

```cpp
    }

    std::string mat_type;
    if (!child->getAttribute("type", mat_type)) {
      std::ostringstream out;
      out << "MPM::ConstitutiveModel:No type provided for elasticity model.";
      throw Uintah::ProblemSetupException(out.str(), __FILE__, __LINE__);
    }

    if (mat_type == "constant")
      return (scinew ElasticModuli_Constant(child));
    .....
    else if (mat_type == "tabular")
      return (scinew ElasticModuli_Tabular(child));
    else {
      std::cerr << "**WARNING** No elasticity model provided. "
                << "Creating default (constant elasticity) model" << std::endl;
      return (scinew ElasticModuli_Constant(child));
    }
}
ElasticModuliModel*
ElasticModuliModelFactory::createCopy(const ElasticModuliModel* smm)
{
    if (dynamic_cast<const ElasticModuli_Constant*>(smm))
      return (scinew ElasticModuli_Constant( dynamic_cast<const
        ElasticModuli_Constant*>(smm)));
    ....
    else if (dynamic_cast<const ElasticModuli_Tabular*>(smm))
      return (scinew ElasticModuli_Tabular( dynamic_cast<const ElasticModuli_Tabular
        *>(smm)));
    else {
      std::cerr << "**WARNING** No elasticity model provided. " << "Creating default
        (constant elasticity) model" << std::endl;
      return (scinew ElasticModuli_Constant( dynamic_cast<const
        ElasticModuli_Constant*>(smm)));
    }
}
```

6. Next we add the new file to the compilation list in CMakeLists.txt in the Models directory:

```
SET(MPM_ConstitutiveModel_Models_SRCS
  ...............
  ${CMAKE_CURRENT_SOURCE_DIR}/ElasticModuli_Tabular.cc
)
```

7. Finally, we add the new model tags in the constitutive_models.xml file in the directory src/Stand-dAlone/inputs/UPS_SPEC:

```xml
<elastic_moduli_model  spec="OPTIONAL NO_DATA" attribute1="type REQUIRED STRING '
    constant  arenisca3  soil_model_brannon  arena  arena_mixture  tabular'"
    need_applies_to="type arena Arenisca3 camclay soil_model_brannon arena
    arena_mixture  tabular_plasticity">
  <!-- type = tabular -->
  <filename           spec="REQUIRED STRING" need_applies_to="type tabular"/>
  <independent_variables spec="REQUIRED STRING" need_applies_to="type tabular"/>
  <dependent_variables   spec="REQUIRED STRING" need_applies_to="type tabular"/>
  <interpolation        spec="REQUIRED NO_DATA" attribute1="type REQUIRED STRING
      'linear, cubic'" need_applies_to="type tabular"/>
  <nu                 spec="REQUIRED DOUBLE" need_applies_to="type tabular" />
</elastic_moduli_model>
```

7.5 Models that use tabular data

The tabular data infrastructure in the Models directory has the potential to be used for many constitutive models. A brief description of the implementation and design choices is given in this section.

The existing tabular data handling classes in VAANGO were designed to deal with multidimensional structured data. However, some MPM models need unstructured tabular data as input. The TabularData class

was designed to handle unstructured tables containing (at most) four independent variables. However, it can be extended to deal with more independent variables at the cost of added complexity.

The input tabular data are expected in the following JSON format:

```
1  {"Vaango_tabular_data": {
2     "Meta" : {
3       "title" : "Test data"
4     },
5     "Data" : {
6       "Salinity" : [0.1, 0.2],
7       "Data" : [{
8         "Temperature" : [100, 200, 300],
9         "Data" : [{
10          "Volume" : [0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8],
11          "Pressure" : [10, 20, 30, 40, 50, 60, 70, 80]
12        }, {
13          "Volume" : [0.15, 0.25, 0.35],
14          "Pressure" : [100, 200, 300]
15        }, {
16          "Volume" : [0.05, 0.45, 0.75],
17          "Pressure" : [1000, 2000, 3000]
18        }]
19      }, {
20        "Temperature" : [0, 400],
21        "Data" : [{
22          "Volume" : [0.1, 0.2, 0.3, 0.4],
23          "Pressure" : [15, 25, 35, 45]
24        }, {
25          "Volume" : [0.1, 0.45, 0.65],
26          "Pressure" : [150, 250, 350]
27        }]
28      }]
29    }
30   }
31 }
```

7.5.1 The TableContainers class

To keep things easier to understand, the TabularData class has been split into two parts: a TableContainers class and the main TabularData class.

The TableContainers class contains the following:

1. The independent and dependent variable name and data are stored in the TableVar structure which has the form:

```
struct TableVar
{
  std::string name;
  std::unordered_map<IndexKey, std::vector<double>, IndexHash, IndexEqual>
    data;
  TableVar() {}
  TableVar(const std::string name) { this->name = name; }
};
```

The data for each variable is stored as a vector of doubles and associated with a key that is determined by the indices of the associated independent variables. For easier reading, the structure TableVar is given the aliases IndependentVar and DependentVar :

```
using IndependentVar = TableVar;
using DependentVar = TableVar;
```

2. The IndexKey object that is used to locate a particular table variable data vector has an associated a IndexHash functor and an IndexEqual equality operator functor as shown below:

```
struct IndexKey
{
  std::uint8_t _ii;
```

```
      std::uint8_t _jj;
      std::uint8_t _kk;
      std::uint8_t _ll;
      IndexKey(std::uint8_t ii, std::uint8_t jj, std::uint8_t kk, std::uint8_t ll)
        : _ii(ii) , _jj(jj) , _kk(kk) , _ll(ll)
      {
      }
   };
```

```
   struct IndexEqual
   {
      bool operator()(const IndexKey& lhs, const IndexKey& rhs) const
      {
         return (lhs._ii == rhs._ii && lhs._jj == rhs._jj && lhs._kk == rhs._kk &&
                 lhs._ll == rhs._ll);
      }
   };
```

```
   struct IndexHash
   {
      std::size_t operator()(const IndexKey& key) const
      {
         std::size_t hashval = key._ii;
         hashval *= 37;
         hashval += key._jj;
         hashval *= 37;
         hashval += key._kk;
         hashval *= 37;
         hashval += key._ll;
         return hashval;
      }
   };
```

7.5.2 The TabularData class

The TabularData class header file uses, among other things, the json.hpp header-only JSON library that is included as a git submodule in Vaango . The current version only allows for linear interpolation and leaves higher order interpolation to the calling programs.

```
#ifndef VAANGO_MPM_CONSTITUTIVE_MODEL_TABULAR_DATA_H
#define VAANGO_MPM_CONSTITUTIVE_MODEL_TABULAR_DATA_H

#include <CCA/Components/MPM/ConstitutiveModel/Models/TableContainers.h>
#include <submodules/json/src/json.hpp>

namespace Vaango {

class TabularData
{
public:
  // Constructors, assignment, destructor
  TabularData() {};
  TabularData(Uintah::ProblemSpecP& ps);
  TabularData(const TabularData& table);
  ~TabularData() = default;
  TabularData& operator=(const TabularData& table);

  // Initialization and output
  void initialize();
  void outputProblemSpec(Uintah::ProblemSpecP& ps);

  // Add variables
  void addIndependentVariable(const std::string& varName);
  std::size_t addDependentVariable(const std::string& varName);

  // A setup method
  void setup();

  // A method to modify the table data if needed
```

```cpp
  template <int dim> void translateAlongNormals(const std::vector<Uintah::Vector>& vec,
      const double& shift);
  template <int dim> void translateIndepVar1ByIndepVar0();

  // File read methods
  template <int dim> void readJSONTableFromFile(const std::string& tableFile);
  template <int dim> void readJSONTable(const nlohmann::json& doc, const std::string&
      tableFile);

  // A method for interpolation
  template <int dim> DoubleVec1D interpolate(const std::array<double, dim>& indepValues)
       const;
  template <int dim> DoubleVec1D interpolateLinearSpline( const std::array<double, dim>&
       indepValues, const IndepVarPArray& indepVars, const DepVarPArray& depVars) const;

  // Some get and set methods
  std::size_t getNumIndependents() const { return d_indepVars.size(); }
  std::size_t getNumDependents() const { return d_depVars.size(); }
  const IndepVarPArray& getIndependentVars() const { return d_indepVars; }
  const DepVarPArray& getDependentVars() const { return d_depVars; }

  DoubleVec1D getIndependentVarData(const std::string& name, const IndexKey& index)
      const;
  DoubleVec1D getDependentVarData(const std::string& name, const IndexKey& index) const;
  void setIndependentVarData(const std::string& name, const IndexKey& index, const
      DoubleVec1D& data);
  void setDependentVarData(const std::string& name, const IndexKey& index, const
      DoubleVec1D& data);
private:
  std::string d_filename;
  std::string d_indepVarNames;
  std::string d_depVarNames;
  std::string d_interpType;

  IndepVarPArray d_indepVars;
  DepVarPArray d_depVars;

  // A parser for variable names
  std::vector<std::string> parseVariableNames(const std::string& vars);

  // JSON related methods
  nlohmann::json loadJSON(std::stringstream& inputStream, const std::string& fileName);
  nlohmann::json getContentsJSON(const nlohmann::json& doc, const std::string& fileName)
      ;
  std::string getTitleJSON(const nlohmann::json& contents, const std::string& fileName);
  nlohmann::json getDataJSON(const nlohmann::json& contents, const std::string& fileName
      );
  DoubleVec1D getVectorJSON(const nlohmann::json& object, const std::string key, const
      std::string& tableFile);
  DoubleVec1D getDoubleArrayJSON(const nlohmann::json& object, const std::string key,
      const std::string& tableFile);

  // Search and interpolation methods
  std::size_t findLocation(const double& value, const DoubleVec1D& varData) const;
  double computeParameter(const double& input, const std::size_t& startIndex, const
      DoubleVec1D& data) const;
  double computeInterpolated(const double& tval, const std::size_t& startIndex, const
      DoubleVec1D& data) const;
};
}
#endif // VAANGO_MPM_CONSTITUTIVE_MODEL_TABULAR_DATA_H
```

The TabularData class uses the following aliases:

```cpp
using IndexKey = TableContainers::IndexKey;

using DoubleVec1D = std::vector<double>;
using DoubleVec2D = std::vector<DoubleVec1D>;

using IndependentVarP = std::unique_ptr<TableContainers::IndependentVar>;
using DependentVarP = std::unique_ptr<TableContainers::DependentVar>;
```

```
using IndepVarPArray = std::vector<IndependentVarP>;
using DepVarPArray = std::vector<DependentVarP>;
```

7.5.3 A TabularData implementation

The current implementation of TabularData was designed with the TabularPlasticity model in mind, specifically, the need to be able to represent and interpolate between several unloading curves at different values of plastic strain.

Construction

A TabularData object is created during input file processing using the following constructor. A copy constructor and an assignment operator are also provided.

```
// Read input file and construct
TabularData::TabularData(ProblemSpecP& ps)
{
  ps->require("independent_variables", d_indepVarNames);
  ps->require("dependent_variables", d_depVarNames);
  ps->require("filename", d_filename);
  ProblemSpecP interp = ps->findBlock("interpolation");
  if (!interp) {
    d_interpType = "linear";
  } else {
    if (!interp->getAttribute("type", d_interpType)) {
      throw ProblemSetupException("**ERROR** Interpolation \
        tag needs type=linear/cubic", __FILE__, __LINE__);
    }
  }
  initialize();
}
```

The independent_variables and dependent_variables have to be specified in the format var1, var2, var3 .

The constructor calls the initialize method that parses the variable names from the input lists and allocates heap memory for the tabular data.

```
void
TabularData::initialize()
{
  std::vector<std::string> indepVarNames = parseVariableNames(d_indepVarNames);
  std::vector<std::string> depVarNames = parseVariableNames(d_depVarNames);
  for (const auto& name : indepVarNames) {
    addIndependentVariable(name);
  }
  for (const auto& name : depVarNames) {
    addDependentVariable(name);
  }
}
```

The variable name parser splits the strings in the input file using "," as a separator. Variable names may contain spaces in the middle but have to be separated by commas.

```
std::vector<std::string>
TabularData::parseVariableNames(const std::string& vars)
{
  std::vector<std::string> varNames = Vaango::Util::split(vars, ',');
  for (auto& name : varNames) {
    name = Vaango::Util::trim(name);
  }
  return varNames;
}
```

Variables are created on the stack and added to the list of independent variables.

```
void
```

```
TabularData::addIndependentVariable(const std::string& varName)
{
  d_indepVars.push_back(std::make_unique<IndependentVar>(varName));
}
```

Each set of independent variables can have more that one dependent variable. These are created next.

```
std::size_t
TabularData::addDependentVariable(const std::string& varName)
{
  // Check if dependent variable already exists
  // and just return index if true
  for (auto ii = 0u; ii < d_depVars.size(); ii++) {
    if (d_depVars[ii]->name == varName) {
      return ii;
    }
  }
  // Name does not exist, add the new variable and return its index
  d_depVars.push_back(std::make_unique<DependentVar>(varName));
  return d_depVars.size() - 1;
}
```

The copy constructor calls the assignment operator and carries out a deep copy of the data.

```
TabularData::TabularData(const TabularData& table) {
  *this = table;
}

TabularData&
TabularData::operator=(const TabularData& table) {
  if (this != &table) {
    d_indepVarNames = table.d_indepVarNames;
    d_depVarNames = table.d_depVarNames;
    d_filename = table.d_filename;
    for (auto ii = 0u; ii < table.d_indepVars.size(); ii++) {
      addIndependentVariable(table.d_indepVars[ii]->name);
      d_indepVars[ii]->data = table.d_indepVars[ii]->data;
    }
    for (auto ii = 0u; ii < table.d_depVars.size(); ii++) {
      addDependentVariable(table.d_depVars[ii]->name);
      d_depVars[ii]->data = table.d_depVars[ii]->data;
    }
  }
  return *this;
}
```

Setup

The TabularData data can be populated either by reading a file or directly using a JSON object. **Warning**: Do not expect the data to be populated automatically just because you have specified a file name in the input. You have to explicitly call the setup method in the calling program after you have created the TabularData object.

The setup method is listed below. At present we allow for only three independent variables, but more can be added with template specializations.

```
void
TabularData::setup()
{
  if (d_indepVars.size() == 1) {
    readJSONTableFromFile<1>(d_filename);
  } else if (d_indepVars.size() == 2) {
    readJSONTableFromFile<2>(d_filename);
  } else if (d_indepVars.size() == 3) {
    readJSONTableFromFile<3>(d_filename);
  } else {
    std::ostringstream out;
    out << "**ERROR**" << " More than three independent variables not allowed in " <<
        d_filename;
```

```
    throw ProblemSetupException(out.str(), __FILE__, __LINE__);
  }
}
```

Reading the JSON data from a file

The process of reading is the JSON data consists of loading the file into an input stream and then converting the stream int a JSON object:

```
template <int dim>
void
TabularData::readJSONTableFromFile(const std::string& tableFile)
{
  std::ifstream inputFile(tableFile);
  if (!inputFile) {
    std::ostringstream out;
    out << "**ERROR**" << " Cannot read tabular input data file " << tableFile;
    throw ProblemSetupException(out.str(), __FILE__, __LINE__);
  }

  std::stringstream inputStream;
  inputStream << inputFile.rdbuf();
  inputFile.close();

  json doc = loadJSON(inputStream, tableFile);
  readJSONTable<dim>(doc, tableFile);
}

json
TabularData::loadJSON(std::stringstream& inputStream, const std::string& tableFile)
{
  json doc;
  try {
    doc << inputStream;
  } catch (std::invalid_argument err) {
    std::ostringstream out;
    out << "**ERROR**" << " Cannot parse tabular input data file " << tableFile << "\n"
        << " Please check that the file is valid JSON using a linter";
    throw ProblemSetupException(out.str(), __FILE__, __LINE__);
  }
  return doc;
}
```

Reading the data from a JSON object

After the JSON object has been created, we can interpret the contents according to the number of independent variables it contains. For example, if there are three independent variables, we specialize the readJSONTable method as follows:

```
template <>
void
TabularData::readJSONTable<3>(const json& doc, const std::string& tableFile)
{
  json contents = getContentsJSON(doc, tableFile);
  std::string title = getTitleJSON(contents, tableFile);
  json data = getDataJSON(contents, tableFile);

  DoubleVec1D indepVar0Data = getDoubleArrayJSON(data, d_indepVars[0]->name, tableFile);
  d_indepVars[0]->data.insert({ IndexKey(0, 0, 0, 0), indepVar0Data });

  json data0 = getDataJSON(data, tableFile);
  for (auto ii = 0u; ii < indepVar0Data.size(); ii++) {

    DoubleVec1D indepVar1Data = getDoubleArrayJSON(data0[ii], d_indepVars[1]->name,
        tableFile);
    d_indepVars[1]->data.insert({ IndexKey(ii, 0, 0, 0), indepVar1Data });

    json data1 = getDataJSON(data0[ii], tableFile);
    for (auto jj = 0u; jj < indepVar1Data.size(); jj++) {
```

```
    int index = 0;
    for (const auto& indepVar : d_indepVars) {
      if (index > 1) {
        DoubleVec1D indepVar2Data = getDoubleArrayJSON(data1[jj], indepVar->name,
            tableFile);
        indepVar->data.insert({ IndexKey(ii, jj, 0, 0), indepVar2Data });
      }
      index++;
    }

    for (const auto& depVar : d_depVars) {
      DoubleVec1D depVar2Data = getDoubleArrayJSON(data1[jj], depVar->name, tableFile)
          ;
      depVar->data.insert({ IndexKey(ii, jj, 0, 0), depVar2Data });
    }
  }
 }
}
```

The helper functions getContentsJSON , getDataJSON , and getDoubleArrayJSON are encapsulate get
methods provided by the JSON library. Notice that the data are expected to be in the order in which the
independent variables have been specified. Checks for consistency are minimal.

Translation of the data

In some material submodels, the data have to be translated by a given amount. These are accomplished by
translateIndepVar1ByIndepVar¡2¿ for the tabular elastic moduli model, and translateAlongNormals¡1¿
for the yield condition model.

Interpolation of the data

The current implementation allows only for linear interpolation of the data. **Warning:** Extrapolation is
strictly not allowed and will cause an exception to be thrown. The interpolation method has the form

```
template <int dim>
DoubleVec1D
TabularData::interpolate(const std::array<double, dim>& indepValues) const
{
  return interpolateLinearSpline<dim>(indepValues, d_indepVars, d_depVars);
}
```

The linear interpolation method for three independent variables is listed below.

```
template <>
DoubleVec1D
TabularData::interpolateLinearSpline<3>( const std::array<double, 3>& indepValues, const
    IndepVarPArray& indepVars, const DepVarPArray& depVars) const
{
  // Get the numbers of vars
  auto numDepVars = depVars.size();

  // First find the segment containing the first independent variable value
  // and the value of parameter s
  auto indepVarData0 = getIndependentVarData(indepVars[0]->name, IndexKey(0, 0, 0, 0));
  auto segLowIndex0 = findLocation(indepValues[0], indepVarData0);
  auto sval = computeParameter(indepValues[0], segLowIndex0, indepVarData0);

  // Choose the two vectors containing the relevant independent variable data
  // and find the segments containing the data
  DoubleVec1D depValsT;
  for (auto ii = segLowIndex0; ii <= segLowIndex0 + 1; ii++) {
    auto indepVarData1 = getIndependentVarData(indepVars[1]->name, IndexKey(ii, 0, 0, 0)
        );
    auto segLowIndex1 = findLocation(indepValues[1], indepVarData1);
    auto tval = computeParameter(indepValues[1], segLowIndex1, indepVarData1);

    // Choose the last two vectors containing the relevant independent variable data
```

```
    // and find the segments containing the data
    DoubleVec1D depValsU;
    for (auto jj = segLowIndex1; jj <= segLowIndex1 + 1; jj++) {
      auto indepVarData2 = getIndependentVarData(indepVars[2]->name, IndexKey(ii, jj, 0,
          0));
      auto segLowIndex2 = findLocation(indepValues[2], indepVarData2);
      auto uval = computeParameter(indepValues[2], segLowIndex2, indepVarData2);

      for (const auto& depVar : depVars) {
        auto depVarData = getDependentVarData(depVar->name, IndexKey(ii, jj, 0, 0));
        auto depvalU = computeInterpolated(uval, segLowIndex2, depVarData);
        depValsU.push_back(depvalU);
      }
    }

    // First interpolation step
    for (auto index = 0u; index < numDepVars; index++) {
      auto depvalT = (1 - tval) * depValsU[index] + tval * depValsU[index + numDepVars];
      depValsT.push_back(depvalT);
    }
  }

  // Second interpolation step
  DoubleVec1D depVals;
  for (auto index = 0u; index < numDepVars; index++) {
    auto depval = (1 - sval) * depValsT[index] + sval * depValsT[index + numDepVars];
    depVals.push_back(depval);
  }

  return depVals;
}
```

In the above, the findLocation, computeParameter, and computeInterpolated methods are used to find the data points that are to be interpolated between. Details can be found in the VAANGO Theory Manual.
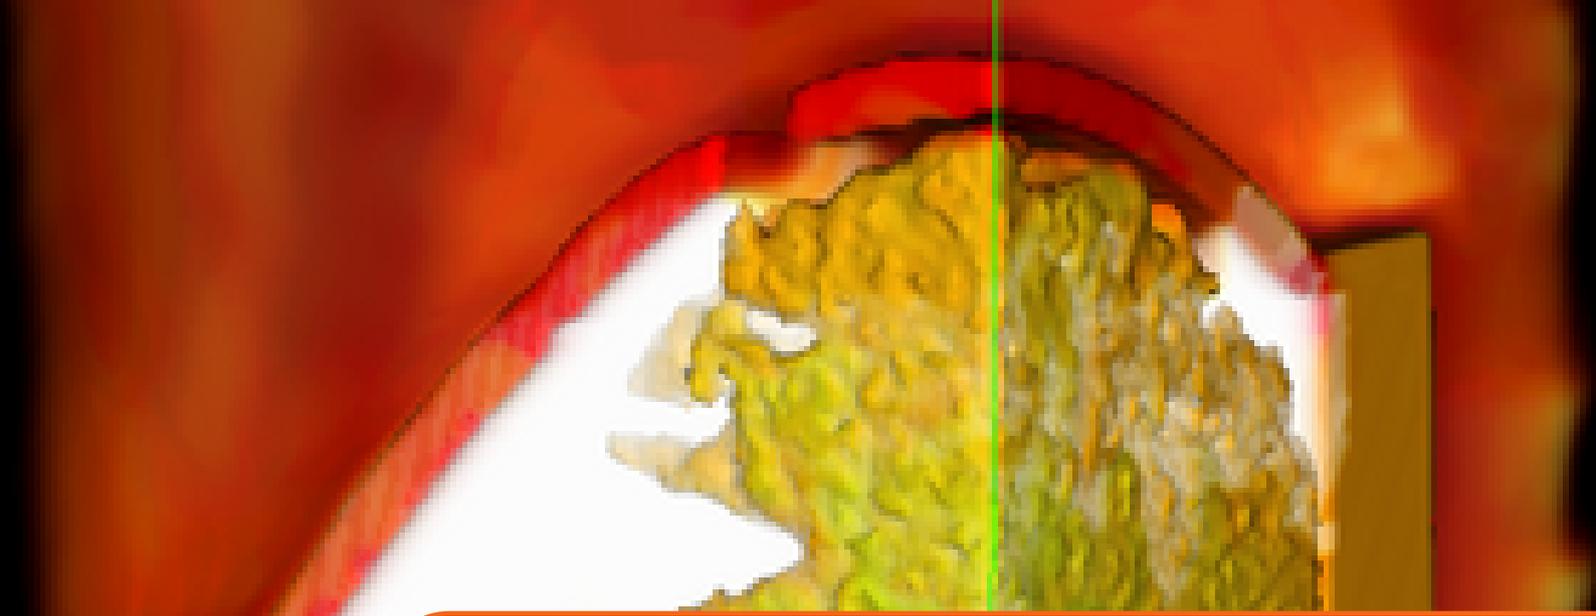
8 — Unit testing in Vaango

Older versions of Vaango used a limited form of CTest-based testing (`https://cmake.org/Wiki/CMake/Testing_With_CTest`) to unit test class APIs and utility functions. Since version 17.0, we have moved to googletest (`https://github.com/google/googletest`) and incorporated a basic version of continuous integration into the Vaango build process.

8.1 Setting up googletest

Vaango now uses the combination of CMake and googletest The reason for the shift is the convenience googletest provides.

8.1.1 Installing googletest

> To make sure that the googletest submodule is checked out when you check out ParSim, you will have to use the –recursive flag:
>
> ```
> git clone --recursive ....
> ```

googletest is installed as a git submodule (`https://git-scm.com/docs/git-submodule`) in the Vaango git repository. The submodule is added using

```
git submodule add git@github.com:google/googletest.git Vaango/src/googletest
```

Note that the typical user will not need to add a submodule to Vaango .

8.1.2 Making sure cmake finds and compiles googletest

> For the typical developer, all that is needed for unit testing to be activated is to run the following commands from a dbg or opt directory:
>
> ```
> cmake -DUSE_CLANG=1 ../src
> make -j4
> ```

To make sure that googletest is found and built during the compile process, we have added the following to root CMakeLists.txt file in the directory Vaango/src :

```
if (USE_CLANG)
```

```
  set(CMAKE_CXX_COMPILER "/usr/local/bin/clang++")
endif ()
set(gtest_force_shared_crt ON CACHE BOOL "" FORCE)
add_subdirectory( googletest)
```

The add_subdirectory command is all that is needed for cmake to compile googletest and produce static libraries. However, the developer may need to specify the absolute path to the clang compiler.

> At present the unit testing and continuous integration function is turned on by default, but only when a clang compiler is used. The absolute path to the clang compiler is needed because I have not been successful in passing the clang compiler name to googletest without specifying the full path.

8.2   Adding local unit tests

The VAANGO convention is to create a UnitTests directory in the directory of interest (e.g., src/CCA/Components/MPM/ConstitutiveModel/Models) and modify the local CMakeLists.txt (e.g., the Models/C-MakeLists.txt) to be able to find the UnitTests directory. For example,

```
SET(MPM_ConstitutiveModel_SRCS
  ${MPM_ConstitutiveModel_SRCS}
  ${CMAKE_CURRENT_SOURCE_DIR}/ModelStateBase.cc
  .....
  PARENT_SCOPE
)
IF(NOT CMAKE_COMPILER_IS_GNUCXX)
  add_subdirectory(UnitTests)
ENDIF()
```

8.2.1   The CMakeLists.txt file in UnitTests

Now we are finally ready to add our unit tests to the build chain. The CMakeLists.txt file in the UnitTests directory has the following components:

1. The tests may need to read some input files. Make sure that these files are copied to the build directory. For example, if you need to copy table_elastic.json, add the following near the top of the CMakeLists.txt file:

   ```
   configure_file(table_elastic.json table_elastic.json COPYONLY)
   ```

2. The tests will probably need to be linked with all the VAANGO shared libraries. Add these to the CMakeLists.txt file:

   ```
   SET(VAANGO_LIBS
           Vaango_Core_Thread
           Vaango_Core_DataArchive
           Vaango_Core_Grid
           Vaango_Core_Parallel
           Vaango_Core_Labels
           Vaango_Core_Util
           Vaango_Core_Math
           Vaango_Core_Disclosure
           Vaango_Core_Exceptions
           Vaango_Core_OS
           Vaango_CCA_Ports
           Vaango_CCA_Components_Parent
           Vaango_CCA_Components_DataArchiver
           Vaango_CCA_Components_LoadBalancers
           Vaango_CCA_Components_Regridder
           Vaango_Core_ProblemSpec
           Vaango_CCA_Components_SimulationController
           Vaango_CCA_Components_Schedulers
           Vaango_CCA_Components_ProblemSpecification
           Vaango_CCA_Components_Solvers
       )
   ```

3. Make sure that the googletest libraries can be found:

```
set(GTEST_INCLUDE_DIR "${CMAKE_SOURCE_DIR}/submodules/googletest/googletest/include
    ")
set(GTEST_LIB gtest_main gtest)
include_directories(${GTEST_INCLUDE_DIR})
```

Note that the two googletest libraries that we use are gtest_main and gtest.

4. Add the actual test compilation and run procedure to the CMakeLists.txt file. The test executable is built and also run during the build process.

5. For a typical unit test (e.g., testTabularData.cc) that does not use MPI, the following will have to be added to the CMakeLists.txt file:

```
add_executable(testTabularData testTabularData.cc)
target_link_libraries(testTabularData
  ${GTEST_LIB}
  ${VAANGO_LIBS}
)
set(UNIT_TEST testTabularData)
add_custom_command(
  TARGET ${UNIT_TEST}
  POST_BUILD
  COMMAND ${UNIT_TEST}
)
```

6. For a unit test (e.g., testMPITabularData.cc) that uses MPI, the following will have to be added to the CMakeLists.txt file:

```
add_executable(testMPITabularData testMPITabularData.cpp)
target_link_libraries(testMPITabularData
  ${GTEST_LIB}
  ${VAANGO_LIBS}
  ${MPI_LIBRARY}
  ....
)
set(UNIT_TEST testMPITabularData)
set(MPI_COMMAND mpirun -np 2 ${UNIT_TEST})
add_custom_command(
  TARGET ${UNIT_TEST}
  POST_BUILD
  COMMAND ${MPI_COMMAND}
)
```

> For the MPI tests, the add_executable line identifies the unit test testMPITabularData.cpp. The target_link_libraries lists the libraries that are needed for this test, i.e., the googletest libraries (GTEST_LIB), the VAANGO libraries (VAANGO_LIBS), and the MPI libraries. Then we set up a custom command (MPI_COMMAND) that uses mpirun. Finally we, add the add_custom_command line that tells cmake to run the MPI_COMMAND after the build is complete (POST_BUILD).

8.2.2 A simple test

Let us look at the test of the reverse function in the utility code testYieldCondUtils.cc. The process is as follows:

1. First, we include the required headers:

```
#include <CCA/Components/MPM/ConstitutiveModel/Models/YieldCondUtils.h>
#include <gtest/gtest.h>
```

2. Then we add the test

```
TEST(YieldCondUtilsTest, reverse)
{
  std::vector<double> array = {{1,2,3,4,5,6,7}};
  std::size_t index = 0;
  for (auto i : Vaango::Util::reverse(array)) {
    if (index == 0) EXPECT_EQ(i, 7);
```

```
      if (index == array.size()-1) EXPECT_EQ(i, 1);
      index++;
  }
}
```

3. When you run make , you will get an output of the form:

```
[ 83%] Building CXX object CCA/Components/MPM/ConstitutiveModel/Models/UnitTests/
    CMakeFiles/testYieldCondUtils.dir/testYieldCondUtils.cc.o
[ 83%] Linking CXX executable testYieldCondUtils
Running main() from gtest_main.cc
[==========] Running 1 tests from 1 test case.
[----------] Global test environment set-up.
[----------] 1 tests from YieldCondUtilsTest
[ RUN      ] YieldCondUtilsTest.reverse
[       OK ] YieldCondUtilsTest.reverse (0 ms)
[----------] 1 tests from YieldCondUtilsTest (1 ms total)
[----------] Global test environment tear-down
[==========] 1 tests from 1 test case ran. (1 ms total)
[  PASSED  ] 1 tests.
[ 83%] Built target testYieldCondUtils
```

This indicates that your unit test has passed the checks you have specified.

8.2.3 A test suite with some initial setup

For some unit tests, a lot of setup work is needed. However, that work does not need to be repeated in each test of a suite. In those situations, we can create a custom googletest test procedure. For example, in the tabular yield condition tests testYieldCond_Tabular.cc , we do the following:

1. Create a class that is derived from ::testing::Test :

```
class YieldCondTabularTest : public ::testing::Test {

protected:

  static void SetUpTestCase() {
    char currPath[2000];
    if (!getcwd(currPath, sizeof(currPath))) {
      std::cout << "Current path not found\n";
    }
    std::string json_file = std::string(currPath) + "/" + "table_yield.json";

    // Create a new document
    xmlDocPtr doc = xmlNewDoc(BAD_CAST "1.0");

    // Create root node
    xmlNodePtr rootNode = xmlNewNode(nullptr, BAD_CAST "plastic_yield_condition");
    xmlNewProp(rootNode, BAD_CAST "type", BAD_CAST "tabular");
    xmlDocSetRootElement(doc, rootNode);

    // Create a child node
    xmlNewChild(rootNode, nullptr, BAD_CAST "filename",
                BAD_CAST "table_yield.json");
    xmlNewChild(rootNode, nullptr, BAD_CAST "independent_variables",
                BAD_CAST "Pressure");
    xmlNewChild(rootNode, nullptr, BAD_CAST "dependent_variables",
                BAD_CAST "SqrtJ2");
    auto interp = xmlNewChild(rootNode, nullptr, BAD_CAST "interpolation",
                              BAD_CAST "");
    xmlNewProp(interp, BAD_CAST "type", BAD_CAST "linear");

    // Print the document to stdout
    xmlSaveFormatFileEnc("-", doc, "ISO-8859-1", 1);

    // Create a ProblemSpec
    ps = scinew ProblemSpec(xmlDocGetRootElement(doc), false);
    if (!ps) {
      std::cout << "**Error** Could not create ProblemSpec." << std::endl;
      std::cout << __FILE__ << ":" << __LINE__ << std::endl;
      exit(-1);
```

```
      }

      // Update the json file and create a circular yield function
      auto doc1 = xmlCopyDoc(doc, 1);
      auto node = xmlDocGetRootElement(doc1);
      auto xpathCtx = xmlXPathNewContext(doc1);
      const xmlChar* xpathExpr = xmlStrncatNew(BAD_CAST ".//", BAD_CAST "filename",
          -1);
      auto xpathObj = xmlXPathNodeEval(node, xpathExpr, xpathCtx);
      xmlXPathFreeContext(xpathCtx);
      auto jsonNode = xpathObj->nodesetval->nodeTab[0];

      xmlNodeSetContent(jsonNode, BAD_CAST "table_yield_circle.json");
      xmlSaveFormatFileEnc("-", doc1, "ISO-8859-1", 1);

      ps_circle = scinew ProblemSpec(xmlDocGetRootElement(doc1), false);
      if (!ps_circle) {
        std::cout << "**Error** Could not create ProblemSpec." << std::endl;
        std::cout << __FILE__ << ":" << __LINE__ << std::endl;
        exit(-1);
      }
   }

   static void TearDownTestCase() {}

   static ProblemSpecP ps;
   static ProblemSpecP ps_circle;
};
```

2. The variables that are used in the tests are declared static and initialized with nullptr:

```
ProblemSpecP YieldCondTabularTest::ps = nullptr;
ProblemSpecP YieldCondTabularTest::ps_circle = nullptr;
```

3. The actual tests are added using the TEST_F tag:

```
TEST_F(YieldCondTabularTest, constructorTest)
{
   // Create a model
   YieldCond_Tabular model(ps);

   // Copy
   YieldCond_Tabular modelCopy(&model);
   YieldCond_Tabular model_circle(ps_circle);

   // Check parameters
   auto params = model.getParameters();
   ASSERT_DOUBLE_EQ(params.at("I1_min"), -19200);
   ASSERT_DOUBLE_EQ(params.at("I1_max"), 30);
   ASSERT_DOUBLE_EQ(params.at("sqrtJ2_max"), 900);
}
```

8.2.4 A test suite with MPI

Consider the unit test testSPHParticleScatter.cpp that tests whether the scatter algorithm is working correctly. In this case we need to set up the MPI environment by deriving an class from the ::testing::Environment class. The procedure is similar to the previous examples (note that we use Boost in this example).

1. First, we include the required headers:

```
#include <SmoothParticleHydro/SmoothParticleHydro.h>
#include <gtest/gtest.h>
#include <boost/mpi.hpp>
```

2. Ideally we would like to return RUN_ALL_TESTS from our main function. However, we cannot use the boost::mpi::environment call to set up MPI (because the environment object is deleted before googletest s are run). Instead, we have to set up a custom environment to run our MPI

tests by creating a MPIEnvironment class that extends the ::testing::Environment class provided by googletest .

```cpp
class MPIEnvironment : public ::testing::Environment
{
public:
  virtual void SetUp() {
    char** argv;
    int argc = 0;
    int mpiError = MPI_Init(&argc, &argv);
    ASSERT_FALSE(mpiError);
  }
  virtual void TearDown() {
    int mpiError = MPI_Finalize();
    ASSERT_FALSE(mpiError);
  }
  virtual ~MPIEnvironment() {}
};
```

Here, the SetUp function calls MPI_Init and sets up the environment while the TearDown function calls MPI_Finalize . All tests are performed when the environment is active.

3. The main function in typically not needed in standard googletest tests and is generated by some internal magic. However, we do need a main function in testSPHParticleScatter.cpp because we are using mpirun to run the test.

```cpp
int main(int argc, char* argv[])
{
  ::testing::InitGoogleTest(&argc, argv);
  ::testing::AddGlobalTestEnvironment(new MPIEnvironment);
  return RUN_ALL_TESTS();
}
```

The MPI specific environment is created by the AddGlobalTestEnvironment function to which we pass a MPIEnvironment object. The object is deleted internally by googletest .

The tests in testSPHParticleScatter.cpp are then run using RUN_ALL_TESTS() . Note that this function returns from main and, therefore, and non-googletest environment objects created in main (e.g., boost::mpi::environment) are deleted before the tests are run.

4. Finally, we add an actual test to testSPHParticleScatter.cpp as follows:

```cpp
TEST(SPHParticleScatterTest, scatter)
{
  // Set up communicator
  boost::mpi::communicator boostWorld;
  // Set up SPH object
  SmoothParticleHydro sph;
  sph.setCommunicator(boostWorld);
  // ... Some code to create particles
  // ....
  sph.setParticles(particles);
  sph.scatterSPHParticle(domain, ghostWidth, domainBuffer);
  if (boostWorld.rank() == 0) {
    EXPECT_EQ(sph.getSPHParticleVec().size(), 100);
  } else {
    EXPECT_EQ(sph.getSPHParticleVec().size(), 200);
  }
}
```

> The Boost MPI environment created by boost::mpi::environment allows MPI calls to fail without throwing an exception. This does not create a problem if the boost::mpi::environment method is used to set up the MPI environment. But when setting up the environment explicitly for the unit tests, we should add
>
> ```cpp
> MPI_Errhandler_set(cartComm, MPI_ERRORS_RETURN);
> ```
>
> before calls that can throw exceptions.

5. Now, if we run make the output has the following form:

```
[ 69%] Built target ellip3D_lib
[ 71%] Built target paraEllip3D
[ 74%] Built target gtest
[ 76%] Built target gtest_main
.......
Scanning dependencies of target testSPHParticleScatter
[ 90%] Building CXX object SmoothParticleHydro/UnitTests/CMakeFiles/
    testSPHParticleScatter.dir/testSPHParticleScatter.cpp.o
[ 91%] Linking CXX executable testSPHParticleScatter
[==========] Running 1 test from 1 test case.
[----------] Global test environment set-up.
[==========] Running 1 test from 1 test case.
[----------] Global test environment set-up.
Set up environment
Set up environment
[----------] 1 test from SPHParticleScatterTest
[----------] 1 test from SPHParticleScatterTest
[ RUN      ] SPHParticleScatterTest.scatter
[ RUN      ] SPHParticleScatterTest.scatter
[       OK ] SPHParticleScatterTest.scatter (5 ms)
[----------] 1 test from SPHParticleScatterTest (5 ms total)

[----------] Global test environment tear-down
[       OK ] SPHParticleScatterTest.scatter (6 ms)
[----------] 1 test from SPHParticleScatterTest (6 ms total)

[----------] Global test environment tear-down
Tore down environment
[==========] 1 test from 1 test case ran. (289 ms total)
[  PASSED  ] 1 test.
Tore down environment
[==========] 1 test from 1 test case ran. (289 ms total)
[  PASSED  ] 1 test.
[ 91%] Built target testSPHParticleScatter
[ 95%] Built target gmock
[100%] Built target gmock_main
```

8.2.5 Unit testing a constitutive model

Constitutive model testing involves a lot of set up work both for the environment and the computational framework. This following example shows how we test the TabularPlasticity model.

1. First set up headers:

```cpp
#include <CCA/Components/MPM/ConstitutiveModel/Models/TabularData.h>

#include <CCA/Components/ProblemSpecification/ProblemSpecReader.h>
#include <CCA/Components/SimulationController/AMRSimulationController.h>
#include <CCA/Components/Regridder/RegridderCommon.h>
#include <CCA/Components/Solvers/SolverFactory.h>
#include <CCA/Components/Parent/ComponentFactory.h>
#include <CCA/Components/LoadBalancers/LoadBalancerCommon.h>
#include <CCA/Components/LoadBalancers/LoadBalancerFactory.h>
#include <CCA/Components/DataArchiver/DataArchiver.h>
#include <CCA/Components/Schedulers/SchedulerCommon.h>
#include <CCA/Components/Schedulers/SchedulerFactory.h>

#include <CCA/Ports/SolverInterface.h>
#include <CCA/Ports/SimulationInterface.h>
#include <CCA/Ports/Output.h>

#include <Core/Parallel/Parallel.h>
#include <Core/Parallel/UintahParallelComponent.h>
#include <Core/Malloc/Allocator.h>
#include <Core/ProblemSpec/ProblemSpec.h>
#include <Core/ProblemSpec/ProblemSpecP.h>
#include <Core/Exceptions/Exception.h>
#include <Core/Exceptions/ProblemSetupException.h>
#include <Core/Exceptions/InvalidValue.h>
```

```cpp
#include <Core/Util/Environment.h>
#include <Core/OS/Dir.h>

#include <libxml/parser.h>
#include <libxml/tree.h>

#include <iostream>
#include <map>
#include <string>
#include <vector>

#include <gtest/gtest.h>
```

2. Set up aliases:

```cpp
using namespace Vaango;
using nlohmann::json;
using Uintah::Dir;
using Uintah::ProblemSpec;
using Uintah::ProblemSpecP;
using Uintah::ProblemSpecReader;
using Uintah::Exception;
using Uintah::ProblemSetupException;
using Uintah::InvalidValue;
using Uintah::ProcessorGroup;
using Uintah::SimulationController;
using Uintah::AMRSimulationController;
using Uintah::RegridderCommon;
using Uintah::SolverInterface;
using Uintah::SolverFactory;
using Uintah::UintahParallelComponent;
using Uintah::ComponentFactory;
using Uintah::SimulationInterface;
using Uintah::LoadBalancerCommon;
using Uintah::LoadBalancerFactory;
using Uintah::DataArchiver;
using Uintah::Output;
using Uintah::SchedulerCommon;
using Uintah::SchedulerFactory;
```

3. Set up a test environment class:

```cpp
class VaangoEnv : public ::testing::Environment {
public:

  int d_argc;
  char** d_argv;
  char** d_env;

  explicit VaangoEnv(int argc, char** argv, char* env[]) {
    d_argc = argc;
    d_argv = argv;
    d_env = env;
  }

  virtual ~VaangoEnv() {}

  virtual void SetUp() {
    Uintah::Parallel::determineIfRunningUnderMPI(d_argc, d_argv);
    Uintah::Parallel::initializeManager(d_argc, d_argv);
    Uintah::create_sci_environment(d_env, 0, true );
  }

  virtual void TearDown() {
    Uintah::Parallel::finalizeManager();
  }

  static ProblemSpecP createInput() {

    char currPath[2000];
    if (!getcwd(currPath, sizeof(currPath))) {
      std::cout << "Current path not found\n";
    }
```

```cpp
//std::cout << "Dir = " << currPath << std::endl;

// Create a new document
xmlDocPtr doc = xmlNewDoc(BAD_CAST "1.0");

// Create root node
xmlNodePtr rootNode = xmlNewNode(nullptr, BAD_CAST "Uintah_specification");
xmlDocSetRootElement(doc, rootNode);

// Meta
auto meta = xmlNewChild(rootNode, nullptr, BAD_CAST "Meta", BAD_CAST "");
xmlNewChild(meta, nullptr, BAD_CAST "title", BAD_CAST "Unit test Tabular
    Plasticity");

// Simulation component
auto simComp = xmlNewChild(rootNode, nullptr, BAD_CAST "SimulationComponent",
                           BAD_CAST "");
xmlNewProp(simComp, BAD_CAST "type", BAD_CAST "mpm");

// Time
auto time = xmlNewChild(rootNode, nullptr, BAD_CAST "Time", BAD_CAST "");
xmlNewChild(time, nullptr, BAD_CAST "maxTime", BAD_CAST "1.0");
xmlNewChild(time, nullptr, BAD_CAST "initTime", BAD_CAST "0.0");
xmlNewChild(time, nullptr, BAD_CAST "delt_min", BAD_CAST "1.0e-6");
xmlNewChild(time, nullptr, BAD_CAST "delt_max", BAD_CAST "0.04");
xmlNewChild(time, nullptr, BAD_CAST "timestep_multiplier", BAD_CAST "0.3");
xmlNewChild(time, nullptr, BAD_CAST "max_Timesteps", BAD_CAST "5");

// DataArchiver
auto da = xmlNewChild(rootNode, nullptr, BAD_CAST "DataArchiver", BAD_CAST "");
xmlNewChild(da, nullptr, BAD_CAST "filebase", BAD_CAST "
    UniaxialStrainRotateTabularPlasticity.uda");
xmlNewChild(da, nullptr, BAD_CAST "outputTimestepInterval", BAD_CAST "1");
auto save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "g.mass");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.x");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.color");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.temperature");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.velocity");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.particleID");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.stress");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "p.deformationGradient");
save = xmlNewChild(da, nullptr, BAD_CAST "save", BAD_CAST "");
xmlNewProp(save, BAD_CAST "label", BAD_CAST "g.acceleration");
save = xmlNewChild(da, nullptr, BAD_CAST "checkpoint", BAD_CAST "");
xmlNewProp(save, BAD_CAST "cycle", BAD_CAST "2");
xmlNewProp(save, BAD_CAST "timestepInterval", BAD_CAST "4000");

// MPM
std::string prescribed =
  std::string(currPath) + "/" + "UniaxialStrainRotate_PrescribedDeformation.inp
      ";
auto mpm = xmlNewChild(rootNode, nullptr, BAD_CAST "MPM", BAD_CAST "");
xmlNewChild(mpm, nullptr, BAD_CAST "time_integrator", BAD_CAST "explicit");
xmlNewChild(mpm, nullptr, BAD_CAST "interpolator", BAD_CAST "linear");
xmlNewChild(mpm, nullptr, BAD_CAST "use_load_curves", BAD_CAST "false");
xmlNewChild(mpm, nullptr, BAD_CAST "minimum_particle_mass", BAD_CAST "1.0e-15")
    ;
xmlNewChild(mpm, nullptr, BAD_CAST "minimum_mass_for_acc", BAD_CAST "1.0e-15");
xmlNewChild(mpm, nullptr, BAD_CAST "maximum_particle_velocity", BAD_CAST "1.0e5
    ");
xmlNewChild(mpm, nullptr, BAD_CAST "artificial_damping_coeff", BAD_CAST "0.0");
xmlNewChild(mpm, nullptr, BAD_CAST "artificial_viscosity", BAD_CAST "true");
xmlNewChild(mpm, nullptr, BAD_CAST "artificial_viscosity_heating", BAD_CAST "
    false");
xmlNewChild(mpm, nullptr, BAD_CAST "do_contact_friction_heating", BAD_CAST "
```

```cpp
      false");
  xmlNewChild(mpm, nullptr, BAD_CAST "create_new_particles", BAD_CAST "false");
  xmlNewChild(mpm, nullptr, BAD_CAST "use_momentum_form", BAD_CAST "false");
  xmlNewChild(mpm, nullptr, BAD_CAST "with_color", BAD_CAST "true");
  xmlNewChild(mpm, nullptr, BAD_CAST "use_prescribed_deformation", BAD_CAST "true
      ");
  xmlNewChild(mpm, nullptr, BAD_CAST "prescribed_deformation_file", BAD_CAST
      prescribed.c_str());
  xmlNewChild(mpm, nullptr, BAD_CAST "minimum_subcycles_for_F", BAD_CAST "-2");
  auto ero = xmlNewChild(mpm, nullptr, BAD_CAST "erosion", BAD_CAST "");
  xmlNewProp(ero, BAD_CAST "algorithm", BAD_CAST "none");

  // Physical constants
  auto pc = xmlNewChild(rootNode, nullptr, BAD_CAST "PhysicalConstants", BAD_CAST
      "");
  xmlNewChild(pc, nullptr, BAD_CAST "gravity", BAD_CAST "[0,0,0]");

  // Material properties
  auto matProp =
    xmlNewChild(rootNode, nullptr, BAD_CAST "MaterialProperties", BAD_CAST "");
  mpm = xmlNewChild(matProp, nullptr, BAD_CAST "MPM", BAD_CAST "");
  auto mat = xmlNewChild(mpm, nullptr, BAD_CAST "material", BAD_CAST "");
  xmlNewProp(mat, BAD_CAST "name", BAD_CAST "TabularPlastic");

  // General properties
  xmlNewChild(mat, nullptr, BAD_CAST "density", BAD_CAST "1050");
  xmlNewChild(mat, nullptr, BAD_CAST "melt_temp", BAD_CAST "3695.0");
  xmlNewChild(mat, nullptr, BAD_CAST "room_temp", BAD_CAST "294.0");
  xmlNewChild(mat, nullptr, BAD_CAST "thermal_conductivity", BAD_CAST "174.0e-7")
      ;
  xmlNewChild(mat, nullptr, BAD_CAST "specific_heat", BAD_CAST "134.0e-8");
  auto cm = xmlNewChild(mat, nullptr, BAD_CAST "constitutive_model", BAD_CAST "")
      ;
  xmlNewProp(cm, BAD_CAST "type", BAD_CAST "tabular_plasticity");

  // Elastic properties
  std::string table_elastic =
    std::string(currPath) + "/" + "tabular_linear_elastic.json";
  auto elastic = xmlNewChild(cm, nullptr, BAD_CAST "elastic_moduli_model",
                             BAD_CAST "");
  xmlNewProp(elastic, BAD_CAST "type", BAD_CAST "tabular");
  xmlNewChild(elastic, nullptr, BAD_CAST "filename",
              BAD_CAST table_elastic.c_str());
  xmlNewChild(elastic, nullptr, BAD_CAST "independent_variables",
              BAD_CAST "PlasticStrainVol, TotalStrainVol");
  xmlNewChild(elastic, nullptr, BAD_CAST "dependent_variables",
              BAD_CAST "Pressure");
  auto interp_elastic = xmlNewChild(elastic, nullptr, BAD_CAST "interpolation",
                                    BAD_CAST "");
  xmlNewProp(interp_elastic, BAD_CAST "type", BAD_CAST "linear");
  xmlNewChild(elastic, nullptr, BAD_CAST "G0",
              BAD_CAST "1.0e4");
  xmlNewChild(elastic, nullptr, BAD_CAST "nu",
            BAD_CAST "0.2");

  // Yield criterion
  std::string table_yield =
    std::string(currPath) + "/" + "tabular_von_mises.json";
  auto yield = xmlNewChild(cm, nullptr, BAD_CAST "plastic_yield_condition",
                           BAD_CAST "");
  xmlNewProp(yield, BAD_CAST "type", BAD_CAST "tabular");
  xmlNewChild(yield, nullptr, BAD_CAST "filename",
              BAD_CAST table_yield.c_str());
  xmlNewChild(yield, nullptr, BAD_CAST "independent_variables",
              BAD_CAST "Pressure");
  xmlNewChild(yield, nullptr, BAD_CAST "dependent_variables",
              BAD_CAST "SqrtJ2");
  auto yield_interp = xmlNewChild(yield, nullptr, BAD_CAST "interpolation",
                                  BAD_CAST "");
  xmlNewProp(yield_interp, BAD_CAST "type", BAD_CAST "linear");

  // Geometry
  auto geom = xmlNewChild(mat, nullptr, BAD_CAST "geom_object", BAD_CAST "");
```

```cpp
      auto box = xmlNewChild(geom, nullptr, BAD_CAST "box", BAD_CAST "");
      xmlNewProp(box, BAD_CAST "label", BAD_CAST "Plate1");
      xmlNewChild(box, nullptr, BAD_CAST "min", BAD_CAST "[0.0,0.0,0.0]");
      xmlNewChild(box, nullptr, BAD_CAST "max", BAD_CAST "[1.0,1.0,1.0]");
      xmlNewChild(geom, nullptr, BAD_CAST "res", BAD_CAST "[1,1,1]");
      xmlNewChild(geom, nullptr, BAD_CAST "velocity", BAD_CAST "[0,0,0]");
      xmlNewChild(geom, nullptr, BAD_CAST "temperature", BAD_CAST "294");
      xmlNewChild(geom, nullptr, BAD_CAST "color", BAD_CAST "0");

      // Contact
      auto contact = xmlNewChild(mpm, nullptr, BAD_CAST "contact", BAD_CAST "");
      xmlNewChild(contact, nullptr, BAD_CAST "type", BAD_CAST "null");
      xmlNewChild(contact, nullptr, BAD_CAST "materials", BAD_CAST "[0]");
      xmlNewChild(contact, nullptr, BAD_CAST "mu", BAD_CAST "0.1");

      // Grid
      auto grid = xmlNewChild(rootNode, nullptr, BAD_CAST "Grid", BAD_CAST "");
      xmlNewChild(grid, nullptr, BAD_CAST "BoundaryConditions", BAD_CAST "");
      auto level = xmlNewChild(grid, nullptr, BAD_CAST "Level", BAD_CAST "");
      box = xmlNewChild(level, nullptr, BAD_CAST "Box", BAD_CAST "");
      xmlNewProp(box, BAD_CAST "label", BAD_CAST "1");
      xmlNewChild(box, nullptr, BAD_CAST "lower", BAD_CAST "[-2,-2,-2]");
      xmlNewChild(box, nullptr, BAD_CAST "upper", BAD_CAST "[3,3,3]");
      xmlNewChild(box, nullptr, BAD_CAST "resolution", BAD_CAST "[5,5,5]");
      xmlNewChild(box, nullptr, BAD_CAST "extraCells", BAD_CAST "[0,0,0]");
      xmlNewChild(box, nullptr, BAD_CAST "patches", BAD_CAST "[1,1,1]");

      // Print the document to stdout
      //xmlSaveFormatFileEnc("-", doc, "ISO-8859-1", 1);
      std::string ups_file = std::string(currPath) + "/" +
                             "UniaxialStrainRotateTabularPlasticity.ups";
      xmlSaveFormatFileEnc(ups_file.c_str(), doc, "ISO-8859-1", 1);

      // Create a ProblemSpec
      ProblemSpecP ps = scinew ProblemSpec(xmlDocGetRootElement(doc), false);
      if (!ps) {
        std::cout << "**Error** Could not create ProblemSpec." << std::endl;
        std::cout << __FILE__ << ":" << __LINE__ << std::endl;
        exit(-1);
      }

      return ps;
  }
};
```

4. Set up the entry function for the test:

```cpp
int main(int argc, char** argv, char* env[]) {
  ::testing::InitGoogleTest(&argc, argv);
  ::testing::AddGlobalTestEnvironment(new VaangoEnv(argc, argv, env));
  return RUN_ALL_TESTS();
}
```

5. Create the test:

```cpp
TEST(TabularPlasticityTest, singleParticleTest)
{
  char currPath[2000];
  if (!getcwd(currPath, sizeof(currPath))) {
    std::cout << "Current path not found\n";
  }
  std::string ups_file = std::string(currPath) + "/" +
                         "UniaxialStrainRotateTabularPlasticity.ups";
  //std::cout << "Created Filename = " << ups_file << std::endl;

  // Remove existing uda
  std::string uda_file = std::string(currPath) + "/" +
                         "UniaxialStrainRotateTabularPlasticity.uda.000";
  Dir::removeDir(uda_file.c_str());

  char * start_addr = (char*)sbrk(0);
  bool thrownException = false;
```

```cpp
  try {
    ProblemSpecP ups = VaangoEnv::createInput();
    ups->getNode()->_private = (void *) ups_file.c_str();
    //std::cout << "Filename = " << static_cast<char*>(ups->getNode()->_private) <<
        std::endl;

    const ProcessorGroup* world = Uintah::Parallel::getRootProcessorGroup();
    SimulationController* ctl = scinew AMRSimulationController(world, false, ups);

    RegridderCommon* reg = 0;
    SolverInterface* solve = SolverFactory::create(ups, world, "");

    UintahParallelComponent* comp = ComponentFactory::create(ups, world, false, "")
        ;
    SimulationInterface* sim = dynamic_cast<SimulationInterface*>(comp);
    ctl->attachPort("sim", sim);
    comp->attachPort("solver", solve);
    comp->attachPort("regridder", reg);

    LoadBalancerCommon* lbc = LoadBalancerFactory::create(ups, world);
    lbc->attachPort("sim", sim);

    DataArchiver* dataarchiver = scinew DataArchiver(world, -1);
    Output* output = dataarchiver;
    ctl->attachPort("output", dataarchiver);
    dataarchiver->attachPort("load balancer", lbc);
    comp->attachPort("output", dataarchiver);
    dataarchiver->attachPort("sim", sim);

    SchedulerCommon* sched = SchedulerFactory::create(ups, world, output);
    sched->attachPort("load balancer", lbc);
    ctl->attachPort("scheduler", sched);
    lbc->attachPort("scheduler", sched);
    comp->attachPort("scheduler", sched);
    sched->setStartAddr( start_addr );
    sched->addReference();

    ctl->run();
    delete ctl;

    sched->removeReference();
    delete sched;
    delete lbc;
    delete sim;
    delete solve;
    delete output;

  } catch (ProblemSetupException& e) {
    std::cout << e.message() << std::endl;
    thrownException = true;
  } catch (Exception& e) {
    std::cout << e.message() << std::endl;
    thrownException = true;
  } catch (...) {
    std::cout << "**ERROR** Unknown exception" << std::endl;
    thrownException = true;
  }
}
```
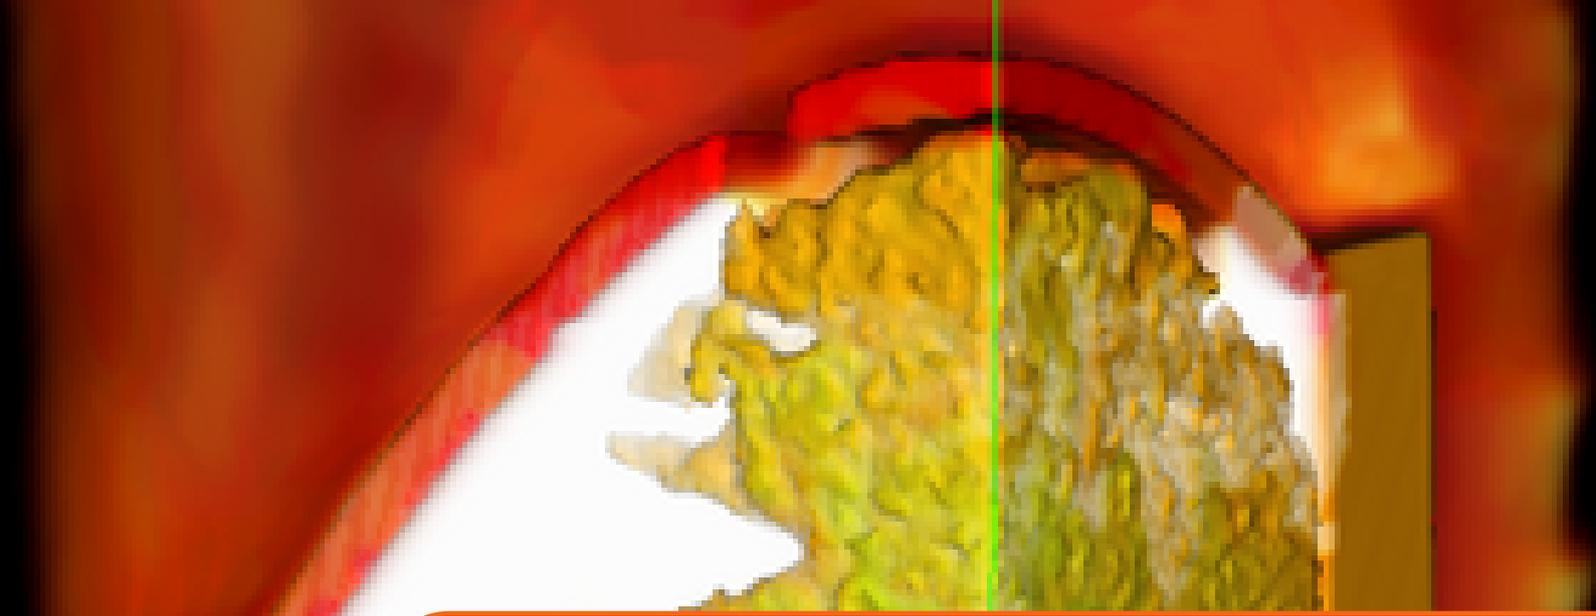
9 — Regression testing in Vaango

This chapter provides documentation for the VAANGO regression testers and the helper scripts that they call. The regression tester code is primarily in Python with a few interspersed shell scripts and is located at

```
Vaango/src/R_tester
```

The files in this structure are listed below:

```
|-- CMakeLists.txt
|-- Examples.py
|-- GPU.py
|-- ICE.py
|-- IMPM.py
|-- Models.py
|-- MPMF.py
|-- MPMICE.py
|-- MPM.py
|-- UCF.py
|-- helpers
|    |-- cleanup
|    |-- compare_dat_files.pl
|    |-- compare_dats
|    |-- compare_vaango_runs
|    |-- compare_vaango_runs_by_udas
|    |-- highwater_percent.pl
|    |-- mem_leak_check
|    |-- modUPS.py
|    |-- performance_check
|    |-- performance_percent.pl
|    |-- plotMemUsage.gp
|    |-- plotRunTimeHistory.gp
|    |-- replace_all_GS
|    |-- replace_gold_standard
|    |-- runComponentTests
|    |-- runTests
|    |-- runVaangoTests.py
|    |-- selectComponents.sh
|-- toplevel
|    |-- generateGoldStandards.py
```

The location of the place where the regression tests are actually run is the <build_dir > which is typically called dbg—opt and the regression tester makefiles and other copied scripts are in

```
Vaango/<build\_dir >/R_tester
```

9.1 Basic operation

9.1.1 Creating gold standards

The first step in setting up a local version of the regression tester is to create gold standards based on a tested and working version of the code. To do that, go to the <build_dir > and run

```
make gold_standard
```

The gold standards are saved in

```
Vaango/\textless build\_dir >/R_tester/gold_standards
```

9.1.2 Running regression tests

To run all the regression tests from the <build_dir >, you can do

```
make localRT
```

will either be in relation to the tester root, which currently is

9.1.3 Viewing regression tester results

Regression tester output is locally stored in the <build_dir >/R_Tester/localRTData directory. For example, MPM related test results will be stored in the MPM-results directory.

There will be many log files saved. They will show you the resulting output of the test, which (hopefully) explains clearly why the test failed.

9.1.4 Adding new regression tests

After a new model has been developed, it is expected that you will add a regression test for the model to the tester. To do that go to the Vaango/src/R_Tester directory and select the component to which you want to add a new test. For example, if you want to add a new test called const_test_tabular that uses two MPI processes to MPM.py , edit the python script and add the following to the NIGHTLYTESTS array:

```
NIGHTLYTESTS = [
  .........
  ("const_test_viscoelastic_fortran.ups",
   "ViscoElastic/const_test_viscoelastic_fortran.ups",
    1,  "Linux", ["exactComparison"] ), \
  .........
  ("const_test_tabular.ups",
   "const_test_tabular.ups",
    4,  "Linux", ["exactComparison"] )
]
```

If you create a new component and want to add a set of tests for that component, follow the procedure used in the other test files. For a better understanding of the details, notice that NIGHTLYTESTS is a list of tests. There can be as many tests as you like, but each test must be a list of 5-6 items. For example, in MPM we find one of the tests:

```
("heatcond2mat", "heatcond2mat.ups", "", 1, "ALL")
```

The first parameter is the name of the test. It doesn't really matter what the name is, most people just use the name of the ups file, but doesn't need to be.

The second parameter is the name of the ups file. The third is any additional flags that you want to pass to vaango. The fourth is the number of processors to use. The fifth is which operating system to run on. Current choices are "Linux", "IRIX64", or "ALL" - these are case-sensitive.

Next, you'll notice the exit(runVaangoTests(...)) call at the very bottom. All you need to do here is change the third argument to <algo >, where algo is the algorithm directly passed into vaango, like "ice" or "mpm", AND the uppercase version of algo must be the name of a directory with ups files (ICE, MPM).

The callback function, if passed into runVaangoTests will be called before executing every test. It must have the same parameters as the runVaangoTest function, found in runVaangoTests.py.

Then add that file to the cvs repository and commit it, and the regression tester will see it and take care of the rest!

9.2 Regression test scripts

The main script that is used in runVaangoTests.py . Scripts that are further down the script hierarchy

9.2.1 cleanup

```
cleanup - Vaango/R_Tester/helpers/cleanup (sh script)
```

——- This script can be called from the command line in the <build_dir > . It tries to remove all Vaango.* in the directory that don't have a *.lock link pointing to it.

Previously, any directory that had a .lock file pointing to it had a file in BUILDROOT/ called lock. Cleanup removes those lock files, and creates one for every folder that is pointed to by a .lock file.

Moves all the (non-locked) Vaango .* dirs to the "trash" dir and then tries to remove them. Tries to delete everything in trask. If not successful (due to not owning the files), send email to people to clean it up.

9.2.2 MPM.py—ICE.py etc.

These are called from runVaangoTests.py . Each test is defined in an array

```
TESTS = [   (nameOfTests, upsInputFile, extraFlags, numProcs, whichOS),
            ...,
    ]
```

Then it calls

```
runVaangoTests(argv, TESTS, algo)
```

You can create your own test file, and it will automatically run in the tester. It will need 755 permissions. Example:

```python
#!/usr/bin/python

from sys import argv,exit
from helpers.runVaangoTests import runVaangoTests

TESTS = [   ("advect", "advect.ups", "", 1, "ALL"),      \
            ("hotBlob2mat", "hotBlob2mat.ups", "", 1, "ALL"),      \
            ("hotBlob2mat_sym", "hotBlob2mat_sym.ups", "", 1, "ALL"),      \
            ("hotBlob8patch", "hotBlob2mat8patch.ups", "", 8, "ALL"),      \
        ]

exit(runVaangoTests(argv, TESTS, "ice"))
```

You can also call modUPS to use the same ups file with different settings:

```
from helpers.modUPS import modUPS
test8patch_ups = modUPS("%s/MPM" % inputs_root(), \
                        "test_table.ups", \
                        ["<patches>[2,2,2]</patches>", \
                         "<maxTime>10.0</maxTime>"])
```

This will change the test_table.ups patches from whatever they were to [2,2,2] and the maxTime to 10.0. Make sure to put the entire field from the beginning of the tag to the close of the tag.

modUPS will scan through the ups file and look for the pattern of the tag secified. I.e., in the example above, it will look for

```
<patches>...</patches>
```

and then replace it with

```
<patches>[2,2,2]</patches>.
```

It can also replace tags in form

```
<varname=something />.
```

9.2.3 runVaangoTests.py

runVaangoTests.py is divided largely into 2 functions, runVaangoTests , and runVaangoTest . runVaangoTests takes the tests passed into it from MPM, ICE, etc. and executes vaango on them. It then compares the results of the output with the saved goldStandard.

usage:

```
runVaangoTests(argv, TESTS, testtype, callback)
```

where argv is the same as the arguments passed into MPM, ICE etc, vaangodir inputs goldStandard mode maxParallelism test, TESTS is a list of tests, testtype is the algorithm to run - mpm, ice, mpmice, examples, etc, callback is an optional parameter. If you pass a callback, you must define that function, which has parameters: test, vaangodir, inputsdir, compare_root, algo, mode, max_parallelism Callback is called before running each test.

Files created by runVaangoTests the the R_Tester/localRTData directory include, the directories <resultsdir > (MPM-results, ICE-results, etc.). In the <resultsdir >, the following files are created:

```
<resultsdir>/<testname>
<resultsdir>/<testname>/<testname>.uda.000
<resultsdir>/<testname>/replace_gold_standard
<resultsdir>/<testname>/compare_vaango_runs.log.txt
<resultsdir>/<testname>/timestamp
<resultsdir>/<testname>/vaango.log.txt
<resultsdir>/<testname>/restart/<testname>.uda.000
<resultsdir>/<testname>/restart/compare_vaango_runs.log.txt
<resultsdir>/<testname>/restart/timestamp
<resultsdir>/<testname>/restart/vaango.log.txt
```

runVaangoTests basically does a lot of error checking. It proceeds to iterate through the tests, and if there was a solotest arg specified, it will check and see if that is the current test. Then it compares the OS specified of the current test, and if they are not compatible it continues with the next test. Then it calls the callback if there is one, and calls runVaangoTest . Then it repeats that process for the restart test (if the Vaango test passed) AND the files created in the resultsdir directory will then be created in the restart directory.

runVaangoTest is responsible for the running and testing of one individual test. It first checks to see if the test requires more processors than the tester's max paralellism setting (passed into startTester with the -j flag). It proceeds to run the tests, noting that at any point if a test fails, it displays where the log message is that will be helpful, and then returns immediately. Vaango is then run with the parameters as set by the test, and sends its output to vaango.log.txt . If it passes, it proceeds to compare the udas generated by vaango with the gold standard, with compare_vaango_runs , which saves its output in compare_vaango_runs.log.txt . If we are in debug mode, it tests for memory leaks with mem_leak_check which saves its output in mem_leak_check.log.txt .

9.2.4   compare_vaango_runs - sh script

This script compares output of vaango with gold standard and can be called using

```
compare_vaango_runs <testname> <testdir> <compare_root> <vaangodir>
```

where testname is the name of test, testdir is the directory where test is, compare_root is the folder where testdir/testname should be - i.e, OS/TestData/mode/ALGO, vaangodir is the folder where vaango and compare_uda are located.

If there is no data for the specified test in compare_root , compare_vaango_runs creates the necessary directories and copies the output from vaango (in testdir/testname) over to be the new gold standard. Otherwise, it calls compare_dats and compare_udas on the input.

9.2.5   mem_leak_check - sh script

This script checks and compares memory usage and can be called with

```
mem_leak_check <testname> <stats_file> <compare_stats_file> <tmpdir>
```

where stats_file is the name of malloc output file -i.e., malloc_stats , compare_stats_file is the name/location of file to compare to.

Looks for memory leaks in stats_file , and outputs them to tmpdir/scinew_malloc_stats . Then calls highwater_percent.pl to check memory usage. If the percentage returned is greater than 10, an error is posted, if it returns a percentage less that -5%, then it replaces the stats in the goldStandard.

9.2.6   compare_dats - sh script

It checks to see if there are any .dat files to compare. If there are, it calls compare_dat_files.pl to compare them.

```
compare_dats <testname> <testdir> <compare_udadir> <tmpdir>
```

where compare_uda_dir is the compare_root/testname/udadir.

9.2.7   compare_udas sh script

Simply calls compare_uda (executable in vaangodir), with the flags given

```
compare_udas <vaangodir> <test> <test_udadir> <compare_udadir> <tmp_dir> <flags>
```

9.2.8   compare_dat_files.pl - perl script

Checks to see if values in dat files and gold standard dat files are within a tolerable percent of error.

```
compare_dat_files.pl <abs error> <rel error> <udadir> <compare_udadir>
                     <dat filenames>
```

where abs error - Absolute error - the max allowable exact difference, rel error - Relative error - the max allowable relative difference - based on significant figures, udadir - Uda directory with data files, compare_udadur - Gold standard uda directory, dat filenames - list of .dat filenames.

9.2.9   highwater_percent.pl - perl script

Compares memory allocated by the two input files and calculates the percent difference.

```
highwater_percent.pl <malloc_stats_file> <compare_malloc_file>
```

9.2.10 **replace_gold_standard - sh script**

Script created by runVaangoTests . To be run at the command line, when deemed necessary. It replaces the files in the gold standard with the ones created by this test.

```
(A) calls (B)
(A)  replace_gold_standard
(B)  replace_gold_standard gold_standard_dir replacement_dir testname
```

where gold_standard_dir is the gold-standard directory, replacement_dir is the directory to replace gold standard with testname is the name of the test

9.2.11 **Other files**

These are items that are not scripts themselves, but may give a little insight to the tester.

```
__init__.py - \Vaango/R_Tester/helpers/__init__.py
```

The existence of this file alone (it is virtually empty), helps to create a python package for helpers. This makes it so the python scripts can call each other's functions.
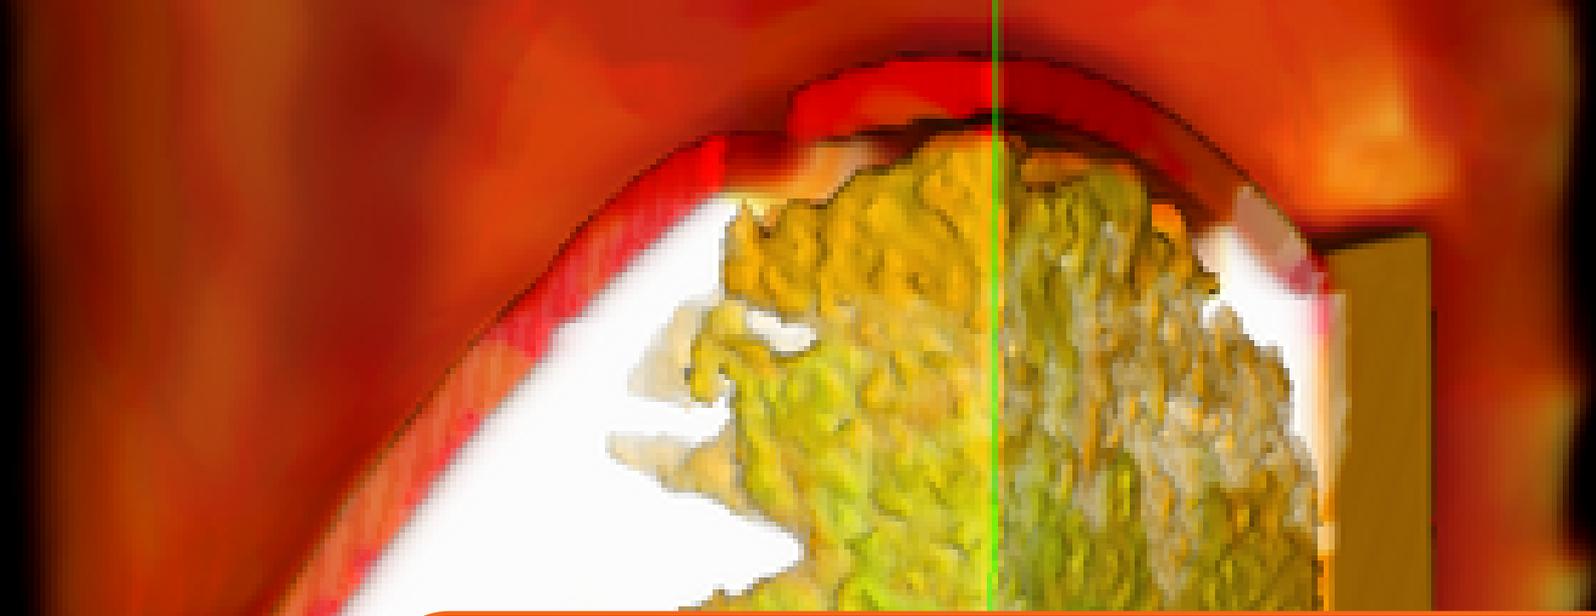
10 — Debugging

Debugging multi-processor software can be very difficult. Below are some hints on how to approach this.

10.1 Debugger

Use a debugger such as GDB to attach to the running program. Note, if you are an Emacs user, running GDB through Emacs' gdb-mode makes debugging even easier!

10.1.1 Serial Debugging

Whenever possible, debug vaango running in serial (ie. no MPI). This is much easier for many reasons (though may also not be possible in many situations). To debug serially, use:

```
gdb vaango <input_file>
```

10.1.2 Parallel Debugging

If it is necessary to use a parallel run of vaango to debug, a few things must take place.

1) The helpful macro `WAIT_FOR_DEBUGGER();` must be inserted into the code (vaango.cc) just before the real execution of the components begins ctl->run():

```
WAIT_FOR_DEBUGGER();
ctl->run();
```

Recompile the code (will only take a few seconds), then run through MPI as normal.

Caveats

1. Unless you know that the error occurs on a specific processor, you will need to attach a debugger to every process... so it behooves the developer to narrow the problem down to as few processors as possible.
2. If you did know the processor, then the `WAIT_FOR_DEBUGGER()` will need to be placed inside an 'if' statement to check for the correct MPI Rank.
3. Try to use only one node. If you need more 'processors', you (most likely) can double up on the same machine. If vaango must span multiple nodes, then you will need to log into each node separately to attach debuggers (see below).

10.1.3   **Running Vaango**

Run vaango :

```
mpirun -np 4 -m mpihosts vaango inputs/MPM/thickCylinderMPM.ups
```

Create a new terminal window for each processor you are using. (In the above case, 4 windows.)

You will notice that the output from vaango will include lines (one for each processor used) like this:

```
updraft2:22793 waiting for debugger
```

Attach to each process in the following manner (where PID is the process id number (22793 in the above case)):

```
cd Uintah/bin/StandAlone
gdb -p <PID>
```

Within GDB, you will then need to break each process out of the WAIT loop in the following manner:

```
up 2
set wait=false
cont
```

The above GDB commands must be run from each GDB session. Once all the WAITs are stopped, vaango will begin to run.
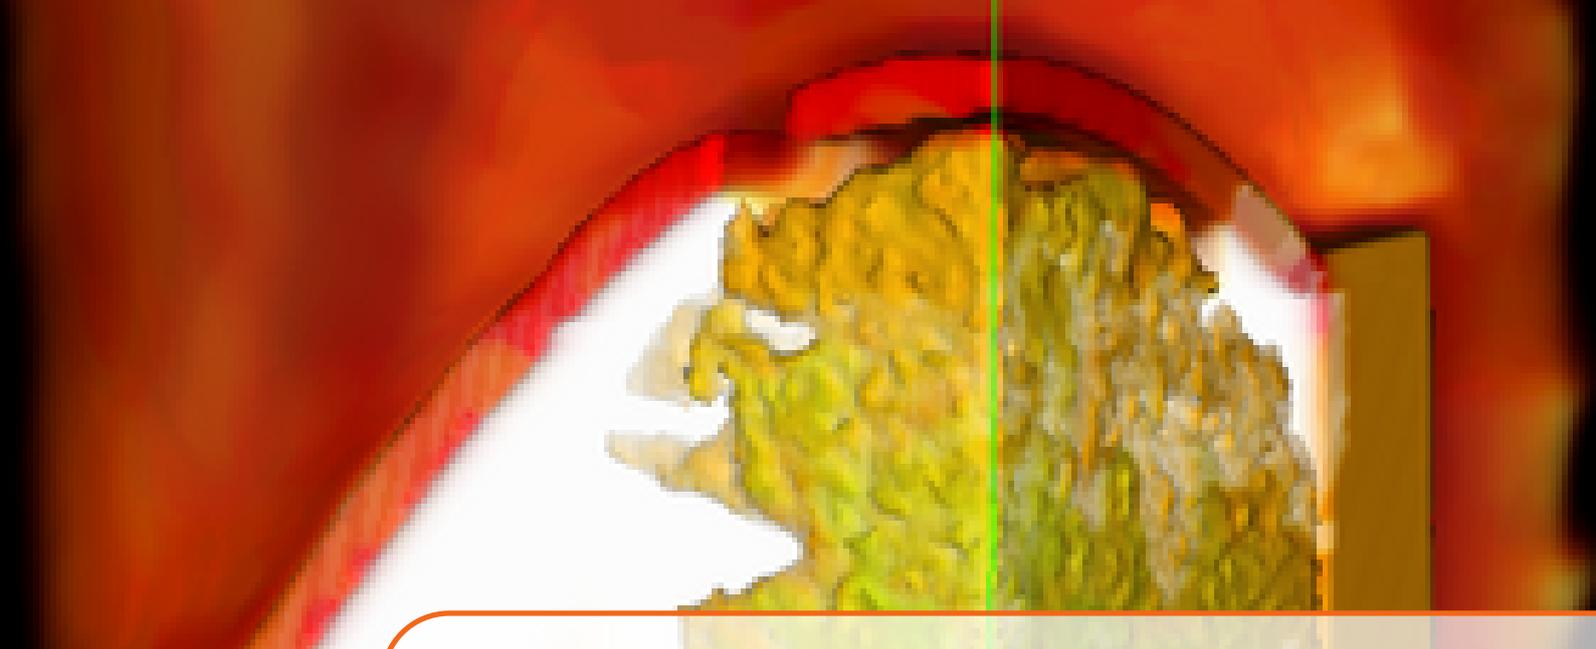
Bibliography

[1] Y Guo and JA Nairn. "Three-dimensional dynamic fracture analysis using the material point method". In: *Computer Modeling in Engineering and Sciences* 6 (2004), pages 295–308 (cited on page 38).

[2] J. E. Guilkey and J. A. Weiss. "Implicit time integration for the material point method: Quantitative and algorithmic comparisons with the finite element method". In: *Int. J. Numer. Meth. Engng.* 57.9 (2003), pages 1323–1338 (cited on page 38).

[3] M. A. Homel, J. E. Guilkey, and R. M. Brannon. "Continuum effective-stress approach for high-rate plastic deformation of fluid-saturated geomaterials with application to shaped-charge jet penetration". In: *Acta Mechanica* (2015), pages 1–32 (cited on pages 51, 52).

[4] J. O. Coplien. *Advanced C++ Programming Styles and Idioms*. Reading, MA: Addison-Wesley, 1992 (cited on page 61).