



CNPq / Inria

PaMPA : Parallel Mesh Partitioning and Adaptation

Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

Some results

Work in progress

Upcoming features

Contents

State of the art

Parallel remeshers

Load balancing

Existing tools

Context

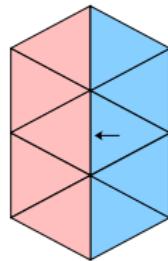
- ▶ Distributed meshes
- ▶ Subdomain decomposition
- ▶ Data exchanges between subdomains

Parallel mesh adaptation algorithms

- ▶ Parallel mesh generation (N. Chrisochoides 2005 [3])
 - ▶ Delaunay triangulation
 - ▶ Frontal method
 - ▶ Refinement by subdivisions (B. G. Larwood 2003 [7])
- ▶ Communication between subdomains:
 - ▶ Data migration
 - ▶ Matching algorithms

Problems induced by parallelism

- ▶ High complexity to parallelize remesh methods
 - ▶ Too much communication on boundaries
 - ▶ Boundaries not remeshed:
 - ▶ Local remeshing



Sequential algorithms

- ▶ Methods:
 - ▶ Moving nodes and remeshing locally (O. Hassan 2006 [6])
 - ▶ Coarse-grain parallel remeshing through multiple successive sequential remeshings (U. Tremel 2006 [8]):
 - ▶ Finding zones to remesh according to error estimator (T. Coupez 2000 [4])
 - ▶ Identifying zones to remesh in parallel
 - ▶ Remeshing zones in sequential
 - ▶ subdomain load balancing
- ▶ Main benefits:
 - ▶ Scalability of algorithms
 - ▶ Re-use of sequential codes

Contents

State of the art

Parallel remeshers

Load balancing

Existing tools

Generic dynamic load balancing

- ▶ General purpose (re)partitioning:
 - ▶ Jostle
 - ▶ Zoltan (K. Devine 2000 [5])
 - ▶ LB_Migrate (R. Chaube 2007 [2])
- ▶ Require a lot of extra coding

Specialized dynamic load balancing softwares

- ▶ Libraries:
 - ▶ DRAMA (A. Basermann 2000 [1])
- ▶ Pros and cons:
 - ▶ Mainly interfaced with solvers
 - ▶ Data structures based on meshes

Contents

State of the art

Parallel remeshers

Load balancing

Existing tools

Existing tools for handling unstructured meshes

- ▶ Partitioners:

- ▶ Chaco
- ▶ MeTiS
- ▶ Mondriaan
- ▶ Patoh
- ▶ Scotch
- ▶ Zoltan

- ▶ Intermediate:

- ▶ DRAMA
- ▶ PaMPA
- ▶ PHG

- ▶ Advanced:

- ▶ Arcane

Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

Some results

Work in progress

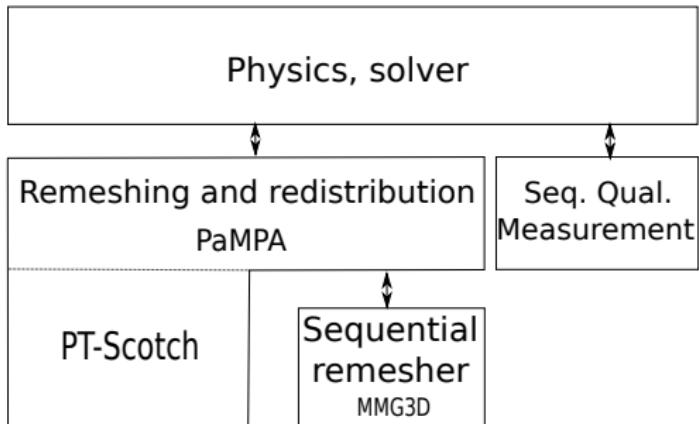
Upcoming features

Common needs of solvers regarding meshes

- ▶ Handling of mesh structures
- ▶ Distribution of meshes across the processors of a parallel architecture
 - ▶ Handling of load balance
- ▶ Data exchange across neighboring entities
- ▶ Iteration on mesh entities
 - ▶ Entities of any kind: e.g. elements, faces, edges, nodes, ...
 - ▶ Entity sub-classes: e.g. regular or boundary faces, ...
 - ▶ Inner or frontier entities with respect to neighboring processors
 - ▶ Maximization of cache effects thanks to proper data reordering
- ▶ Dynamic modification of mesh structure
 - ▶ Dynamic redistribution
- ▶ Adaptive remeshing

What is PaMPA

- ▶ PaMPA: “Parallel Mesh Partitioning and Adaptation”
- ▶ Middleware library managing the parallel repartitioning and remeshing of unstructured meshes modeled as interconnected valuated entities
- ▶ The user can focus on his/her “core business”:
 - ▶ Solver
 - ▶ Sequential remesher
 - ▶ Coupling with MMG3D provided for tetrahedra



Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

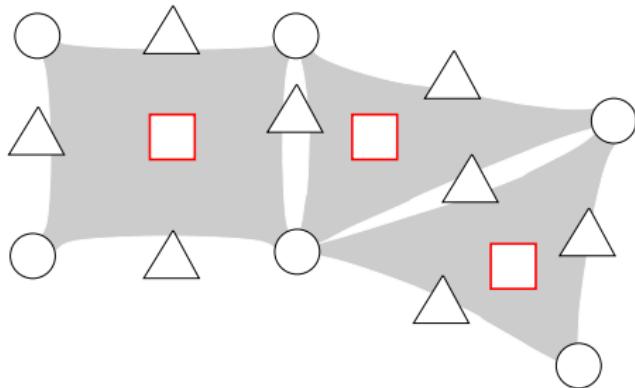
Some results

Work in progress

Upcoming features

Definitions

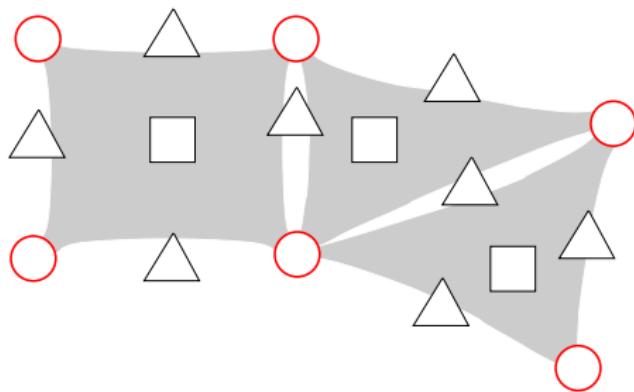
- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
 - ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**
- Top-level mesh entity
May bear some data (volume,
pressure, etc.)



Definitions

- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**

May bear some data (geometry, etc.)



Definitions

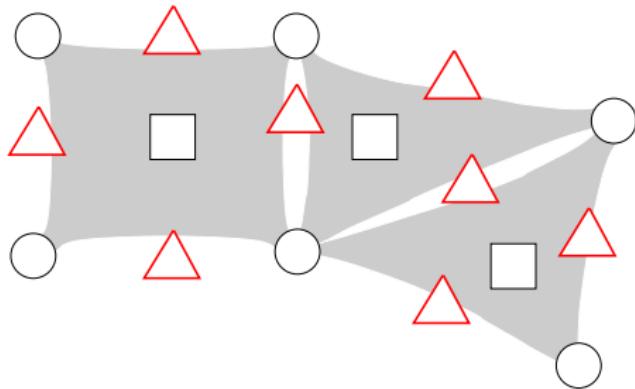
- ▶ Mesh:

- ▶ **Element**
- ▶ **Node**
- ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary

- ▶ PaMPA Mesh:

- ▶ **Vertex**
- ▶ **Relation**
- ▶ **Entity**
- ▶ **Sub-entity**
- ▶ **Enriched graph**

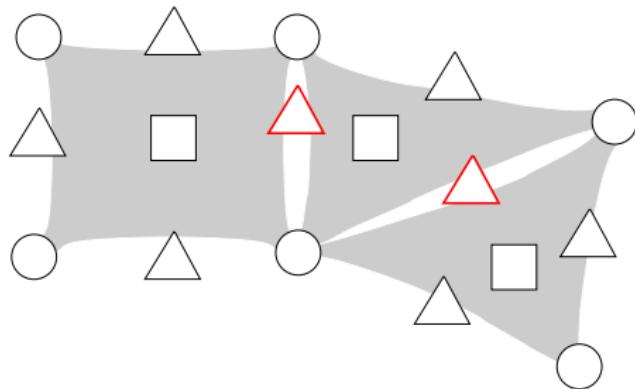
May bear some data (flux, etc.)



Definitions

- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**

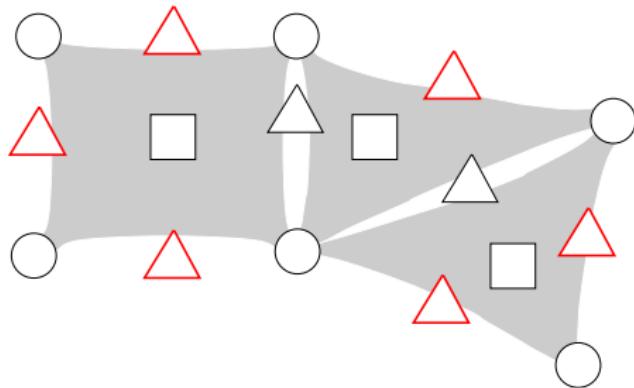
Regular mesh edge



Definitions

- ▶ Mesh:
 - ▶ Element
 - ▶ Node
 - ▶ Edge
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ Vertex
 - ▶ Relation
 - ▶ Entity
 - ▶ Sub-entity
 - ▶ Enriched graph

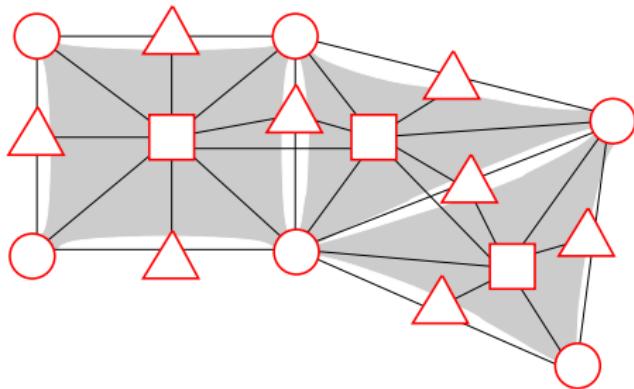
Boundary mesh edge



Definitions

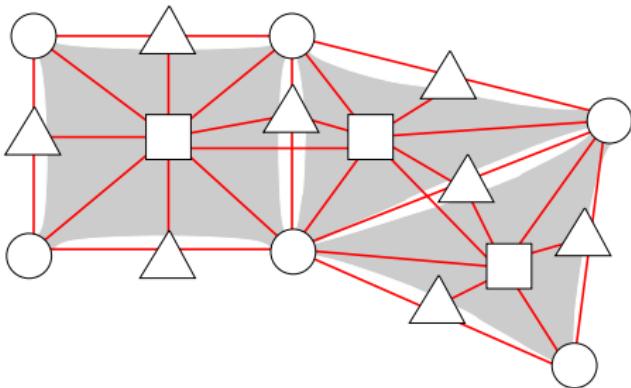
What all entities are in fact...

- ▶ Mesh:
 - ▶ Element
 - ▶ Node
 - ▶ Edge
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ Vertex
 - ▶ Relation
 - ▶ Entity
 - ▶ Sub-entity
 - ▶ Enriched graph



Definitions

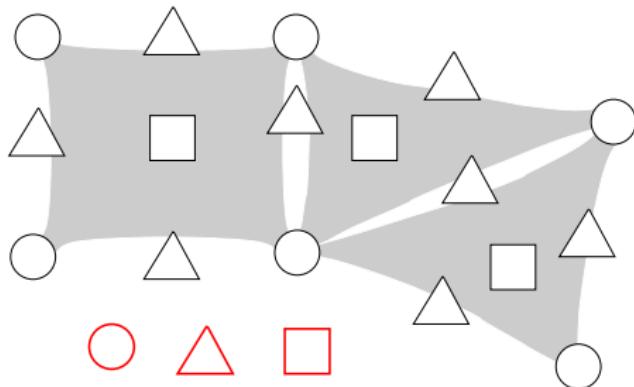
- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**



Definitions

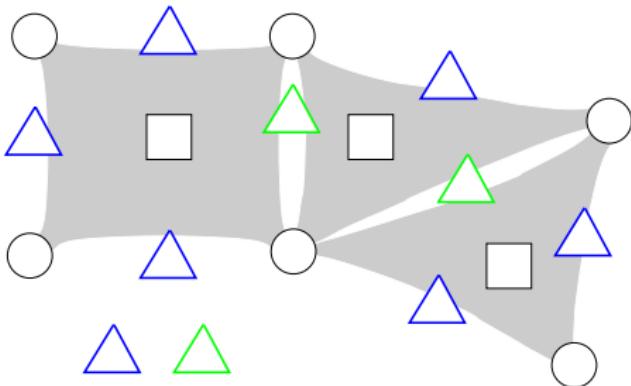
- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**

Subset of vertices bearing the same data



Definitions

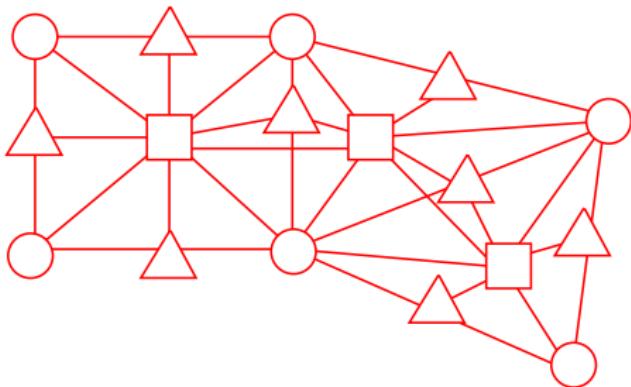
- ▶ Mesh:
 - ▶ **Element**
 - ▶ **Node**
 - ▶ **Edge**
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ **Vertex**
 - ▶ **Relation**
 - ▶ **Entity**
 - ▶ **Sub-entity**
 - ▶ **Enriched graph**



Definitions

- ▶ Mesh:
 - ▶ Element
 - ▶ Node
 - ▶ Edge
 - ▶ Internal
 - ▶ Boundary
- ▶ PaMPA Mesh:
 - ▶ Vertex
 - ▶ Relation
 - ▶ Entity
 - ▶ Sub-entity
 - ▶ Enriched graph

Whole set of vertices and relations
Every vertex belongs to one and only
one entity (and sub-entity)



Global vue

- ▶ All vertices have a global unique number

baseval

1

enttgbnbr

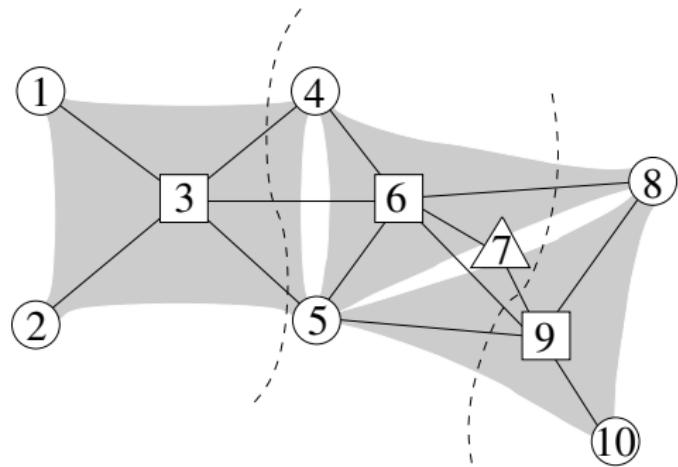
3

procctnttab

3 4 3

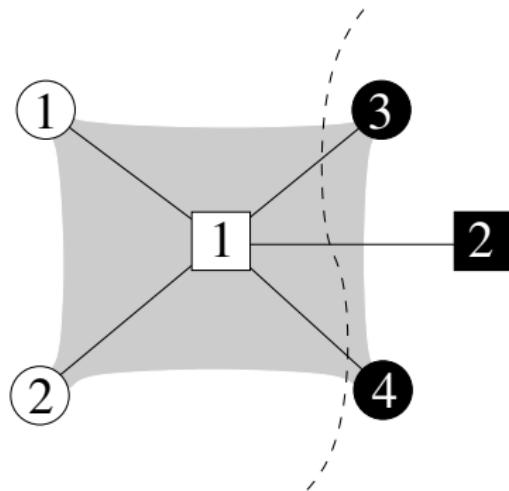
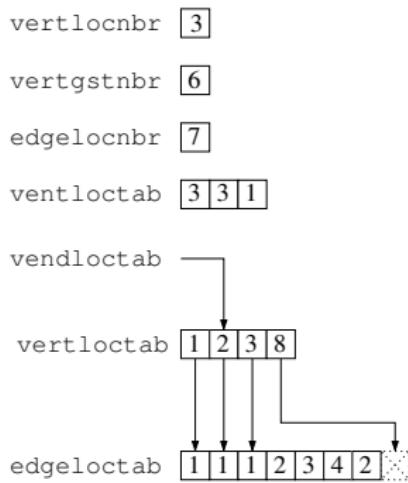
procvrttab

1 4 8 11



Local vision of process 0

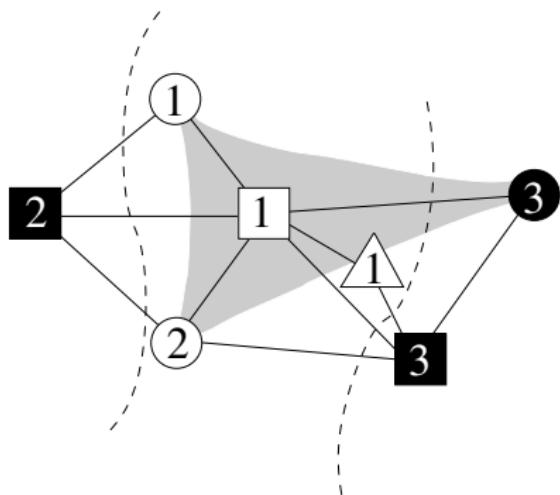
- ▶ All local and ghost vertices have a compact local index
 - ▶ Per-entity numbering



Local vision of process 1

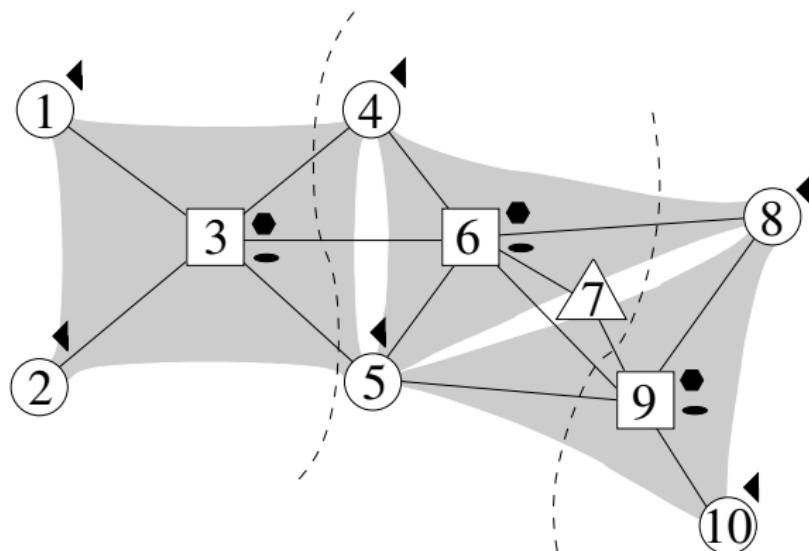
- ▶ All local and ghost vertices have a compact local index
 - ▶ Per-entity numbering

| | |
|-------------|-----------------------------|
| vertlocnbr | [4] |
| vertgstnbr | [7] |
| edgelocnbr | [13] |
| ventloctab | [3 3 1 2] |
| vendloctab | |
| vertloctab | [1 3 6 12 14] |
| edgelocatab | [2 1 2 1 3 2 1 2 1 3 3 1 3] |



Linking values to entities

- ▶ Multiple value types can be associated with each entity
 - ▶ Value data structures can be split to improve cache usage
- ▶ Entities without values serve as iterators



Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

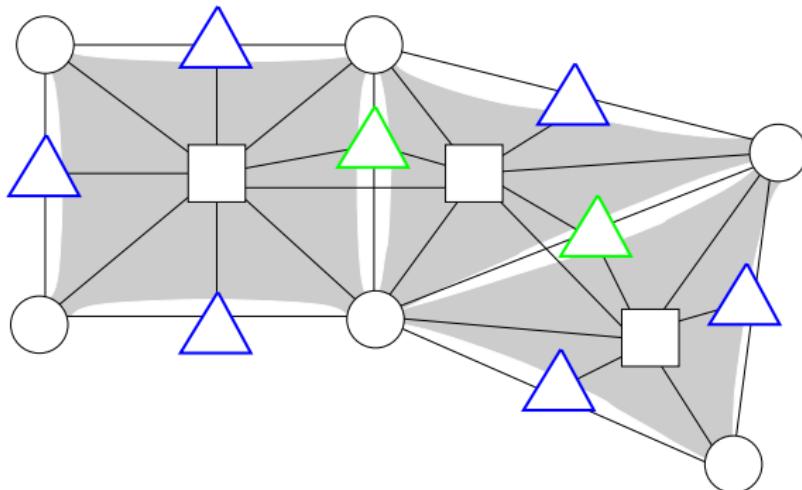
Some results

Work in progress

Upcoming features

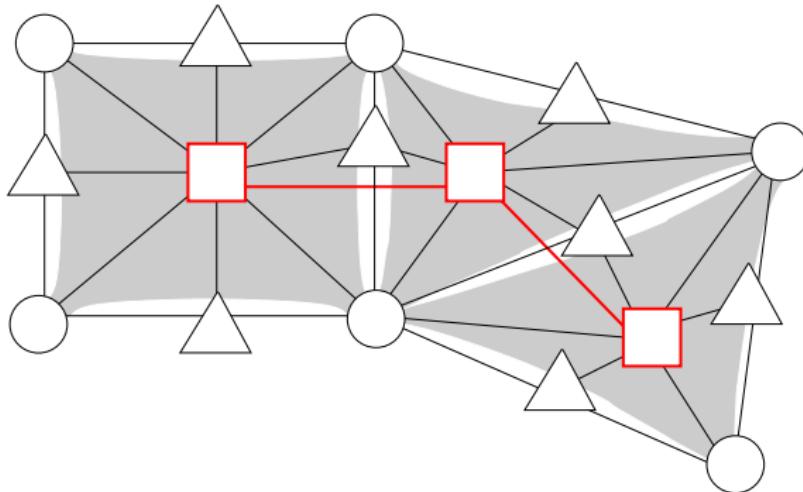
Entities

- ▶ Based
 - ▶ Smallest array index depends on language
 - ▶ 0 in **C** and 1 in **Fortran**
- ▶ Stable
 - ▶ Neighbor ordering preserved with respect to sub-entities



Partitioning

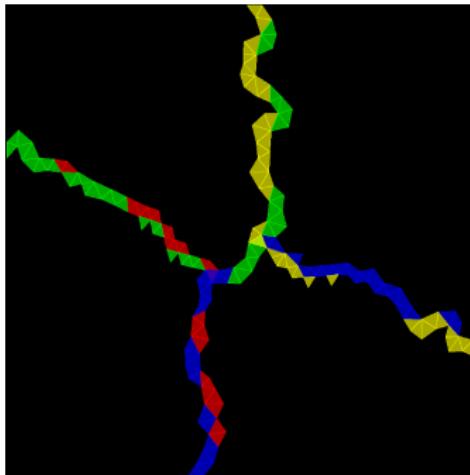
- ▶ Performed with respect to top-level entity
- ▶ Cache oblivious ordering
 - ▶ Overdecomposition into small groups of vertices
 - ▶ Vertices in groups ordered consecutively



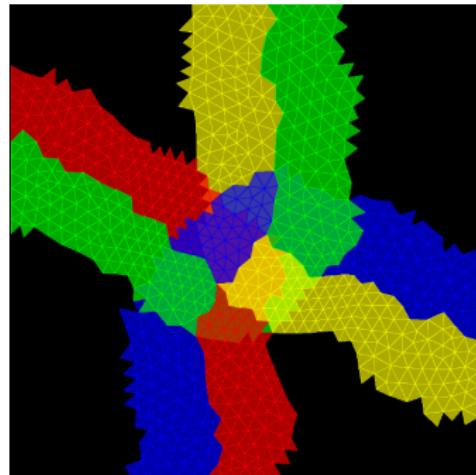
Overlap

- ▶ Of size n
- ▶ Requires top-level relations

Halo of size 1

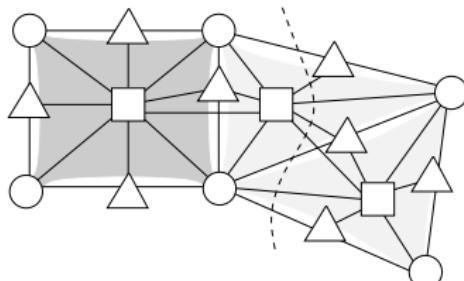
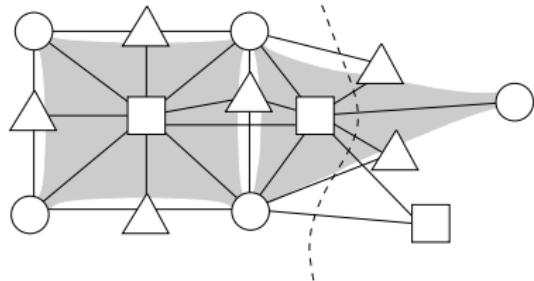


Halo of size 10



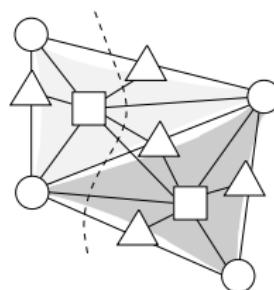
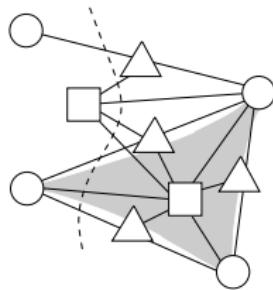
Local vue of process 0

- ▶ Halo: local copy of remote vertices linked to local vertices
- ▶ Overlap: local copy of remote vertices needed to complete split mesh units
 - ▶ Extensible to distance n
- ▶ Halo of size 1
- ▶ Overlap of size 1



Local view of process 1

- ▶ Halo: local copy of remote vertices linked to local vertices
- ▶ Overlap: local copy of remote vertices needed to complete split mesh units
 - ▶ Extensible to distance n
- ▶ Halo of size 1
- ▶ Overlap of size 1



Iterators

- ▶ Operate on entities or sub-entities

```
PAMPA_dmeshItInitStart (mesh, ventlocnum, PAMPA_VERT_ANY, &it_vrt) ;
PAMPA_dmeshItInit (mesh, ventlocnum, nentlocnum, &it_ngb) ;

while (PAMPA_itHasMore (&it_vrt)) {
    vertlocnum = PAMPA_itCurEnttVertNum (&it_vrt);
    printf ("vertlocnum : %d\n", vertlocnum);
    PAMPA_itStart (&it_ngb, vertlocnum) ;

    while (PAMPA_itHasMore (&it_ngb)) {
        PAMPA_Num sngblocnum, engblocnum, mngblocnum, nsblobcnum;
        mngblocnum = PAMPA_itCurMeshVertNum (&it_ngb);           // Global number of vertex
        engblocnum = PAMPA_itCurEnttVertNum (&it_ngb);          // Local number in entity
        sngblocnum = PAMPA_itCurSubEnttVertNum (&it_ngb);       // Local number in sub-entity
        nsblobcnum = PAMPA_itCurSubEnttNum (&it_ngb);           // Number of sub-entity
        printf ("%d_ngb_of_%d_with_entity_%d : %d_(sub : %d)_%d_(ent : %d)_%d_(glb)\n",
               rank, vertlocnum, ventlocnum, sngblocnum, nsblobcnum,
               engblocnum, nentlocnum, mngblocnum);
        PAMPA_itNext (&it_ngb);
    }
    PAMPA_itNext (&it_vrt);
}
```

PaMPA software distribution

- ▶ Project still in development stage
 - ▶ Soon licensed under GPL
 - ▶ Objective: create a community
- ▶ Already used in development of 2 fluid mechanic solvers:
 - ▶ Plato, INRIA Sophia-Antipolis
 - ▶ A platform for tokamak simulations
 - ▶ Mesh reading in sequential, partitioning and building distributed mesh
 - ▶ AeroSol, INRIA Bordeaux
 - ▶ High-order methods
 - ▶ Sequential and parallel executions
 - ▶ Great number of entities, mesh reading and redistributing in parallel

Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

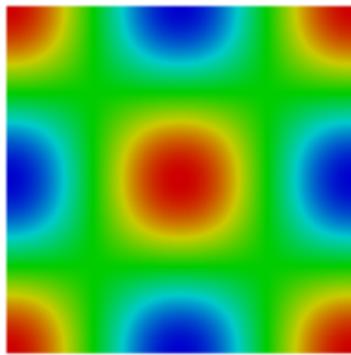
Some results

Work in progress

Upcoming features

Definition

- ▶ Solving 2D Poisson equation:
 - ▶ $\Delta u(x, y) = f(x, y)$
 - ▶ $g(x, y) = u(x, y)$ on the boundary Γ
- ▶ Test case:
 - ▶ $f(x, y) = -2 * \cos(x) * \cos(y)$
in the domain Ω
 - ▶ $g(x, y) = \cos(x) * \cos(y)$ on the boundary Γ
 - ▶ $u(x, y) = \cos(x) * \cos(y)$



Mesh properties

- ▶ Entities:
 - ▶ Elements
 - ▶ Nodes
 - ▶ Boundary edges
- ▶ Relations:
 - ▶ Element to element
 - ▶ Element to node
 - ▶ Element to boundary edge
 - ▶ Node to node
- ▶ Overlap of size 1
- ▶ Values:
 - ▶ Coordinates and solution on nodes
 - ▶ Type on boundary edges
 - ▶ Area, volume on elements

All steps

```
! Mesh loading and distribution
!
! Processor 0 only:
if (rank == 0) THEN
    CALL ReadMesh()           ! Read mesh on file
    CALL GraphPampa()          ! Build graph of vertex neighbors for PaMPA
    CALL CentralizedMesh()     ! Build PaMPA global non distributed mesh
END IF

! On all processors:
CALL DistributedMesh() ! Build PaMPA distributed mesh:
!
!                                     1— Scatter centralized mesh on all processors
!
!                                     2— Call PaMPA mesh partitionner
!
!                                     3— Redistribute distribute mesh
CALL ElementVolume()
CALL InitializeMatrixCSR()

! Solution computation
!
CALL InitSol()
CALL FillMatrix()
CALL SolveSystem()
CALL WriteDistributedMeshAndSolFiles()
```

FillMatrix

```

RHS = 0.
CALL PAMPAF_dmshItInitStart (dm, ENTITY_ELEM, PAMPAF_VERT_ANY, it_vrt, ierr)
CALL PAMPAF_dmshItInit (dm, ENTITY_ELEM, ENTITY_NODE, it_ngb, ierr)
DO WHILE (PAMPAF_itHasMore (it_vrt))
    jt = PAMPAF_itCurEntVertNum (it_vrt)
    Volt = VoIEI (jt)
    ngb = 0
    CALL PAMPAF_itStart (it_ngb, jt, ierr)
    DO WHILE (PAMPAF_itHasMore (it_ngb))
        ngb = ngb + 1
        is = PAMPAF_itCurEntVertNum (it_ngb)
        NuElemt(ngb) = is
        CoorElemt(:,ngb) = Coor(:,is)
        PAMPAF_itNext (it_ngb)
    END DO
    CALL GradPhi (CoorElemt(:,1), CoorElemt(:,2), CoorElemt(:,3), GrdPhi)
    DO i = 1, Nsmplx
        is = NuElemt(i)
        DO j = 1, Nsmplx
            js = NuElemt(j)
            JJac = Volt * Sum (GrdPhi(:,i) * GrdPhi(:,j))
            CALL assembly_addCSR (JJac, is, js)
        END DO ! loop on j
        RHS(is) = RHS(is) - Volt * SourceTerm (Coor(1,is), Coor(2,is)) / Nsmplx
    END DO ! loop on i
    PAMPAF_itNext (it_vrt)
END DO

```

Solve system: Jacobi (1/2)

```

UaPrec = 0.           ! Suppose  $A = L + D + U$ , system to solve :  $A \cdot x = b$ 
CALL PAMPAF_dmshItInit(dm, ENTITY_NODE, ENTITY_NODE, it_ngb, ierr)
DO irelax = 1, Nrelax
  res      = 0.
  CALL PAMPAF_dmshItInitStart(dm, ENTITY_NODE, PAMPAF_VERT_BOUNDARY, it_vrt, ierr)
  DO WHILE (PAMPAF_itHasMore(it_vrt))
    is = PAMPAF_itCurEntVertNum(it_vrt)
    CALL PAMPAF_dmshMatLineData(dm, ENTITY_NODE, is, I1, I1Fin, ierr)
    CALL PAMPAF_itStart( it_ngb, is, ierr )
    res0     = RHS(is)                                ! res0 = b
    iv = i1
    DO WHILE (PAMPAF_itHasMore(it_ngb))
      js = PAMPAF_itCurEntVertNum(it_ngb)
      PAMPAF_itNext(it_ngb)
      res0 = res0 - MatCSR%Vals(iv) * UaPrec(js) ! res0 = b - (L + U) \cdot x^n
      iv = iv + 1
    END DO
    Ua(is) = res0 / MatCSR%Diag(is)                ! x^{n+1} = ( b - (L + U) \cdot x^n ) / D
    PAMPAF_itNext(it_vrt)
  END DO
  CALL PAMPAF_dmshHaloValueAsync(dm, ENTITY_NODE, PAMPA_TAG_SOL, req, ierr)

```

Solve system: Jacobi (2/2)

```

CALL PAMPAF_dmeshItInitStart(dm, ENTITY_NODE, PAMPAF_VERT_INTERNAL, it_vrt, ierr)
DO WHILE (PAMPAF_itHasMore(it_vrt))
    is = PAMPAF_itCurEnttVertNum(it_vrt)
    CALL PAMPAF_dmeshMatLineData(dm, ENTITY_NODE, is, I1, I1Fin, ierr)
    CALL PAMPAF_itStart( it_ngb, is, ierr)
    res0 = RHS(is)                                ! res0 = b

    iv = i1
    DO WHILE (PAMPAF_itHasMore(it_ngb))
        js = PAMPAF_itCurEnttVertNum(it_ngb)
        PAMPAF_itNext(it_ngb)
        res0 = res0 - MatCSR%Vals(iv) * UaPrec(js) ! res0 = b - (L + U) x^n
        iv = iv + 1
    END DO
    Ua(is) = res0 / MatCSR%Diag(is)                ! x^{n+1} = ( b - (L + U) x^n )/D
    PAMPAF_itNext(it_vrt)
END DO

CALL PAMPAF_dmeshHaloWait(req, ierr)

UaPrec = Ua
END DO ! end loop on irelax

```

Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

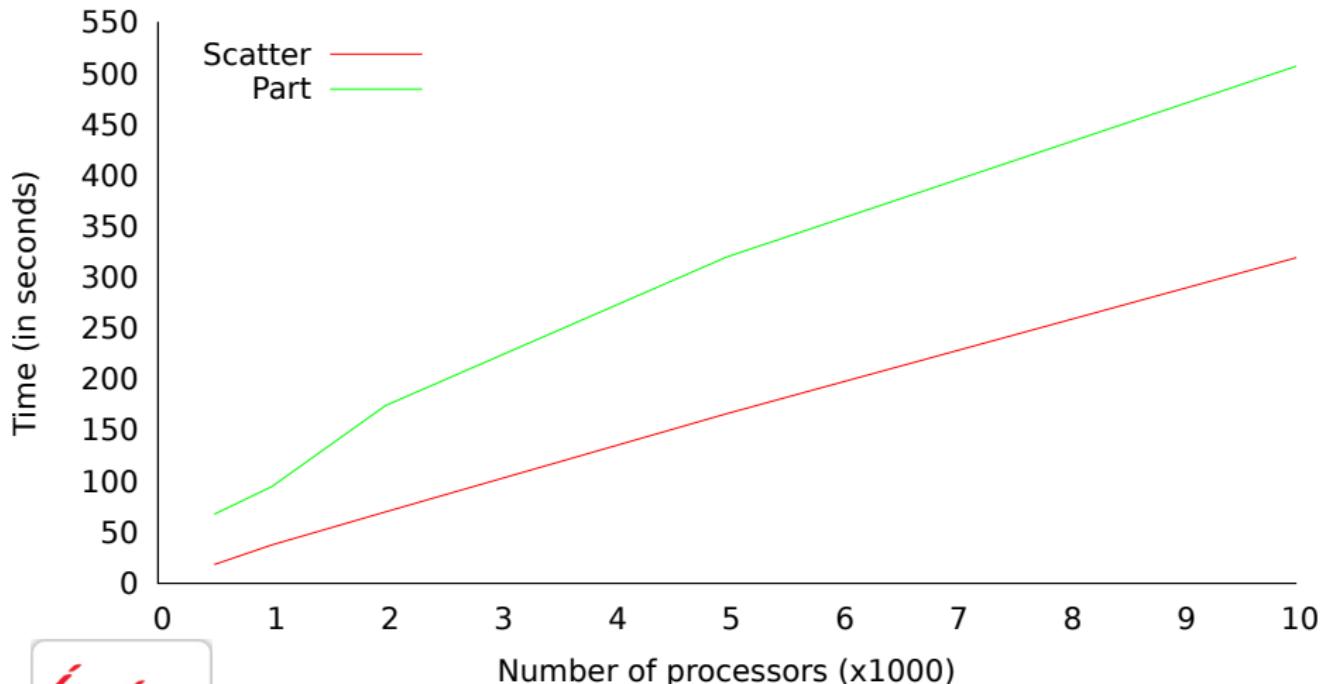
Some results

Work in progress

Upcoming features

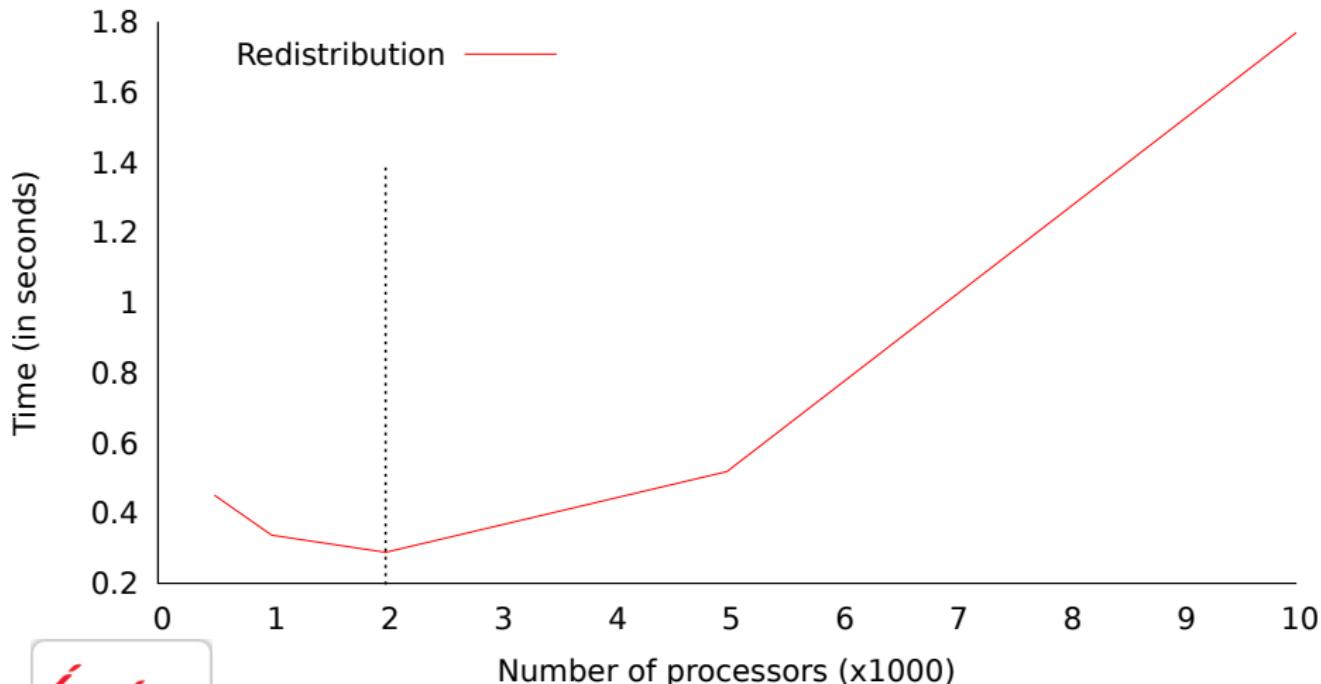
First results (1/2)

Time for distributing centralized mesh and partitionning distributed mesh
(3D cube with 3.3M nodes and 20M tetrahedra)



First results (2/2)

Time for redistributing distributed mesh according to the partition
(3D cube with 3.3M nodes and 20M tetrahedra)



Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

Some results

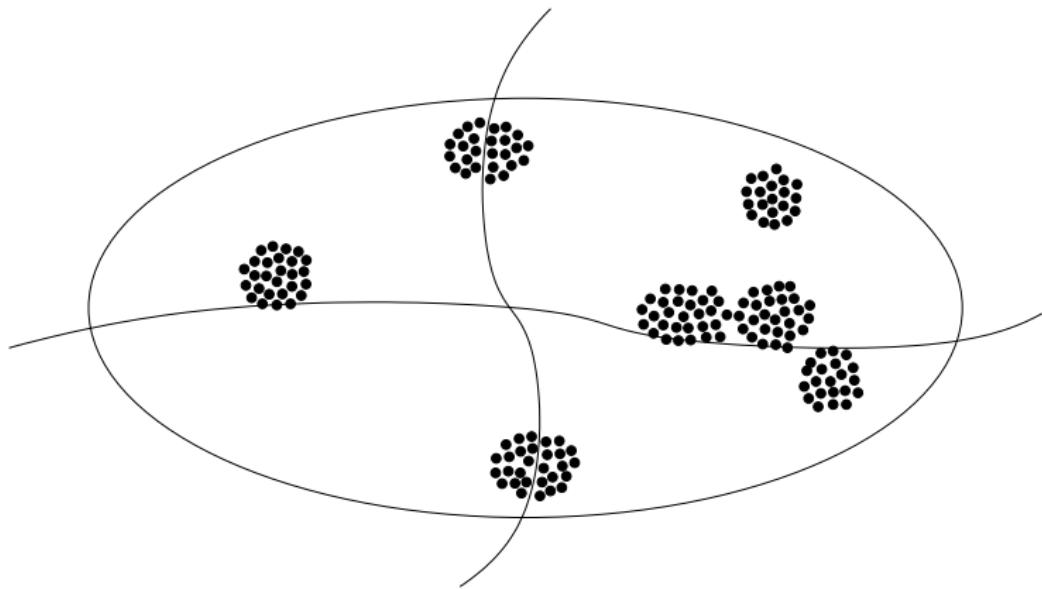
Work in progress

Upcoming features

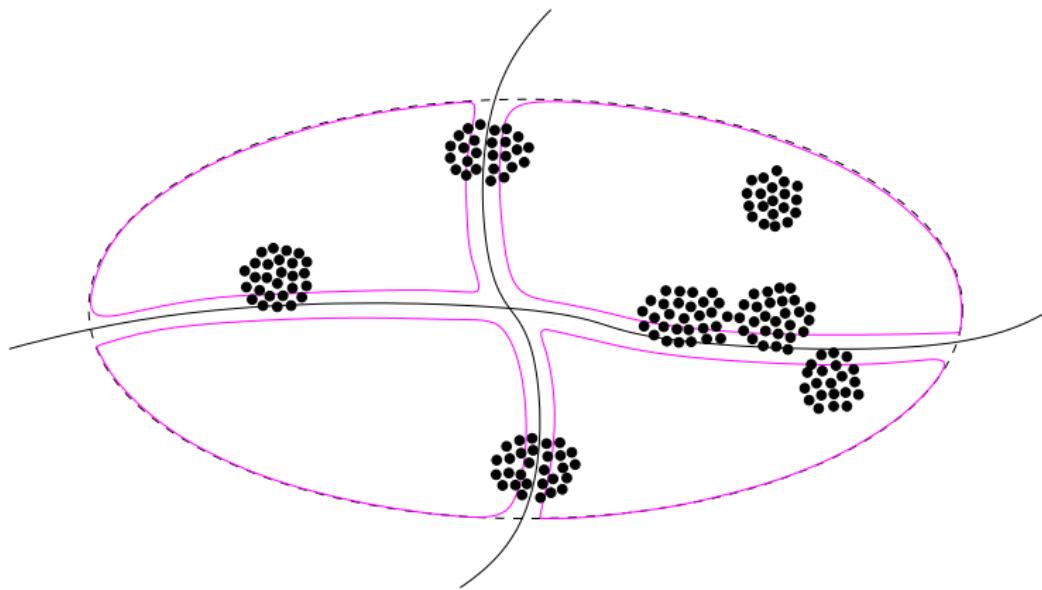
Work in progress

- ▶ Release of version 0.1
 - ▶ Available soon from Inria Gforge
 - ▶ Licensed under GPL
- ▶ Parallel mesh adaptation based on sequential remesher

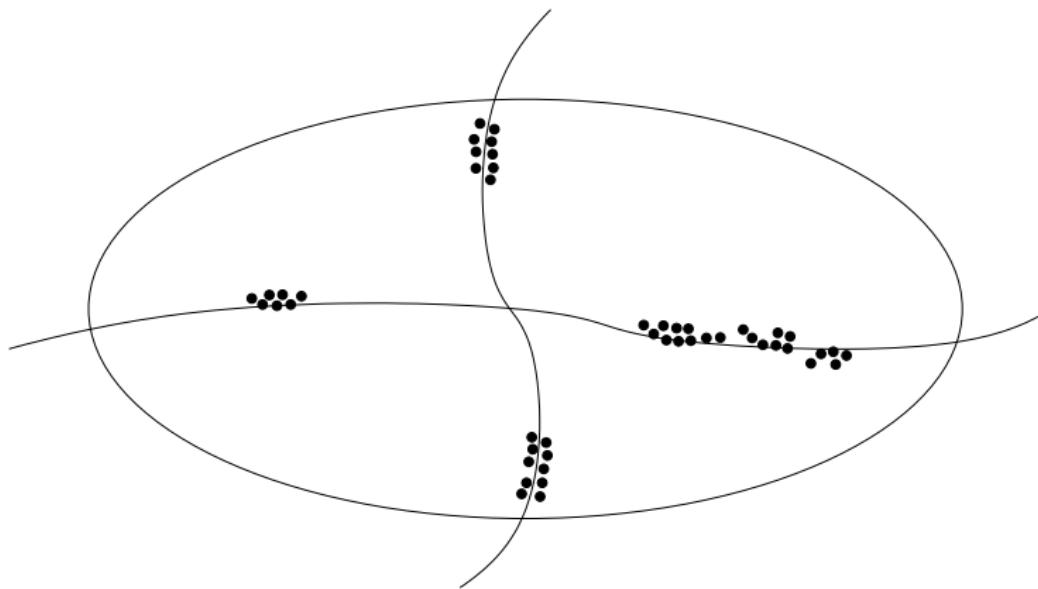
Adaptation: zones to remesh



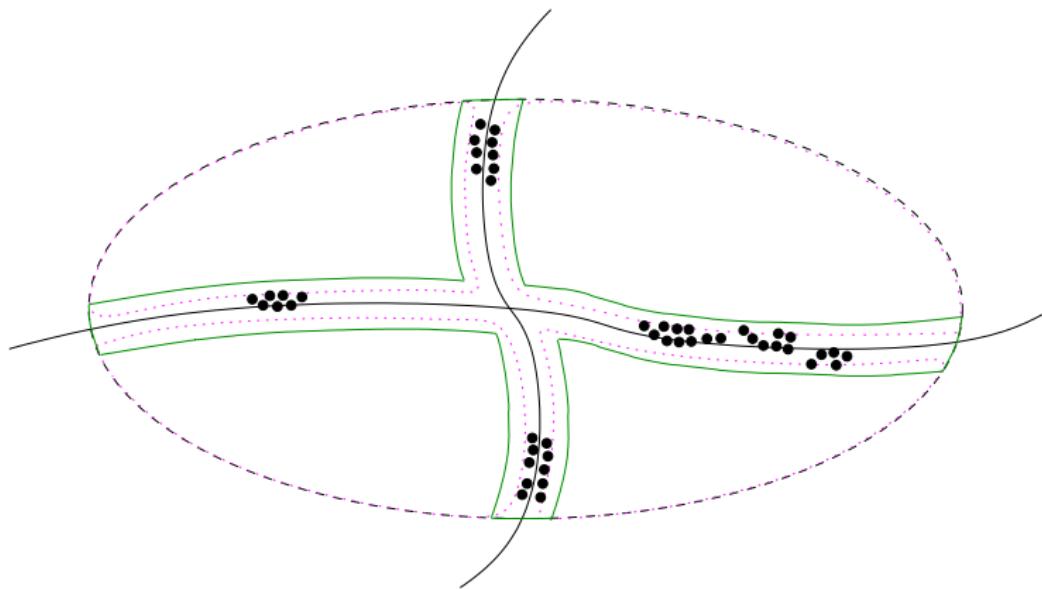
Adaptation: local remeshing



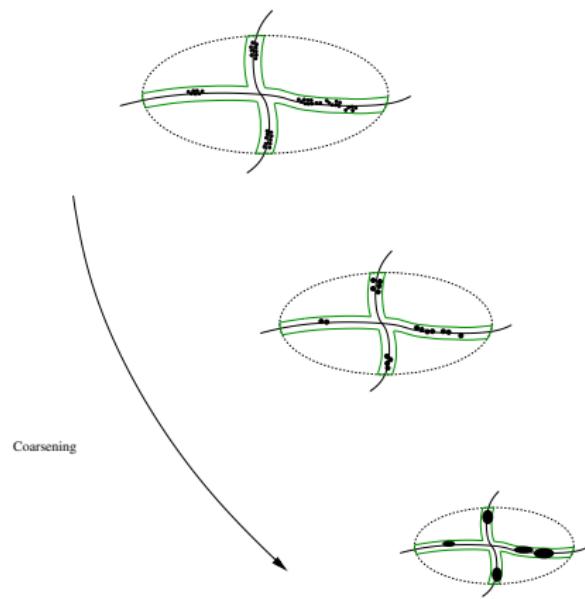
Adaptation: remaining zones



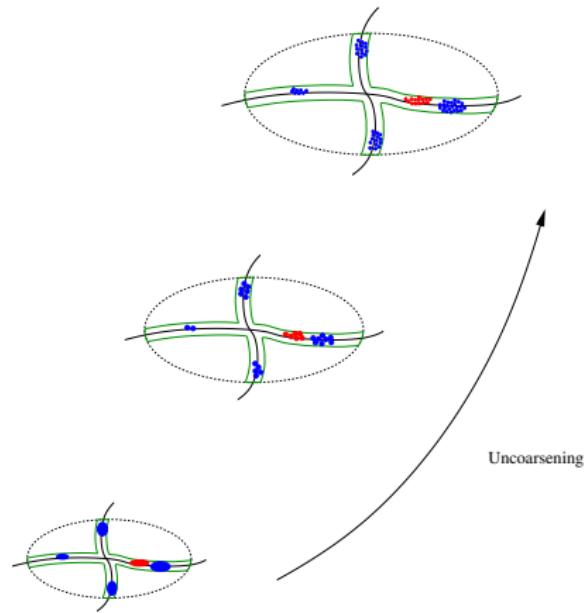
Adaptation: remeshing of frontier zones



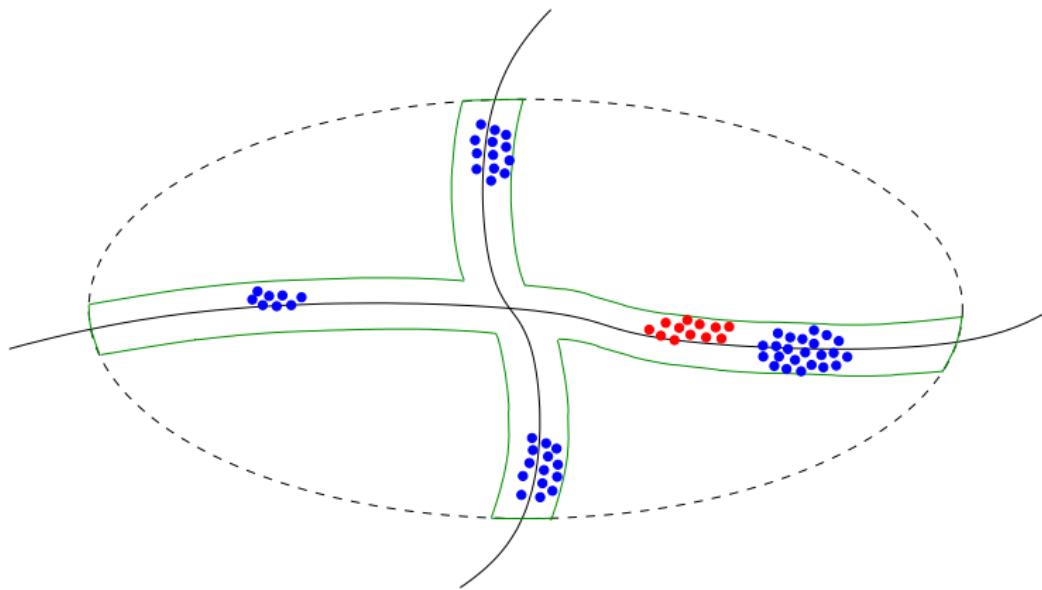
Adaptation: coarsening (iteration 1)



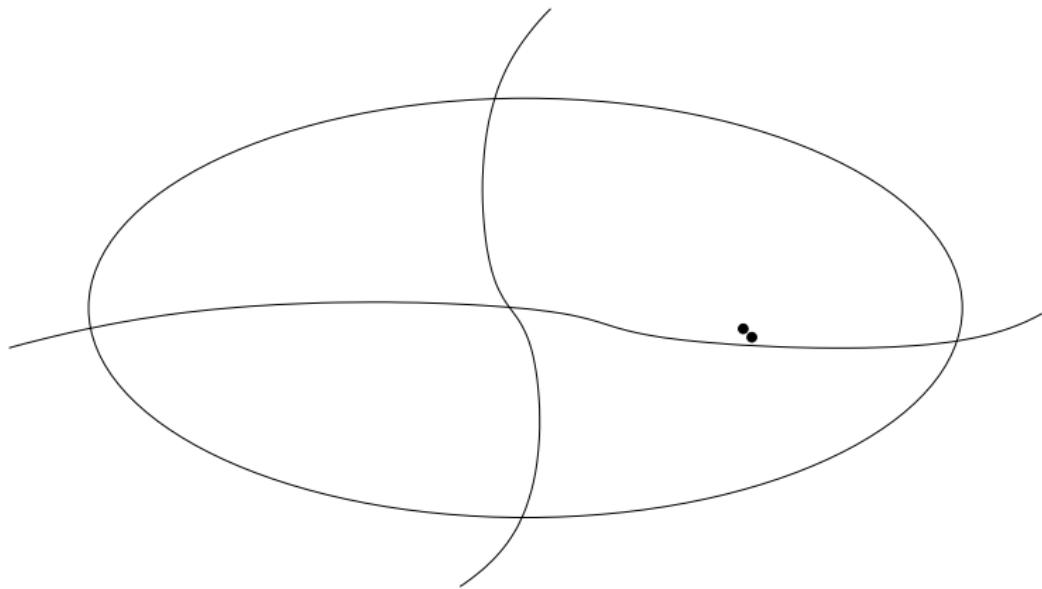
Adaptation: uncoarsening (iteration 1)



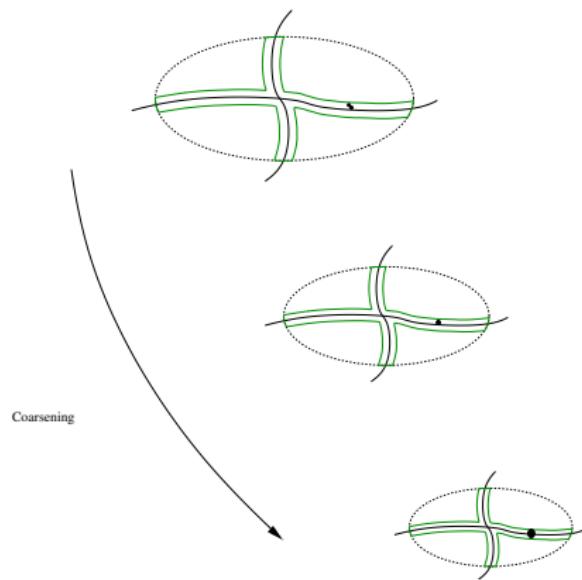
Adaptation: induced subgraphs (iteration 1)



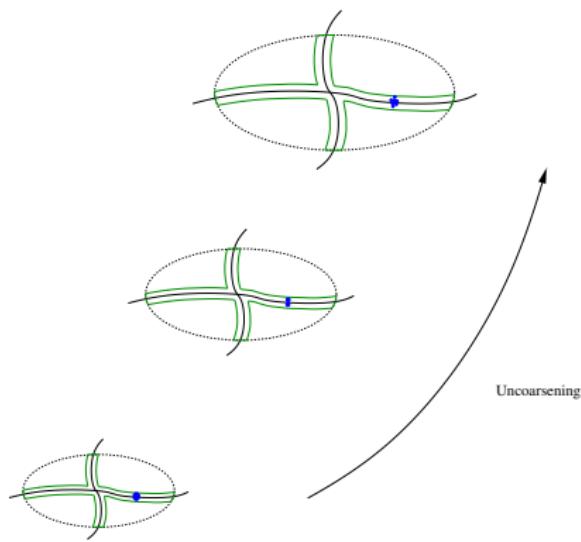
Adaptation: remaining zones



Adaptation: coarsening (iteration 2)



Adaptation: uncoarsening (iteration 2)



Contents

State of the art

Common needs of solvers regarding meshes

Data structures

Version 0.1

Example: Laplacian equation using P_1 finite element method

Some results

Work in progress

Upcoming features

Upcoming features

- ▶ Parallel I/O with HDF5
- ▶ Unbreakable relations
 - ▶ Partitioner will not cut these edges
 - ▶ E.g. to implement DG methods
- ▶ Source code parser to infer automatically:
 - ▶ Overlap and halo sizes
 - ▶ Entity and relation types

THANK YOU



-  A. Basermann, J. Clinckemaillie, T. Coupez, J. Fingberg, H. Digonnet, R. Ducloux, J.-M. Gratien, U. Hartmann, G. Lonsdale, B. Maerten, D. Roose, and C. Walshaw.
Dynamic load balancing of finite element applications with the drama library, 2000.
-  R. Chaube, R. L. Carino, and I. Banicescu.
Effectiveness of a dynamic load balancing library for scientific applications.
In *ISPDC '07: Proceedings of the Sixth International Symposium on Parallel and Distributed Computing*, page 32, Washington, DC, USA, 2007. IEEE Computer Society.
-  N. Chrisochoides.
Parallel mesh generation.
In *Numerical Solution of Partial Differential Equations on Parallel Computers*. Springer, 2005.

-  T. Coupez, H. Digonnet, and R. Ducloux.
Parallel meshing and remeshing.
Applied Mathematical Modelling, 25(2):153 – 175, 2000.
Dynamic load balancing of mesh-based applications on parallel.
-  K. Devine, B. Hendrickson, E. Boman, M. St. John, and C. Vaughan.
Design of dynamic load-balancing tools for parallel applications.
In *ICS '00: Proceedings of the 14th international conference on Supercomputing*, pages 110–118, New York, NY, USA, 2000. ACM.

-  O. Hassan, K. A. Sørensen, K. Morgan, and N. P. Weatherill.
A method for time accurate turbulent compressible fluid flow simulation with moving boundary components employing local remeshing.
In *International Journal for Numerical Methods in Fluids*, volume 53, pages 1243–1266. John Wiley & Sons, 2007.
-  B. G. Larwood, N. Weatherhill, O. Hassan, and K. Morgan.
Domain decomposition approach for parallel unstructured mesh generation.
In *International Journal for Numerical Methods in Engineering*, volume 58, pages 177–188, 2003.

-  U. Tremel, K. A. Sørensen, S. Hitzel, H. Rieger, O. Hassan, and N. P. Weatherill.
Parallel remeshing of unstructured volume grids for cfd applications.
In *International Journal for Numerical Methods in Fluids*, volume 53, pages 1361–1379. John Wiley & Sons, 2007.