

Finite Element Assembly on Arbitrary Meshes

Matthew Knepley

Computation Institute
University of Chicago

Department of Molecular Biology and Physiology
Rush University Medical Center

Department of Applied Mathematics and Computational Science
King Abdullah University of Science and Technology

Apr 5, 2010



Outline

- 1 Introduction
- 2 Mesh Distribution
- 3 Unifying Paradigm
- 4 Finite Element Assembly

Collaborators

- Automated FEM
 - Andy Terrel (UT Austin)
 - Ridgway Scott (UChicago)
 - Rob Kirby (Texas Tech)
- Sieve
 - Dmitry Karpeev (ANL)
 - Peter Brune (UChicago)
 - Anders Logg (Simula)
- PyLith
 - Brad Aagaard (USGS)
 - Charles Williams (NZ)

Outline

- 1 Introduction
- 2 Mesh Distribution**
 - Sieve
 - Section
 - Completion
 - Distribution
 - Interfaces
- 3 Unifying Paradigm
- 4 Finite Element Assembly

Main Point

Rethinking meshes

produces a simple FEM
interface

and good code reuse.

Main Point

Rethinking meshes
produces a simple FEM
interface
and good code reuse.

Main Point

Rethinking meshes
produces a simple FEM
interface
and good code reuse.

Problems

The biggest problem in scientific computing is **programmability**:

- Lack of usable implementations of modern algorithms
 - Unstructured Multigrid
 - Fast Multipole Method
- Lack of comparison among classes of algorithms
 - Meshes
 - Discretizations

We should reorient thinking from

- characterizing the solution (FEM)
 - “what is the convergence rate (in h) of this finite element?”

to

- characterizing the computation (FErari)
 - “how many digits of accuracy per flop for this finite element?”

Problems

The biggest problem in scientific computing is **programmability**:

- Lack of widespread implementations of modern algorithms
 - Unstructured Multigrid
 - Fast Multipole Method
- Lack of comparison among classes of algorithms
 - Meshes
 - Discretizations

We should reorient thinking from

- characterizing the solution (FEM)
 - “what is the convergence rate (in h) of this finite element?”

to

- characterizing the computation (FERari)
 - “how many digits of accuracy per flop for this finite element?”

Outline

2 Mesh Distribution

- Sieve
- Section
- Completion
- Distribution
- Interfaces

Sieve

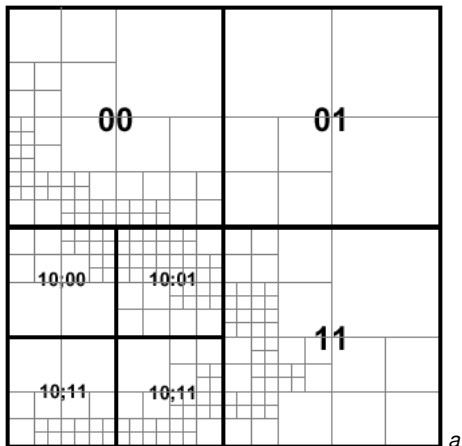
Sieve is an interface for

- general topologies
- functions over these topologies (bundles)
- traversals

One relation handles all hierarchy

- Vast reduction in complexity
 - Dimension independent code
 - A single communication routine to optimize
- Expansion of capabilities
 - Partitioning and distribution
 - Hybrid meshes
 - Complicated structures and embedded boundaries
 - Unstructured multigrid

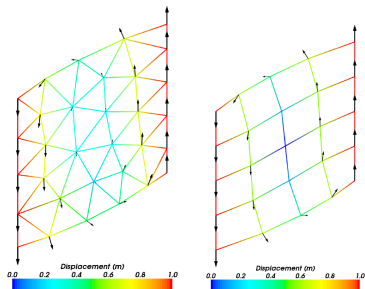
Mesh Databases



- “Most”
 - Sp
 - Str
 - Co
- Topolog
 - Sin
 - Sin
 - Ca

^aLawler, Kalé

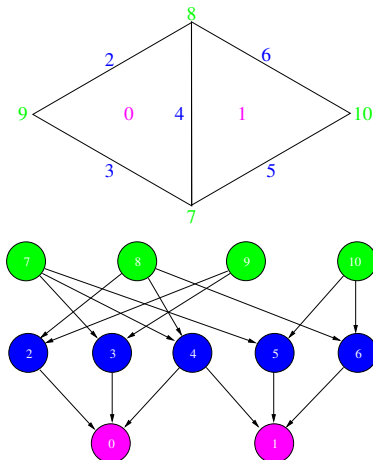
Mesh Databases



- “Most”
 - Sp
 - Str
 - Co
- Topolog
 - Sin
 - Sin
 - Ca

^aAagaard, Knepley, Williams

Doublet Mesh

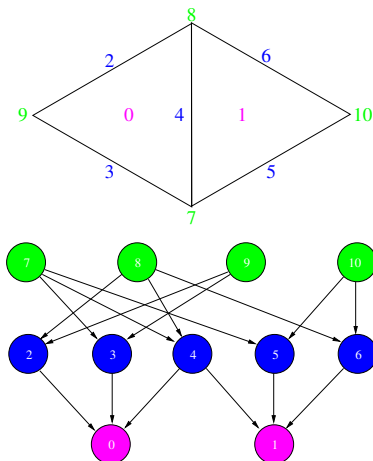


- Incidence/covering arrows

- $\text{cone}(0) = \{2, 3, 4\}$

- $\text{support}(7) = \{2, 3\}$

Doublet Mesh

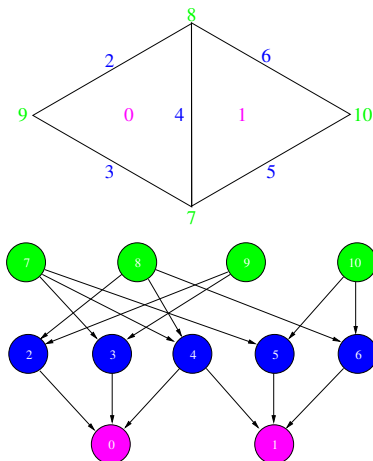


- Incidence/covering arrows

- $cone(0) = \{2, 3, 4\}$

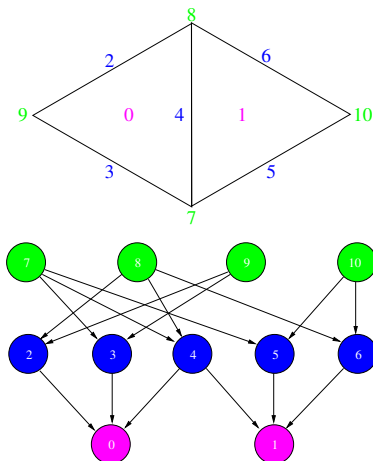
- $support(7) = \{2, 3\}$

Doublet Mesh



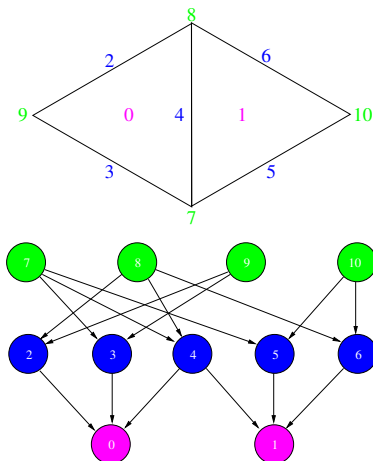
- Incidence/covering arrows
- $\text{cone}(0) = \{2, 3, 4\}$
- $\text{support}(7) = \{2, 3\}$

Doublet Mesh



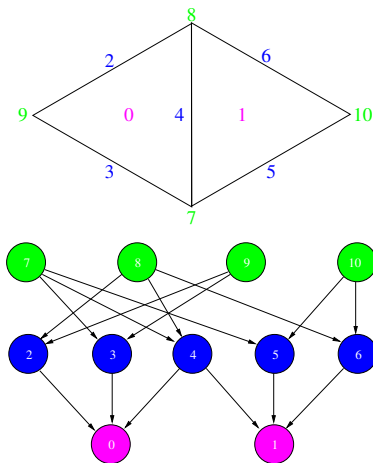
- Incidence/covering arrows
- $\text{closure}(0) = \{0, 2, 3, 4, 7, 8, 9\}$
- $\text{star}(7) = \{7, 2, 3, 0\}$

Doublet Mesh



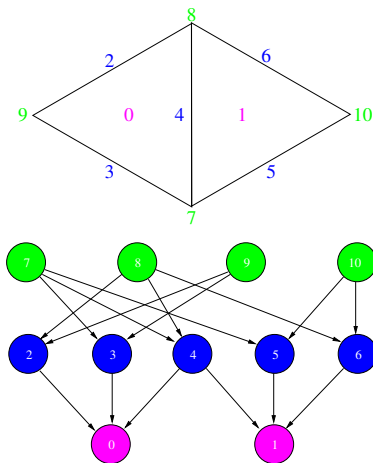
- Incidence/covering arrows
- $\text{closure}(0) = \{0, 2, 3, 4, 7, 8, 9\}$
- $\text{star}(7) = \{7, 2, 3, 0\}$

Doublet Mesh



- Incidence/covering arrows
- $meet(0, 1) = \{4\}$
- $join(8, 9) = \{4\}$

Doublet Mesh



- Incidence/covering arrows
- $meet(0, 1) = \{4\}$
- $join(8, 9) = \{4\}$

Sieve Definition

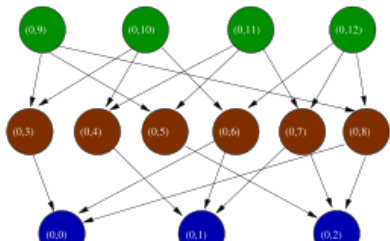
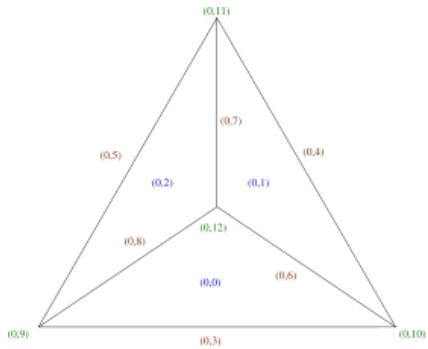
Definition

A Sieve consists of points, and arrows.

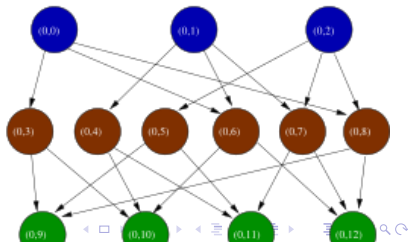
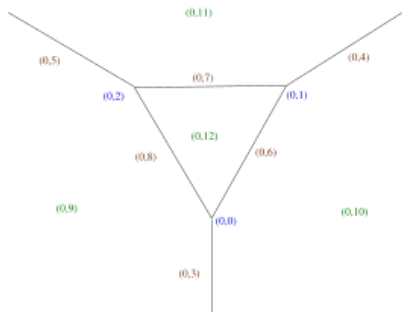
Each arrow connects a point to another which it covers.

$\text{cone}(p)$	sequence of points which cover a given point p
$\text{closure}(p)$	transitive closure of cone
$\text{support}(p)$	sequence of points which are covered by a given point p
$\text{star}(p)$	transitive closure of support
$\text{meet}(p,q)$	minimal separator of $\text{closure}(p)$ and $\text{closure}(q)$
$\text{join}(p,q)$	minimal separator of $\text{star}(p)$ and $\text{star}(q)$

The Mesh Dual



M. Knepley (UC)



FEM

KAUST

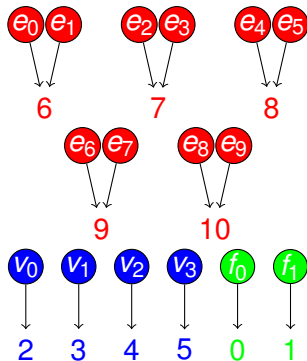
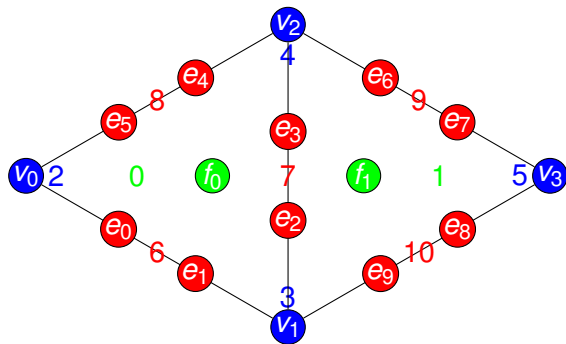
13 / 89

Outline

2 Mesh Distribution

- Sieve
- **Section**
- Completion
- Distribution
- Interfaces

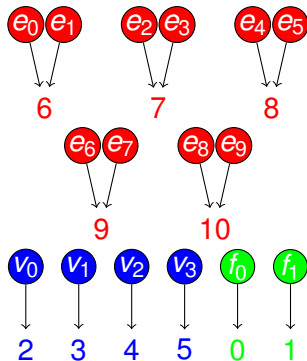
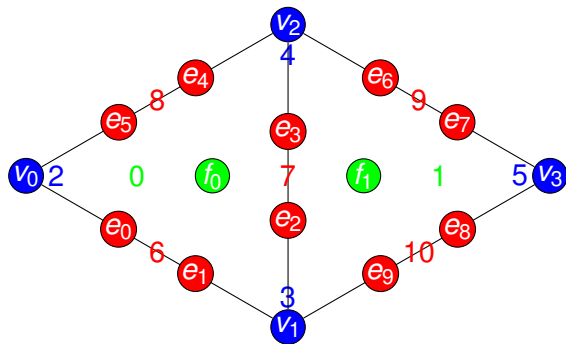
Doublet Section



Section interface

- $restrict(0) = \{f_0\}$
- $restrict(2) = \{v_0\}$
- $restrict(6) = \{e_0, e_1\}$

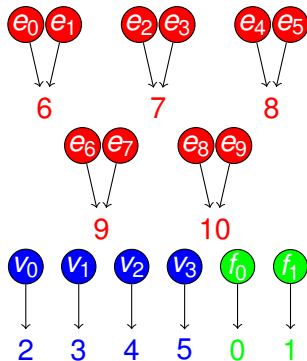
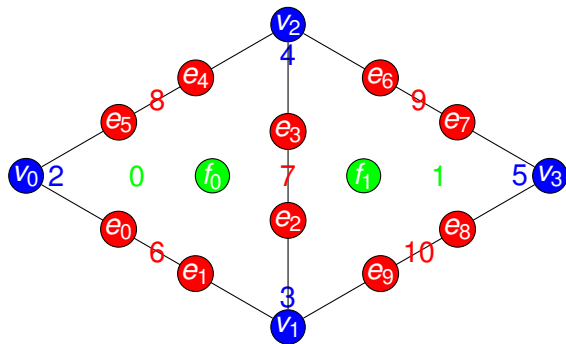
Doublet Section



• Section interface

- $restrict(0) = \{f_0\}$
- $restrict(2) = \{v_0\}$
- $restrict(6) = \{e_0, e_1\}$

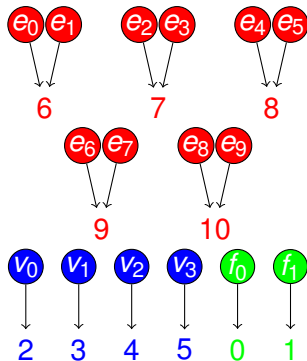
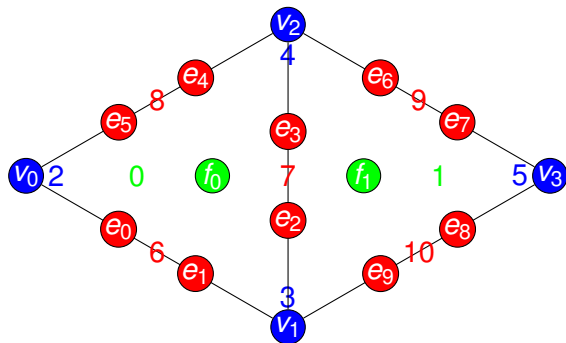
Doublet Section



• Section interface

- $restrict(0) = \{f_0\}$
- $restrict(2) = \{v_0\}$
- $restrict(6) = \{e_0, e_1\}$

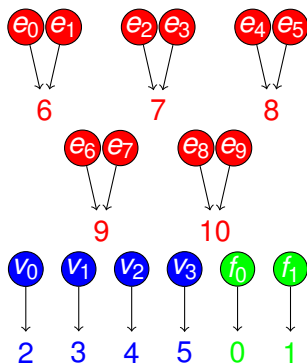
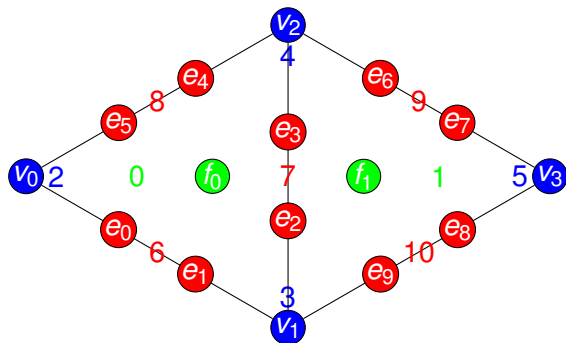
Doublet Section



• Section interface

- $restrict(0) = \{f_0\}$
- $restrict(2) = \{v_0\}$
- $restrict(6) = \{e_0, e_1\}$

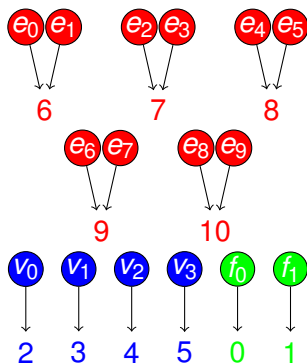
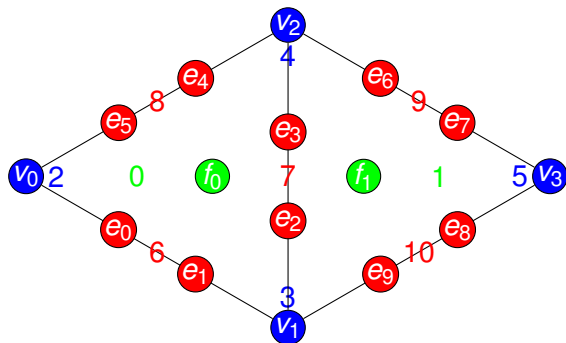
Doublet Section



- Topological traversals: follow connectivity

- $restrictClosure(0) = \{f_0, e_0, e_1, e_2, e_3, e_4, e_5, v_0, v_1, v_2\}$
- $restrictStar(7) = \{v_0, e_0, e_1, e_4, e_5, f_0\}$

Doublet Section

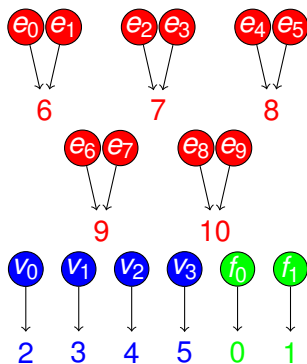
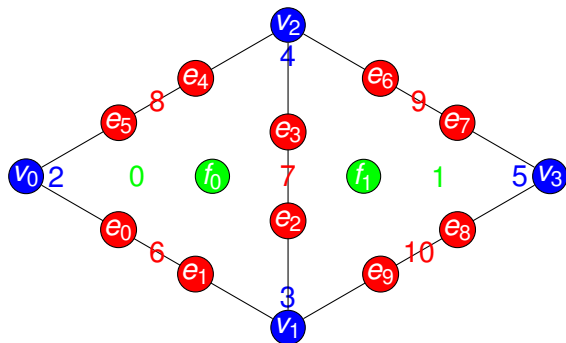


- Topological traversals: follow connectivity

- $restrictClosure(0) = \{f_0 e_0 e_1 e_2 e_3 e_4 e_5 V_0 V_1 V_2\}$

- $restrictStar(7) = \{V_0 e_0 e_1 e_4 e_5 f_0\}$

Doublet Section



- Topological traversals: follow connectivity

- $restrictClosure(0) = \{f_0 e_0 e_1 e_2 e_3 e_4 e_5 v_0 v_1 v_2\}$
- $restrictStar(7) = \{v_0 e_0 e_1 e_4 e_5 f_0\}$

Section Definition

Definition

Section is a mapping from sieve points to a vector of values.

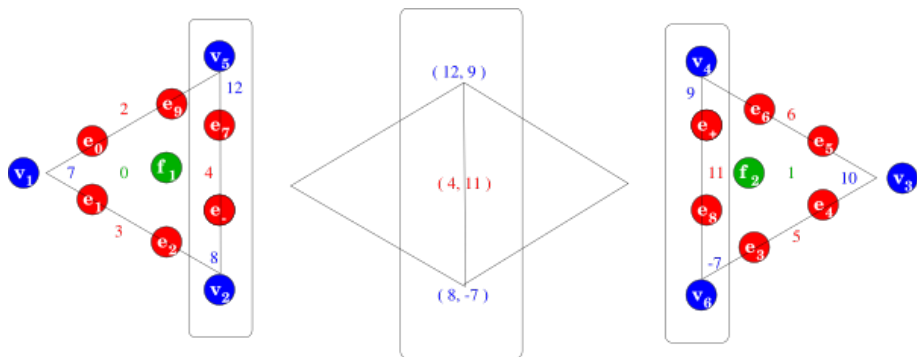
restrict	return all the values on given subdomain
update	inject subdomain values into global section
completion	operation to enforce coherence over sieve

Outline

2 Mesh Distribution

- Sieve
- Section
- **Completion**
- Distribution
- Interfaces

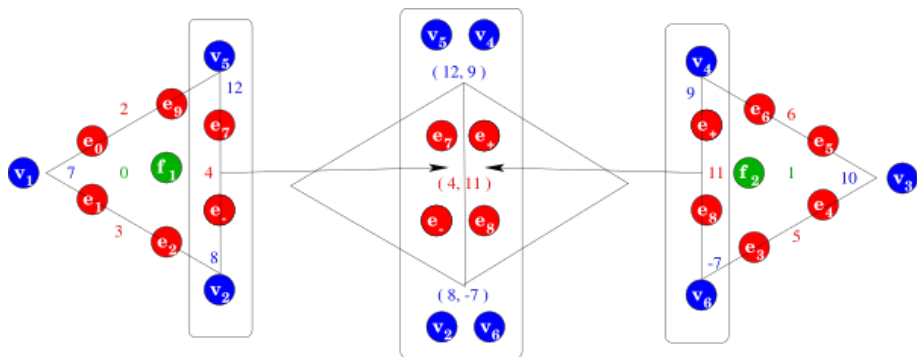
Restriction



- Localization

- Restrict to patches (here an edge closure)
- Compute locally

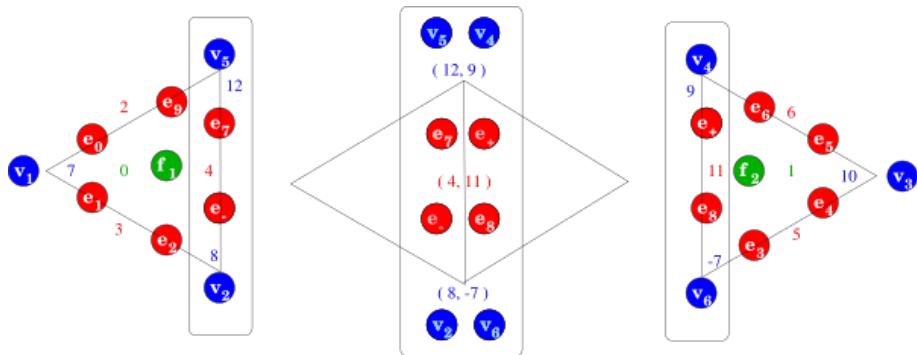
Delta



- Delta

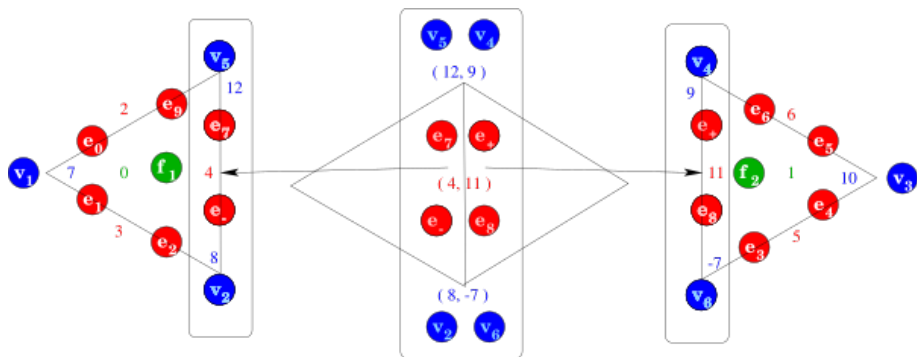
- Restrict further to the overlap
- Overlap now carries twice the data

Fusion



- Merge/reconcile data on the overlap
 - Addition (FEM)
 - Replacement (FD)
 - Coordinate transform (Sphere)
 - Linear transform (MG)

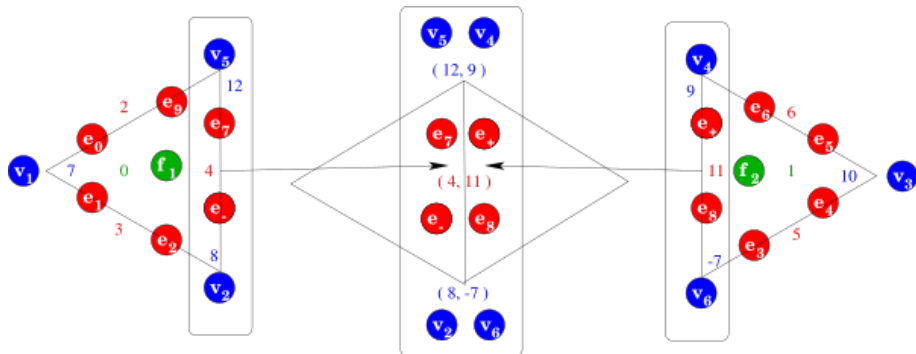
Update



- Update

- Update local patch data
- Completion = restrict \rightarrow fuse \rightarrow update, in parallel

Completion



- A ubiquitous parallel form of *restrict* \rightarrow *fuse* \rightarrow *update*
- Operates on Sections
 - Sieves can be "downcast" to Sections
- Based on two operations
 - Data exchange through overlap
 - Fusion of shared data

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Uses

Completion has many uses:

FEM accumulating integrals on shared faces

FVM accumulating fluxes on shared cells

FDM setting values on ghost vertices

- distributing mesh entities after partition
- redistributing mesh entities and data for load balance
- accumulating matvec for a partially assembled matrix

Section Completion

Completion can be broken into 4 phases:

- 1 `restrict()` to an overlap section
- 2 `copy()` data to the remote overlap section
- 3 `fuse()` data with existing point data
- 4 `update()` remote section with fused overlap section data

It is common to combine phases 1 & 2, and also 3 & 4

- Data is moved directly between communication buffers and storage

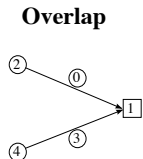
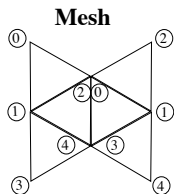
Section Completion

5
17
2
10
7

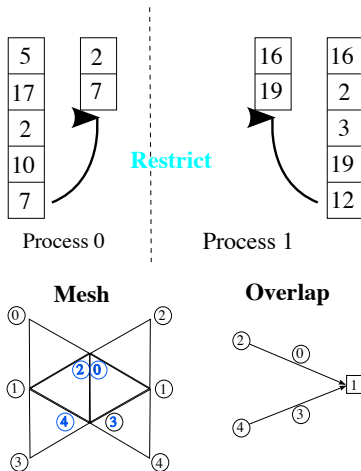
Process 0

16
2
3
19
12

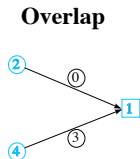
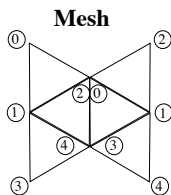
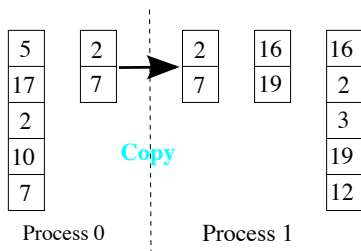
Process 1



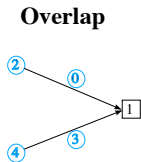
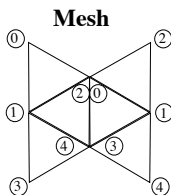
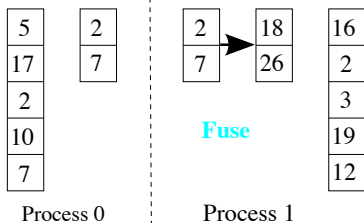
Section Completion



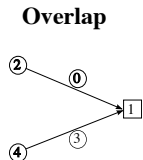
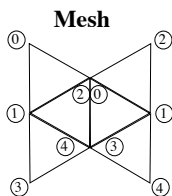
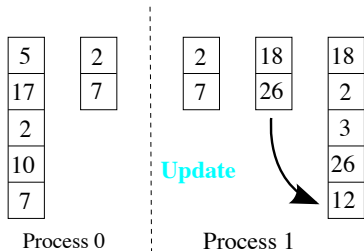
Section Completion



Section Completion



Section Completion



Section Hierarchy

We have a hierarchy of section types of increasing complexity

- `GeneralSection`
 - An arbitrary number of values for each domain point
 - Constrain arbitrary values
 - Atlas is a `UniformSection`
- `UniformSection`
 - A fixed number of values for each domain point
 - Atlas is a `ConstantSection`
- `ConstantSection`
 - The same single value for all domain points
 - Only the domain must be completed

Outline

2 Mesh Distribution

- Sieve
- Section
- Completion
- **Distribution**
- Interfaces

Section Distribution

Section distribution consists of

- Creation of the local Section
- Distribution of the Atlas (layout Section)
- Completion of the Section

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Sieve Distribution

- 1 Construct local mesh from partition
- 2 Construct partition overlap
- 3 Complete () the **partition section**
 - This distributes the cells
- 4 Update Overlap with new points
- 5 Complete () the **cone section**
 - This distributes the remaining sieve points
- 6 Update local Sieves

Mesh Distribution

Distributing a mesh means

- distributing the topology (Sieve)
- distributing data (Section)

However, a Sieve can be interpreted as a Section of `cone()` s!

Mesh Distribution

Distributing a mesh means

- distributing the topology (Sieve)
- distributing data (Section)

However, a Sieve can be interpreted as a Section of `cone()` s!

Mesh Distribution

Distributing a mesh means

- distributing the topology (Sieve)
- distributing data (Section)

However, a Sieve can be interpreted as a Section of `cone()` s!

Mesh Distribution

Distributing a mesh means

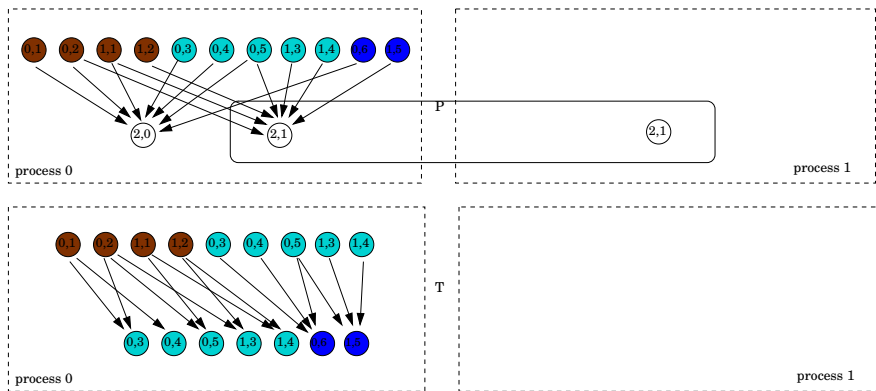
- distributing the topology (Sieve)
- distributing data (Section)

However, a Sieve can be interpreted as a Section of `cone()`s!

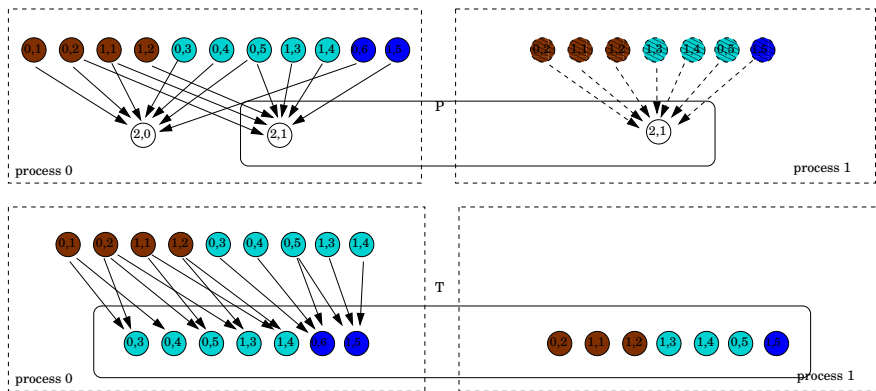
Mesh Partition

- 3rd party packages construct a vertex partition
- For FEM, partition dual graph vertices
- For FVM, construct hyperpgraph dual with faces as vertices
- Assign $\text{closure}(v)$ and $\text{star}(v)$ to same partition

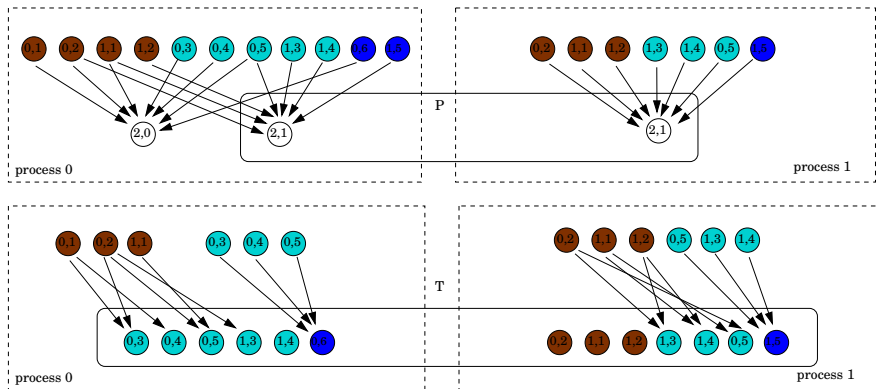
Doublet Mesh Distribution



Doublet Mesh Distribution

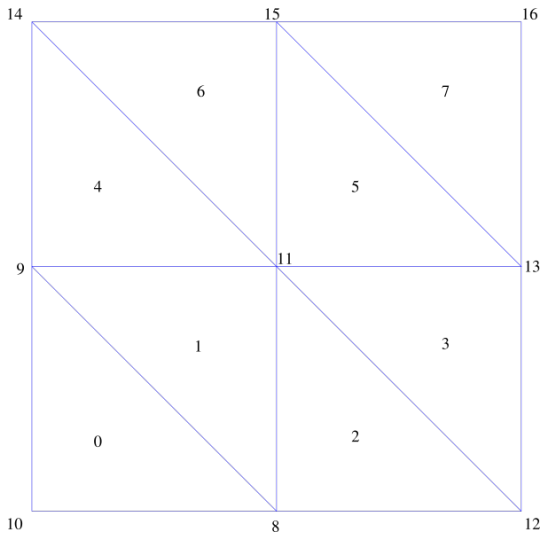


Doublet Mesh Distribution



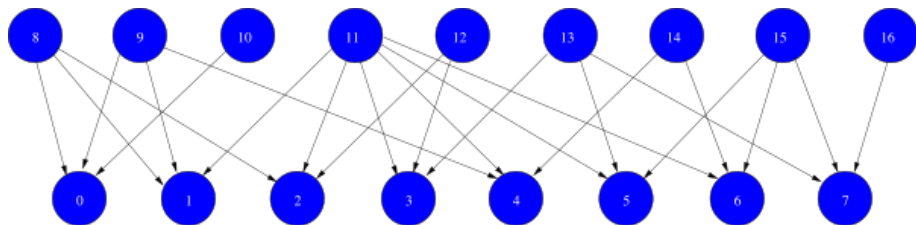
2D Example

A simple triangular mesh



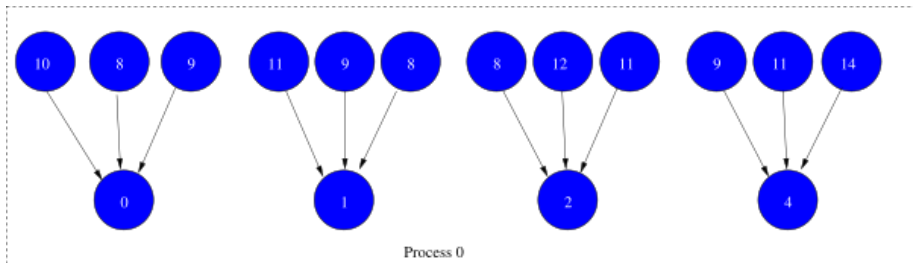
2D Example

Sieve for the mesh



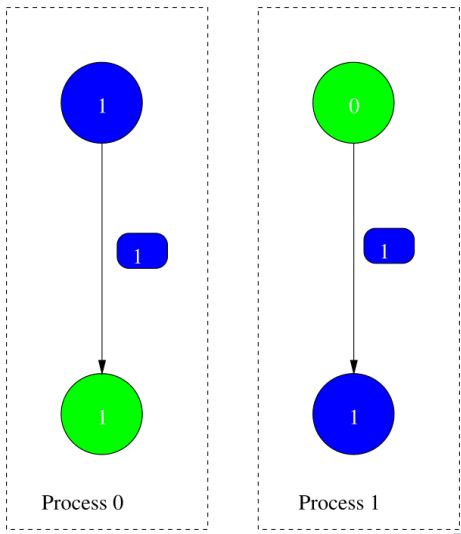
2D Example

Local sieve on process 0



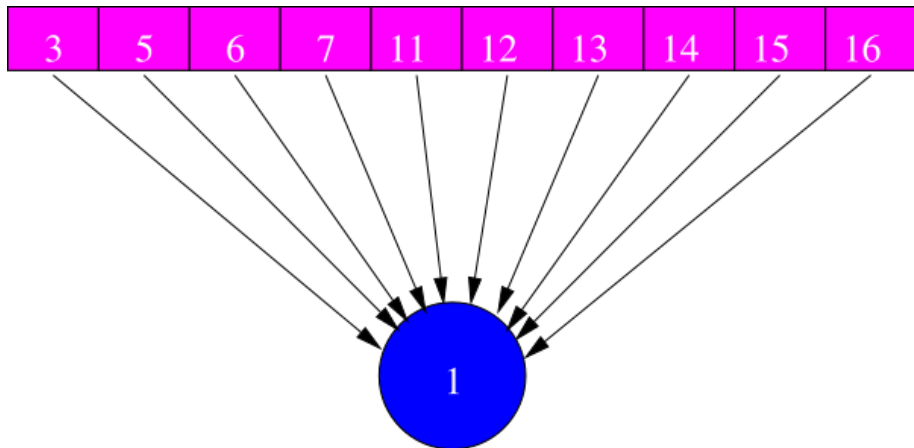
2D Example

Partition Overlap



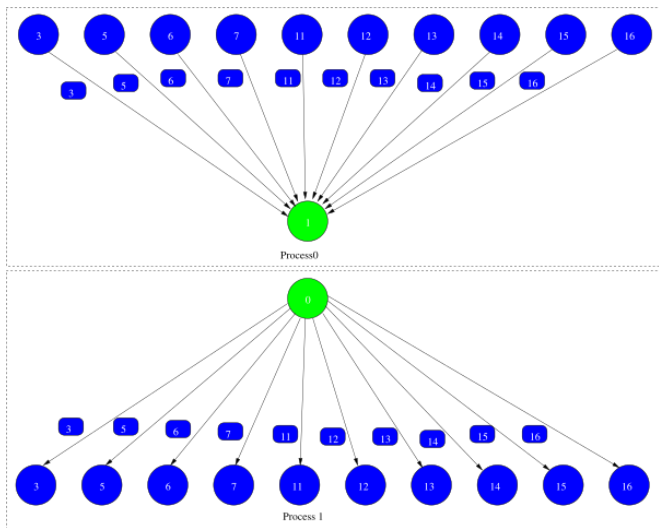
2D Example

Partition Section



2D Example

Updated Sieve Overlap



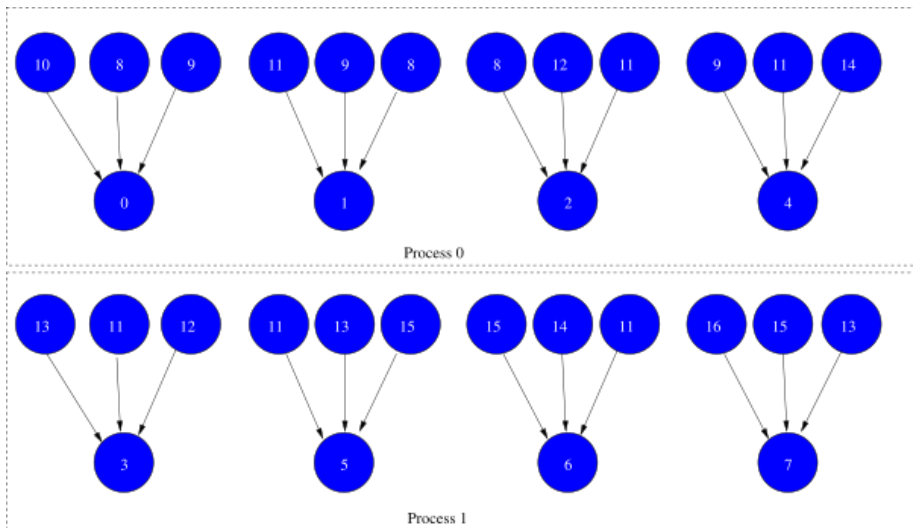
2D Example

Cone Section



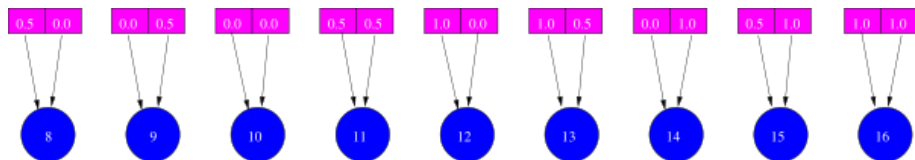
2D Example

Distributed Sieve



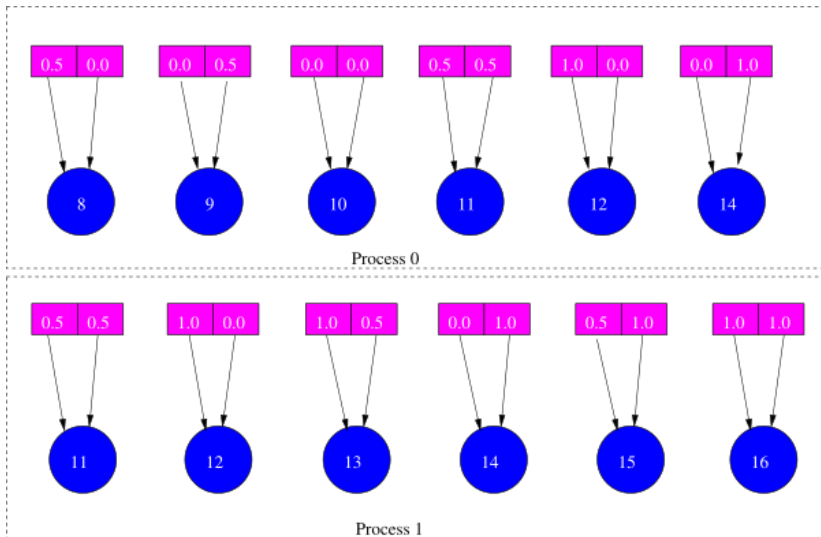
2D Example

Coordinate Section



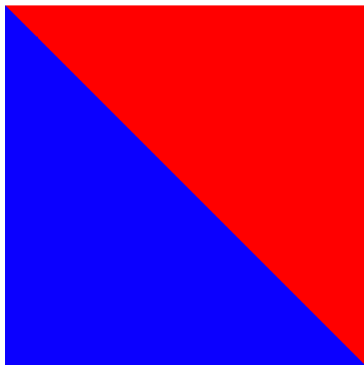
2D Example

Distributed Coordinate Section



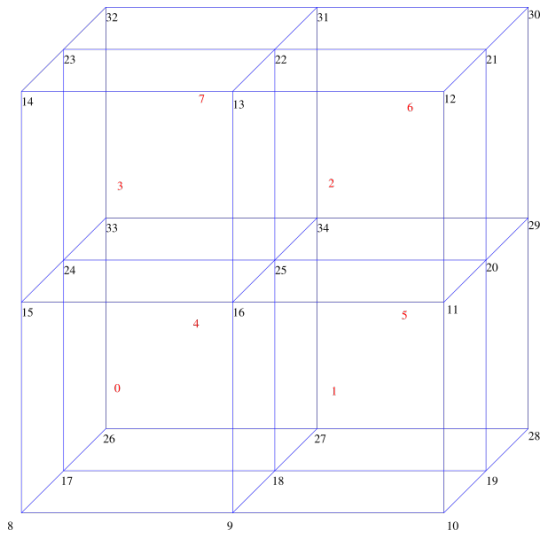
2D Example

Distributed Mesh



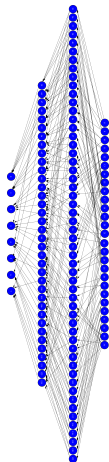
3D Example

A simple hexahedral mesh



3D Example

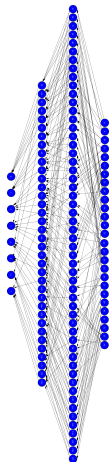
Sieve for the mesh



Its complicated!

3D Example

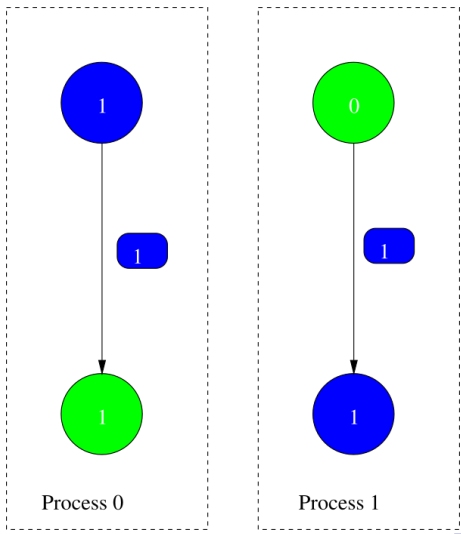
Sieve for the mesh



Its complicated!

3D Example

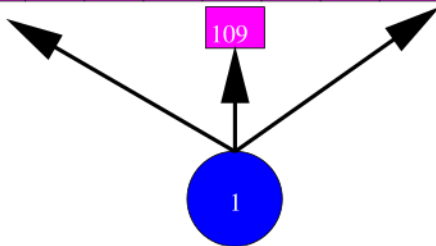
Partition Overlap



3D Example

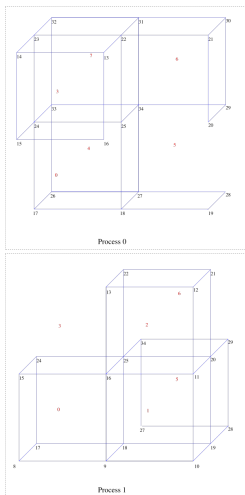
Partition Section

0	1	2	3	5	6	8	9	10	11	12	13	15
16	17	18	19	20	21	22	24	25	27	28	29	34
35	36	37	38	39	40	41	42	43	45	46	47	48
49	50	51	52	53	54	55	56	57	58	59	60	61
62	63	64	65	66	67	68	69	70	71	72	73	74
75	76	77	78	89	96	101	102	103	104	105	107	108



3D Example

Distributed Mesh



Notice cells are ghosted

Outline

2 Mesh Distribution

- Sieve
- Section
- Completion
- Distribution
- **Interfaces**

Sieve Overview

- Hierarchy is the centerpiece
 - Strip out unneeded complexity (dimension, shape, ...)
- Single relation, **covering**, handles all hierarchy
 - Rich enough for FEM
- Single operation, **completion**, for parallelism
 - Enforces consistency of the relation

Sieve Overview

- Hierarchy is the centerpiece
 - Strip out unneeded complexity (dimension, shape, ...)
- Single relation, **covering**, handles all hierarchy
 - Rich enough for FEM
- Single operation, **completion**, for parallelism
 - Enforces consistency of the relation

Sieve Overview

- Hierarchy is the centerpiece
 - Strip out unneeded complexity (dimension, shape, ...)
- Single relation, **covering**, handles all hierarchy
 - Rich enough for FEM
- Single operation, **completion**, for parallelism
 - Enforces consistency of the relation

Global and Local

Local (analytical)

- Discretization/Approximation
 - FEM integrals
 - FV fluxes
- Boundary conditions
- Largely dim dependent (e.g. quadrature)

Global (topological)

- Data management
 - Sections (local pieces)
 - Completions (assembly)
- Boundary definition
- Multiple meshes
 - Mesh hierarchies
- Largely dim independent (e.g. mesh traversal)

Global and Local

Local (analytical)

- Discretization/Approximation
 - FEM integrals
 - FV fluxes
- Boundary conditions
- Largely dim dependent (e.g. quadrature)

Global (topological)

- Data management
 - Sections (local pieces)
 - Completions (assembly)
- Boundary definition
- Multiple meshes
 - Mesh hierarchies
- Largely dim independent (e.g. mesh traversal)

Global and Local

Local (analytical)

- Discretization/Approximation
 - FEM integrals
 - FV fluxes
- Boundary conditions
- Largely dim dependent (e.g. quadrature)

Global (topological)

- Data management
 - Sections (local pieces)
 - Completions (assembly)
- Boundary definition
- Multiple meshes
 - Mesh hierarchies
- Largely dim independent (e.g. mesh traversal)

Global and Local

Local (analytical)

- Discretization/Approximation
 - FEM integrals
 - FV fluxes
- Boundary conditions
- Largely dim dependent (e.g. quadrature)

Global (topological)

- Data management
 - Sections (local pieces)
 - Completions (assembly)
- Boundary definition
- Multiple meshes
 - Mesh hierarchies
- Largely dim independent (e.g. mesh traversal)

Global and Local

Local (analytical)

- Discretization/Approximation
 - FEM integrals
 - FV fluxes
- Boundary conditions
- Largely dim dependent (e.g. quadrature)

Global (topological)

- Data management
 - Sections (local pieces)
 - Completions (assembly)
- Boundary definition
- Multiple meshes
 - Mesh hierarchies
- Largely dim independent (e.g. mesh traversal)

Hierarchical Interfaces

Global/Local Dichotomy is the **Heart** of DD
Software interfaces do not adequately reflect this

- PETSc DA is too specialized
 - Basically 1D methods applied to Cartesian products
- PETSc Index Sets and VecScatters are too fine
 - User “does everything”, no abstraction
- PETSc Linear Algebra (Vec & Mat) is too coarse
 - No access to the underlying connectivity structure

Unstructured Interface (before)

- Explicit references to element type
 - `getVertices(edgeID)`, `getVertices(faceID)`
 - `getAdjacency(edgeID, VERTEX)`
 - `getAdjacency(edgeID, dim = 0)`
- No interface for transitive closure
 - Awkward nested loops to handle different dimensions
- Have to recode for meshes with different
 - dimension
 - shapes

Unstructured Interface (before)

- Explicit references to element type
 - `getVertices(edgeID)`, `getVertices(faceID)`
 - `getAdjacency(edgeID, VERTEX)`
 - `getAdjacency(edgeID, dim = 0)`
- No interface for transitive closure
 - Awkward nested loops to handle different dimensions
- Have to recode for meshes with different
 - dimension
 - shapes

Unstructured Interface (before)

- Explicit references to element type
 - `getVertices(edgeID)`, `getVertices(faceID)`
 - `getAdjacency(edgeID, VERTEX)`
 - `getAdjacency(edgeID, dim = 0)`
- No interface for transitive closure
 - Awkward nested loops to handle different dimensions
- Have to recode for meshes with different
 - dimension
 - shapes

Go Back to the Math

Combinatorial Topology gives us a framework for geometric computing.

- Abstract to a relation, **covering**, on sieve points
 - Points can represent any mesh element
 - Covering can be thought of as adjacency
 - Relation can be expressed in a DAG (Hasse Diagram)
- Simple query set:
 - provides a general API for geometric algorithms
 - leads to simpler implementations
 - can be more easily optimized

Go Back to the Math

Combinatorial Topology gives us a framework for geometric computing.

- Abstract to a relation, **covering**, on sieve points
 - Points can represent any mesh element
 - Covering can be thought of as adjacency
 - Relation can be expressed in a DAG (Hasse Diagram)
- Simple query set:
 - provides a general API for geometric algorithms
 - leads to simpler implementations
 - can be more easily optimized

Go Back to the Math

Combinatorial Topology gives us a framework for geometric computing.

- Abstract to a relation, **covering**, on sieve points
 - Points can represent any mesh element
 - Covering can be thought of as adjacency
 - Relation can be expressed in a DAG (Hasse Diagram)
- Simple query set:
 - provides a general API for geometric algorithms
 - leads to simpler implementations
 - can be more easily optimized

Unstructured Interface (after)

- **NO** explicit references to element type
 - A point may be any mesh element
 - `getCone(point)`: adjacent $(d-1)$ -elements
 - `getSupport(point)`: adjacent $(d+1)$ -elements
- Transitive closure
 - `closure(cell)`: The computational unit for FEM
- Algorithms independent of mesh
 - dimension
 - shape (even hybrid)
 - global topology
 - data layout

Unstructured Interface (after)

- **NO** explicit references to element type
 - A point may be any mesh element
 - `getCone(point)`: adjacent $(d-1)$ -elements
 - `getSupport(point)`: adjacent $(d+1)$ -elements
- Transitive closure
 - `closure(cell)`: The computational unit for FEM
- Algorithms independent of mesh
 - dimension
 - shape (even hybrid)
 - global topology
 - data layout

Unstructured Interface (after)

- **NO** explicit references to element type
 - A point may be any mesh element
 - `getCone(point)`: adjacent $(d-1)$ -elements
 - `getSupport(point)`: adjacent $(d+1)$ -elements
- Transitive closure
 - `closure(cell)`: The computational unit for FEM
- Algorithms independent of mesh
 - dimension
 - shape (even hybrid)
 - global topology
 - data layout

Outline

- 1 Introduction
- 2 Mesh Distribution
- 3 Unifying Paradigm**
 - DA
 - Mesh
 - DMMG
 - PCFieldSplit
- 4 Finite Element Assembly

Hierarchy Abstractions

- Generalize to a set of linear spaces
 - `Sieve` provides topology, can also model `Mat`
 - `Section` generalizes `Vec`
 - Spaces interact through an `Overlap` (just a `Sieve`)
- Basic operations
 - Restriction to finer subspaces, `restrict()/update()`
 - Assembly to the subdomain, `complete()`
- Allow reuse of geometric and multilevel algorithms

Outline

3 Unifying Paradigm

- DA
- Mesh
- DMMG
- PCFieldSplit

Residual Evaluation

The **DM** interface is based upon *local* callback functions

- FormFunctionLocal()
- FormJacobianLocal()

Callbacks are registered using

- SNESSetDM(), TSSetDM()
- DMSNESSetFunctionLocal(), DMTSSetJacobianLocal()

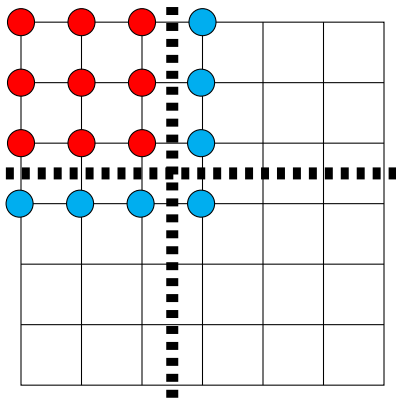
When PETSc needs to evaluate the nonlinear residual **F(x)**,

- Each process evaluates the local residual
- PETSc assembles the global residual automatically
 - Uses DMLocalToGlobal() method

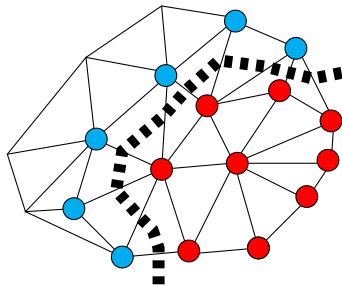
Ghost Values

To evaluate a local function $f(x)$, each process requires

- its local portion of the vector x
- its **ghost values**, bordering portions of x owned by neighboring processes



- Local Node
- Ghost Node



DMDA Local Function

User provided function calculates the nonlinear residual (in 2D)

```
(* If)(DMDALocalInfo *info, PetscScalar**x, PetscScalar**r, void *ctx)
```

`info`: All layout and numbering information

`x`: The current solution (a multidimensional array)

`r`: The residual

`ctx`: The user context passed to `DMDASNESSetFunctionLocal()`

The local DMDA function is activated by calling

```
DMDASNESSetFunctionLocal(dm, INSERT_VALUES, lfunc, &ctx)
```

Bratu Residual Evaluation

$$\Delta u + \lambda e^u = 0$$

```

ResLocal(DMDALocalInfo *info, PetscScalar **x, PetscScalar **f, void *ctx)
for(j = info->ys; j < info->ys+info->ym; ++j) {
  for(i = info->xs; i < info->xs+info->xm; ++i) {
    u = x[j][i];
    if (i==0 || j==0 || i == M || j == N) {
      f[j][i] = 2.0*(hydhx+hxhdy)*u; continue;
    }
    u_xx    = (2.0*u - x[j][i-1] - x[j][i+1])*hydhx;
    u_yy    = (2.0*u - x[j-1][i] - x[j+1][i])*hxhdy;
    f[j][i] = u_xx + u_yy - hx*hy*lambda*exp(u);
  }}

```

[\\$PETSC_DIR/src/snes/examples/tutorials/ex5.c](#)

DMDA Local Jacobian

User provided function calculates the Jacobian (in 2D)

```
(* ljac )(DMDALocalInfo *info, PetscScalar**x, Mat J, void *ctx)
```

`info`: All layout and numbering information

`x`: The current solution

`J`: The Jacobian

`ctx`: The user context passed to `DASetLocalJacobian()`

The local DMDA function is activated by calling

```
DMDASNESSetJacobianLocal(dm, ljac, &ctx)
```

DMDA Vectors

- The **DMDA** object contains only layout (topology) information
 - All field data is contained in PETSc **Vecs**
- Global vectors are parallel
 - Each process stores a unique local portion
 - `DMCreateGlobalVector(DM da, Vec *gvec)`
- Local vectors are sequential (and usually temporary)
 - Each process stores its local portion plus ghost values
 - `DMCreateLocalVector(DM da, Vec *lvec)`
 - includes ghost and boundary values!

Updating Ghosts

Two-step process enables overlapping computation and communication

- `DMGlobalToLocalBegin(da, gvec, mode, lvec)`
 - `gvec` provides the data
 - `mode` is either `INSERT_VALUES` or `ADD_VALUES`
 - `lvec` holds the local and ghost values
- `DMGlobalToLocalEnd(da, gvec, mode, lvec)`
 - Finishes the communication

The process can be reversed with `DALocalToGlobalBegin/End()`.

Outline

3 Unifying Paradigm

- DA
- **Mesh**
- DMMG
- PCFieldSplit

Mesh Paradigm

The **DMMesh** interface also uses *local* callback functions

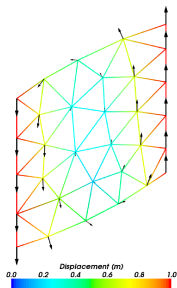
- maps between **global** Vec and **local** Vec
- Local vectors are structured using a **PetscSection**

When PETSc needs to evaluate the nonlinear residual $F(x)$,

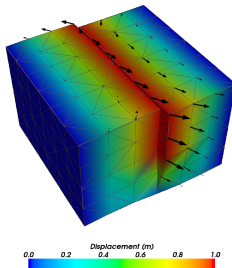
- Each process evaluates the local residual for each element
- PETSc assembles the global residual automatically
 - `DMLocalToGlobal()` works just as in the structured case

Multiple Mesh Types

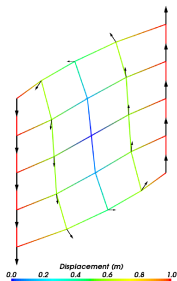
Triangular



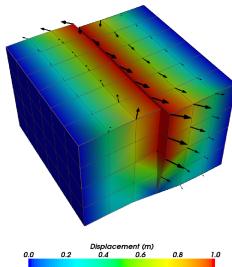
Tetrahedral



Rectangular



Hexahedral



Outline

3 Unifying Paradigm

- DA
- Mesh
- **DMMG**
- PCFieldSplit

Multigrid Paradigm

The **DM** interface uses the *local* callback functions to

- assemble global functions/operators from local pieces
- assemble functions/operators on coarse grids

Then **PCMG** organizes

- control flow for the multilevel solve, and
- projection and smoothing operators at each level.

DM Integration with SNES

- DM supplies global residual and Jacobian to SNES
 - User supplies local version to DM
 - The `Rhs_*()` and `Jac_*()` functions in the example
- Allows automatic parallelism
- Allows grid hierarchy
 - Enables multigrid once interpolation/restriction is defined
- Paradigm is developed in unstructured work
 - Solve needs scatter into contiguous global vectors (initial guess)
- Handle Neumann BC using `KSPSetNullSpace()`

Multigrid with DM

Allows multigrid with some simple command line options

- `-pc_type mg, -pc_mg_levels`
- `-pc_mg_type, -pc_mg_cycle_type, -pc_mg_galerkin`
- `-mg_levels_1_ksp_type, -mg_levels_1_pc_type`
- `-mg_coarse_ksp_type, -mg_coarse_pc_type`
- `-da_refine, -ksp_view`

Interface also works with GAMG and 3rd party packages like ML

Outline

3 Unifying Paradigm

- DA
- Mesh
- DMMG
- **PCFieldSplit**

MultiPhysics Paradigm

The **PCFieldSplit** interface

- extracts functions/operators corresponding to each physics
 - **VecScatter** and `MatGetSubMatrix()` for efficiency
- assemble functions/operators over all physics
 - Generalizes `LocalToGlobal()` mapping
- is composable with **ANY** PETSc solver and preconditioner
 - This can be done recursively

MultiPhysics Paradigm

The **PCFieldSplit** interface

- extracts functions/operators corresponding to each physics
 - **VecScatter** and `MatGetSubMatrix()` for efficiency
- assemble functions/operators over all physics
 - Generalizes `LocalToGlobal()` mapping
- is composable with **ANY** PETSc solver and preconditioner
 - This can be done recursively

FieldSplit provides the **building blocks** for multiphysics preconditioning.

MultiPhysics Paradigm

The **PCFieldSplit** interface

- extracts functions/operators corresponding to each physics
 - **VecScatter** and `MatGetSubMatrix()` for efficiency
- assemble functions/operators over all physics
 - Generalizes `LocalToGlobal()` mapping
- is composable with **ANY** PETSc solver and preconditioner
 - This can be done recursively

Notice that this works in exactly the same manner as

- multiple resolutions (MG, FMM, Wavelets)
- multiple domains (Domain Decomposition)
- multiple dimensions (ADI)

Preconditioning

Several varieties of preconditioners can be supported:

- Block Jacobi or Block Gauss-Siedel
- Schur complement
- Block ILU (approximate coupling and Schur complement)
- Dave May's implementation of Elman-Wathen type PCs

which only require actions of individual operator blocks

Notice also that we may have any combination of

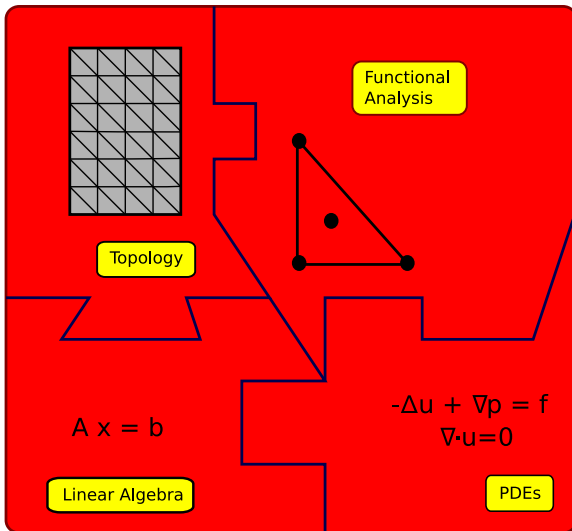
- “canned” PCs (ILU, AMG)
- PCs needing special information (MG, FMM)
- custom PCs (physics-based preconditioning, Born approximation)

since we have access to an algebraic interface

Outline

- 1 Introduction
- 2 Mesh Distribution
- 3 Unifying Paradigm
- 4 Finite Element Assembly**
 - Layout
 - Integration
 - Assembly
 - Examples

Mathematics Puzzle



FEM Components

- Section definition
- Integration
- Assembly and Boundary conditions

Outline

4 Finite Element Assembly

- **Layout**
- Integration
- Assembly
- Examples

Section Allocation

We only need the
 fiber dimension (# of unknowns)
of each
 sieve point (piece of the mesh)

- Determined by discretization
- By symmetry, only depend on point depth
- Obtained from FIAT
- Modified by BC
- Decouples storage and parallelism from discretization

Section Allocation

We only need the
 fiber dimension (# of unknowns)
of each
 sieve point (piece of the mesh)

- Determined by discretization
- By symmetry, only depend on point depth
- Obtained from FIAT
- Modified by BC
- Decouples storage and parallelism from discretization

Section Allocation

We only need the
 fiber dimension (# of unknowns)
of each
 sieve point (piece of the mesh)

- Determined by discretization
- By symmetry, only depend on point depth
- Obtained from FIAT
- Modified by BC
- Decouples storage and parallelism from discretization

Section Allocation

We only need the
 fiber dimension (# of unknowns)
of each
 sieve point (piece of the mesh)

- Determined by discretization
- By symmetry, only depend on point depth
- Obtained from FIAT
- Modified by BC
- Decouples storage and parallelism from discretization

Section Allocation

We only need the
 fiber dimension (# of unknowns)
of each
 sieve point (piece of the mesh)

- Determined by discretization
- By symmetry, only depend on point depth
- Obtained from FIAT
- Modified by BC
- Decouples storage and parallelism from discretization

Outline

4 Finite Element Assembly

- Layout
- **Integration**
- Assembly
- Examples

FIAT

Finite Element Integrator And Tabulator by Rob Kirby

<http://fenicsproject.org/>

FIAT understands

- Reference element shapes (line, triangle, tetrahedron)
- Quadrature rules
- Polynomial spaces
- Functionals over polynomials (dual spaces)
- Derivatives

Can build arbitrary elements by specifying the Ciarlet triple (K, P, P')

FIAT is part of the FEniCS project

FIAT

Finite Element Integrator And Tabulator by Rob Kirby

<http://fenicsproject.org/>

FIAT understands

- Reference element shapes (line, triangle, tetrahedron)
- Quadrature rules
- Polynomial spaces
- Functionals over polynomials (dual spaces)
- Derivatives

Can build arbitrary elements by specifying the Ciarlet triple (K, P, P')

FIAT is part of the FEniCS project

FIAT Integration

The `quadrature.fiat` file contains:

- An element (usually a family and degree) defined by FIAT
- A quadrature rule

It is run

- automatically by `make`, or
- independently by the user

It can take arguments

- `-element_family` and `-element_order`, or
- `make` takes variables `ELEMENT` and `ORDER`

Then `make` produces `quadrature.h` with:

- Quadrature points and weights
- Basis function and derivative evaluations at the quadrature points
- Integration against dual basis functions over the cell
- Local dofs for Section allocation

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

Kinds of Unknowns

We must map local unknowns to the global basis

- FIAT reports the kind of unknown
- Scalars are invariant
 - Lagrange
- Vectors transform as J^{-T}
 - Hermite
- Normal vectors require Piola transform and a choice of orientation
 - Raviart-Thomas
- Moments transform as $|J^{-1}|$
 - Nedelec
- May involve a transformation over the entire closure
 - Argyris
- Conjecture by Kirby relates transformation to affine equivalence
- We have not yet automated this step (FFC, Mython)

FFC

FFC is a compiler for variational forms by Anders Logg.

Here is a mixed-form Poisson equation:

$$a((\tau, w), (\sigma, u)) = L((\tau, w)) \quad \forall (\tau, w) \in V$$

where

$$\begin{aligned} a((\tau, w), (\sigma, u)) &= \int_{\Omega} \tau \sigma - \nabla \cdot \tau u + w \nabla \cdot u \, dx \\ L((\tau, w)) &= \int_{\Omega} w f \, dx \end{aligned}$$

FFC

Mixed Poisson

```
shape = "triangle"
```

```
BDM1 = FiniteElement("Brezzi–Douglas–Marini", shape, 1)
```

```
DG0 = FiniteElement("Discontinuous Lagrange", shape, 0)
```

```
element = BDM1 + DG0
```

```
(tau, w) = TestFunctions(element)
```

```
(sigma, u) = TrialFunctions(element)
```

```
a = (dot(tau, sigma) - div(tau)*u + w*div(sigma))*dx
```

```
f = Function(DG0)
```

```
L = w*f*dx
```

FFC

Here is a discontinuous Galerkin formulation of the Poisson equation:

$$a(v, u) = L(v) \quad \forall v \in V$$

where

$$\begin{aligned} a(v, u) &= \int_{\Omega} \nabla u \cdot \nabla v \, dx \\ &+ \sum_S \int_S - \langle \nabla v \rangle \cdot [[u]]_n - [[v]]_n \cdot \langle \nabla u \rangle - (\alpha/h)vu \, dS \\ &+ \int_{\partial\Omega} -\nabla v \cdot [[u]]_n - [[v]]_n \cdot \nabla u - (\gamma/h)vu \, ds \\ L(v) &= \int_{\Omega} vf \, dx \end{aligned}$$

FFC

DG Poisson

```
DG1 = FiniteElement("Discontinuous Lagrange", shape, 1)
v = TestFunctions(DG1)
u = TrialFunctions(DG1)
f = Function(DG1)
g = Function(DG1)
n = FacetNormal("triangle")
h = MeshSize("triangle")
a = dot(grad(v), grad(u))*dx
  - dot(avg(grad(v)), jump(u, n))*dS
  - dot(jump(v, n), avg(grad(u)))*dS
  + alpha/h*dot(jump(v, n) + jump(u, n))*dS
  - dot(grad(v), jump(u, n))*ds
  - dot(jump(v, n), grad(u))*ds
  + gamma/h*v*u*ds
L = v*f*dx + v*g*ds
```

Outline

4 Finite Element Assembly

- Layout
- Integration
- **Assembly**
- Examples

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    coords = mesh->restrict(coordinates, c);
    v0, J, invJ, detJ = computeGeometry(coords);
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
```


Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
  <Compute cell geometry>
  inputVec = mesh->restrict(U, c);
  for(q = 0; q < numQuadPoints; ++q) {
    <Transform coordinates>
    for(f = 0; f < numBasisFuncs; ++f) {
      <Constant term>
      <Linear term>
      <Nonlinear term>
      elemVec[f] *= weight[q]*detJ;
    }
  }
  <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        realCoords = J*refCoords[q] + v0;
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
  <Compute cell geometry>
  <Retrieve values from input vector>
  for(q = 0; q < numQuadPoints; ++q) {
    <Transform coordinates>
    for(f = 0; f < numBasisFuncs; ++f) {
      elemVec[f] += basis[q,f]*rhsFunc(realCoords);
      <Linear term>
      <Nonlinear term>
      elemVec[f] *= weight[q]*detJ;
    }
  }
  <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            for(d = 0; d < dim; ++d)
                for(e) testDerReal[d] += invJ[e,d]*basisDer[q,
            for(g = 0; g < numBasisFuncs; ++g) {
                for(d = 0; d < dim; ++d)
                    for(e) basisDerReal[d] += invJ[e,d]*basisDer
                elemMat[f,g] += testDerReal[d]*basisDerReal[
                elemVec[f] += elemMat[f,g]*inputVec[g];
            }
        }
    }
}
```


Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
  <Compute cell geometry>
  <Retrieve values from input vector>
  for(q = 0; q < numQuadPoints; ++q) {
    <Transform coordinates>
    for(f = 0; f < numBasisFuncs; ++f) {
      <Constant term>
      <Linear term>
      <Nonlinear term>
      elemVec[f] *= weight[q]*detJ;
    }
  }
  <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            elemVec[f] += basis[q,f]*lambda*exp(inputVec[f])
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
  <Compute cell geometry>
  <Retrieve values from input vector>
  for(q = 0; q < numQuadPoints; ++q) {
    <Transform coordinates>
    for(f = 0; f < numBasisFuncs; ++f) {
      <Constant term>
      <Linear term>
      <Nonlinear term>
      elemVec[f] *= weight[q]*detJ;
    }
  }
  mesh->updateAdd(F, c, elemVec);
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
<Aggregate updates>
```

Integration

```
cells = mesh->heightStratum(0);
for(c = cells->begin(); c != cells->end(); ++c) {
    <Compute cell geometry>
    <Retrieve values from input vector>
    for(q = 0; q < numQuadPoints; ++q) {
        <Transform coordinates>
        for(f = 0; f < numBasisFuncs; ++f) {
            <Constant term>
            <Linear term>
            <Nonlinear term>
            elemVec[f] *= weight[q]*detJ;
        }
    }
    <Update output vector>
}
Distribution<Mesh>::completeSection(mesh, F);
```

Boundary Conditions

Dirichlet conditions may be expressed as

$$u|_{\Gamma} = g$$

and implemented by constraints on dofs in a Section

- The user provides a function.

Neumann conditions may be expressed as

$$\nabla u \cdot \hat{n}|_{\Gamma} = h$$

and implemented by explicit integration along the boundary

- The user provides a weak form.

Dirichlet Values

- Topological boundary is marked during generation
- Cells bordering boundary are marked using `markBoundaryCells()`
- To set values:
 - 1 Loop over boundary cells
 - 2 Loop over the element closure
 - 3 For each boundary point i , apply the functional N_i to the function g
- The functionals are generated with the quadrature information
- Section allocation applies Dirichlet conditions automatically
 - Values are stored in the Section
 - `restrict()` behaves normally, `update()` ignores constraints

Dual Basis Application

We would like the action of a dual basis vector (functional)

$$\langle \mathcal{N}_i, f \rangle = \int_{\text{ref}} N_i(x) f(x) dV$$

- Projection onto \mathcal{P}
- Code is generated from FIAT specification
 - Python code generation package inside PETSc
- Common interface for all elements

Assembly with Dirichlet Conditions

The original equation may be partitioned into

- unknowns in the interior (I)
- unknowns on the boundary (Γ)

so that we obtain

$$\begin{pmatrix} A_{II} & A_{I\Gamma} \\ A_{\Gamma I} & A_{\Gamma\Gamma} \end{pmatrix} \begin{pmatrix} u_I \\ u_\Gamma \end{pmatrix} = \begin{pmatrix} f_I \\ f_\Gamma \end{pmatrix}$$

However u_Γ is known, so we may reduce this to

$$A_{II}u_I = f_I - A_{I\Gamma}u_\Gamma$$

We will show that our scheme automatically constructs this extra term.

Assembly with Dirichlet Conditions

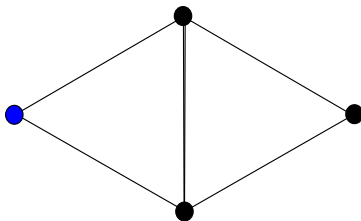
Residual Assembly

u

5	1	3	7
----------	----------	----------	----------

f

5	0	0	0
----------	----------	----------	----------



Assembly with Dirichlet Conditions

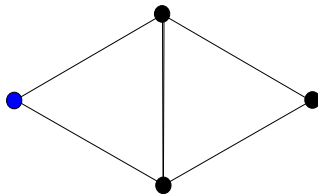
Residual Assembly

$$\mathbf{u} \begin{array}{|c|c|c|c|} \hline 5 & 1 & 3 & 7 \\ \hline \end{array}$$

$$\mathbf{f} \begin{array}{|c|c|c|c|} \hline 5 & 0 & 0 & 0 \\ \hline \end{array}$$

Restrict

$$\begin{array}{|c|} \hline 5 \\ \hline 1 \\ \hline 3 \\ \hline \end{array}$$

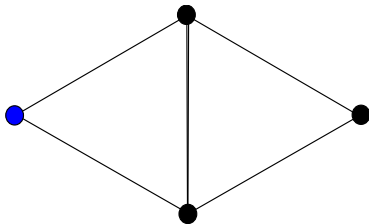


Assembly with Dirichlet Conditions

Residual Assembly

$$\mathbf{u} \quad \begin{array}{|c|c|c|c|} \hline \mathbf{5} & \mathbf{1} & \mathbf{3} & \mathbf{7} \\ \hline \end{array}$$

$$\mathbf{f} \quad \begin{array}{|c|c|c|c|} \hline \mathbf{5} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \hline \end{array}$$



Compute

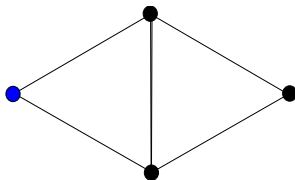
$$\begin{array}{|c|c|c|} \hline \mathbf{0.5} & \mathbf{0.0} & \mathbf{-0.5} \\ \hline \mathbf{0.0} & \mathbf{0.5} & \mathbf{-0.5} \\ \hline \mathbf{-0.5} & \mathbf{-0.5} & \mathbf{1.0} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \mathbf{5} \\ \hline \mathbf{1} \\ \hline \mathbf{3} \\ \hline \end{array} = \begin{array}{|c|} \hline \mathbf{1} \\ \hline \mathbf{-1} \\ \hline \mathbf{0} \\ \hline \end{array}$$

Assembly with Dirichlet Conditions

Residual Assembly

$$\mathbf{u} \quad \begin{bmatrix} 5 & 1 & 3 & 7 \end{bmatrix}$$

$$\mathbf{f} \quad \begin{bmatrix} 5 & 0 & 0 & 0 \end{bmatrix}$$



Compute

$$\begin{bmatrix} \mathbf{A}_{\text{rr}} & \mathbf{A}_{\text{rI}} \\ \mathbf{A}_{\text{Ir}} & \mathbf{A}_{\text{II}} \end{bmatrix} \begin{bmatrix} 5 \\ 1 \\ 3 \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \left. \vphantom{\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}} \right\} \text{This piece contains rhs}$$

Assembly with Dirichlet Conditions

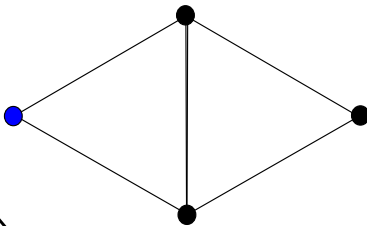
Residual Assembly

\mathbf{u}

5	1	3	7
----------	---	---	---

\mathbf{f}

5	-1	0	0
----------	----	---	---



Update



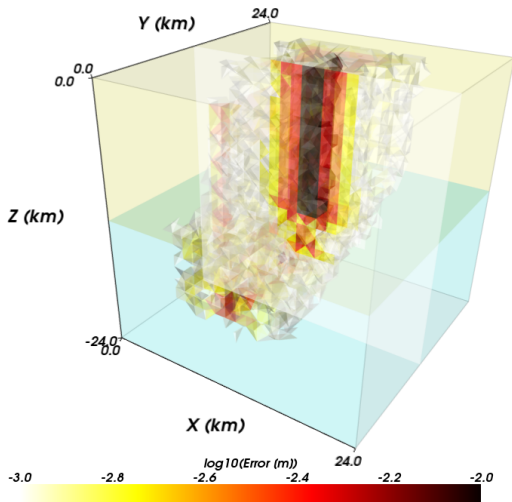
Outline

4 Finite Element Assembly

- Layout
- Integration
- Assembly
- **Examples**

PyLith

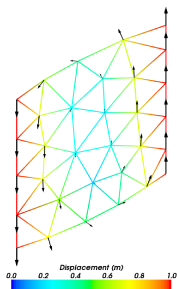
- Multiple problems
 - Dynamic rupture
 - Quasi-static relaxation
- Multiple models
 - Nonlinear visco-plastic
 - Finite deformation
 - Fault constitutive models
- Multiple meshes
 - 1D, 2D, 3D
 - Hex and tet meshes
- Parallel
 - PETSc solvers
 - DMPlex mesh management



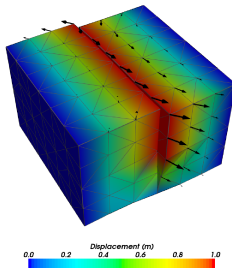
^aAagaard, Knepley, Williams

Multiple Mesh Types

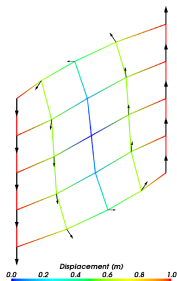
Triangular



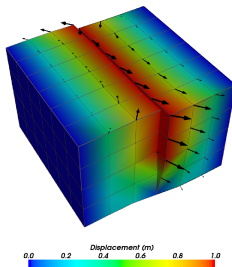
Tetrahedral



Rectangular

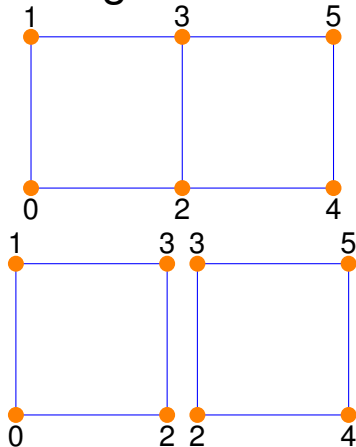


Hexahedral

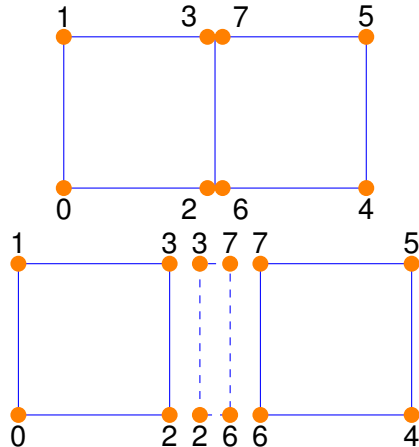


Cohesive Cells

Original Mesh



Mesh with Cohesive Cell



Exploded view of meshes



Cohesive Cells

Cohesive cells are used to enforce slip conditions on a fault

- Demand complex mesh manipulation
 - We allow specification of only fault vertices
 - Must “sew” together on output
- Use Lagrange multipliers to enforce constraints
 - Forces illuminate physics
- Allow different fault constitutive models
 - Simplest is enforced slip
 - Now have fault constitutive models

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

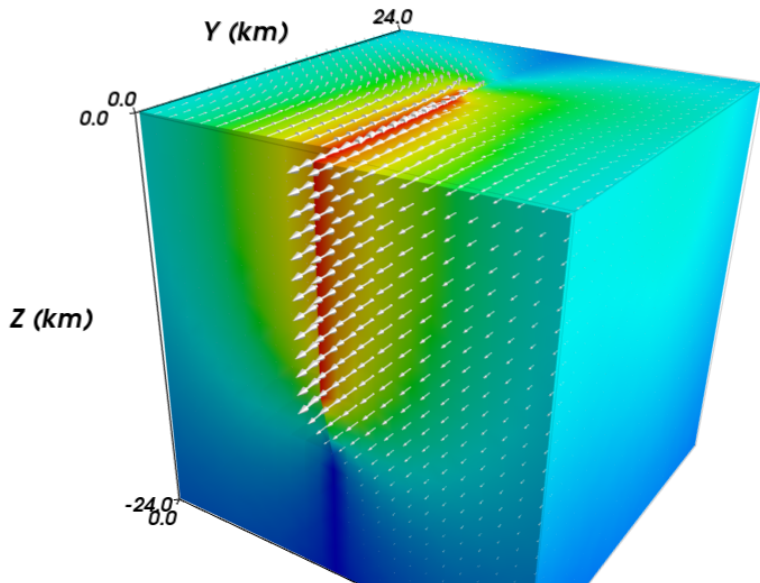
Splitting the Mesh

- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Splitting the Mesh

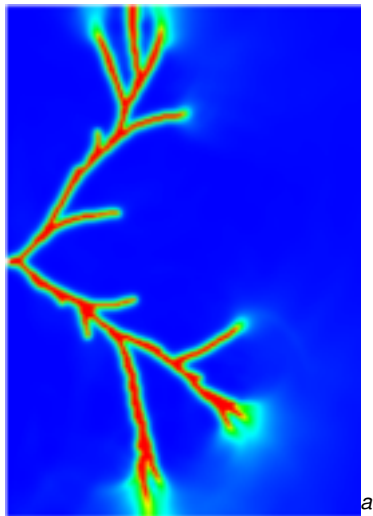
- In order to create a fault, the generator provides
 - a set of fault vertices, or
 - a set of fault faces.
- Fault vertices, unlike fault faces, must be
 - combined into faces on a fault mesh, and
 - oriented
- The fault mesh is used to
 - split vertices along the fault
 - introduce prism elements between adjacent fault faces
- Sieve code works for
 - any dimension
 - any element shape

Reverse-slip Benchmark



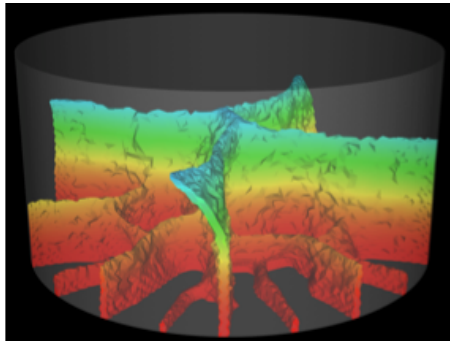
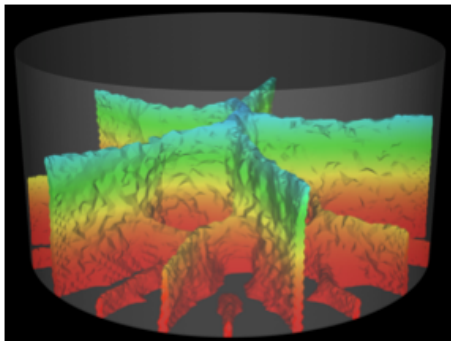
Fracture Mechanics

- Full variational formulation
 - Phase field
 - Linear or Quadratic penalty
- Uses TAO optimization
 - Necessary for linear penalty
 - Backtacking
- No prescribed cracks (**movie**)
 - Arbitrary crack geometry
 - Arbitrary intersections
- Multiple materials
 - Composite toughness



^aBourdin

Fracture Mechanics



¹Bourdin

Conclusions

Better mathematical abstractions bring concrete benefits

- Vast reduction in complexity
 - Dimension and mesh independent code
 - Complete serial code reuse
- Opportunities for optimization
 - Higher level operations missed by traditional compilers
 - Single communication routine to optimize
- Expansion of capabilities
 - Arbitrary elements
 - Unstructured multigrid
 - Multilevel algorithms

References

- **FEniCS Documentation:**

http://www.fenics.org/wiki/FEniCS_Project

- Project documentation
- Users manuals
- Repositories, bug tracking
- Image gallery

- **Publications:**

http://www.fenics.org/wiki/Related_presentations_and_publications

- Research and publications that make use of FEniCS

- **PETSc Documentation:**

<http://www.mcs.anl.gov/petsc/docs>

- PETSc Users manual
- Manual pages
- Many hyperlinked examples
- FAQ, Troubleshooting info, installation info, etc.
- Publication using PETSc

Experimentation is Essential!

Proof is not currently enough to examine solvers

- N. M. Nachtigal, S. C. Reddy, and L. N. Trefethen, *How fast are nonsymmetric matrix iterations?*, SIAM J. Matrix Anal. Appl., **13**, pp.778–795, 1992.
- Anne Greenbaum, Vlastimil Ptak, and Zdenek Strakos, *Any Nonincreasing Convergence Curve is Possible for GMRES*, SIAM J. Matrix Anal. Appl., **17** (3), pp.465–469, 1996.