Optimal Partitioning Of Multi-Block Structured Grids

Mohammad Akhavan[1] Patrice Castonguay[2] Marina Chugunova[3] Julio C. Góez[4] Mikko Kivelä[5] Odile Marcotte[6] Christophe Meyer[7] Tina Mitre[8] Ragheb Rahmaniani[4]

[1]Université de Montréal, [2]Bombardier, [3]Claremont Graduate University, [4]École Polytechnique de Montréal, [5]University of Oxford, [6]CRM and UQAM, [7]UQAM, [8]McGill University

August 21, 2015

- The Advanced Aerodynamics Department uses the FANSC software to solve the Navier-Stokes equations.
- These are coupled non-linear partial differential equations that govern the motion of compressible, viscous fluids such as air.
- The Navier-Stokes equations are solved on a fluid volume region discretized into a large number of cells.
- The solution of the equations yields the values of the density, velocity, and pressure within each cell.
- The computation is distributed over many processors (1000 today, maybe 10000 in a few years from now...).

The discretization produces regions that “look like” grids.

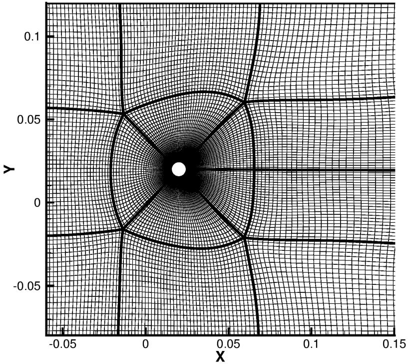

The union of several regions may (or may not) look like a grid.

Allocating Computations to Processors: the Objectives

One must find a way of allocating groups of cells to processors. A group of cells will consist of several grid-like regions called "blocks".

- The computational load must be balanced, i.e., the difference between the numbers of cells allocated to any two processors must be as small as possible.
- The maximum communication cost between a processor and all the other processors must be as small as possible.
- The number of blocks allocated to a given processor must not be too large.
- None of the blocks allocated to a given processor must be too small or too "thin."

Decomposition of the Problem

Since the problem is too complicated to be solved as a whole, we have decomposed it into subproblems.

1. Refine the original partition (produced by an engineer) so as to obtain small blocks and create a graph with weights on nodes and edges.
2. Partition the graph into clusters (i.e., groups of blocks assigned to a given processor).
3. Reorganize the blocks within a given cluster so as to obtain relatively large blocks that are grid-like and are not too "thin."

```
# BLOCKS
    5
=====Block   Imax    Jmax    Kmax    Total surfaces================
    1      65      65      97        6
  ---BCtype       Ista    Jsta    Ksta    Iend    Jend    Kend    icomp   iiconnect---
    CON            1      65       1      65      65      97       0
     2            97       1       1     161       1      97               2  0
    SYM            1       1       1       1      65      97       0
    SYM           65       1       1      65      65      97       0
    SYM            1       1       1      65       1      97       0
    WAL            1       1       1      65      65       1     100
    FAR            1       1      97      65      65      97       0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    2     257      65      97       10
  ---BCtype       Ista    Jsta    Ksta    Iend    Jend    Kend    icomp   iiconnect---
    CON           97       1       1     161       1      97       0
     1             1      65       1      65      65      97               2  0
    CON            1       1       1       1      65      97       0
     2           257       1       1     257      65      97               1  2
    CON          257       1       1     257      65      97       0
     2             1       1       1       1      65      97               1  2
    CON            1       1       1      33       1      97       0
     5            33       1       1       1       1      97               2  0
    CON          225       1       1     257       1      97       0
     5            65       1       1      33       1      97               2  0
    CON            1      65       1     257      65      97       0
     3             1       1       1     257       1      97               2  0
    SYM           33       1       1      97       1      97       0
    SYM          161       1       1     225       1      97       0
    WAL            1       1       1     257      65       1     100
    FAR            1       1      97     257      65      97       0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    3     257      33      97        6
  ---BCtype       Ista    Jsta    Ksta    Iend    Jend    Kend    icomp   iiconnect---
    CON            1       1       1     257       1      97       0
     2             1      65       1     257      65      97               2  0
    CON            1       1       1       1      33      97       0
     3           257       1       1     257      33      97               1  2
    CON          257       1       1     257      33      97       0
     3             1       1       1       1      33      97               1  2
    CON            1      33       1     257      33      97       0
     4             1       1       1     257       1      97               2  0
    WAL            1       1       1     257      33       1       1
    FAR            1       1      97     257      33      97       0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    4     257      65      97        7
  ---BCtype       Ista    Jsta    Ksta    Iend    Jend    Kend    icomp   iiconnect---
    CON            1       1       1     257       1      97       0
     3             1      33       1     257      33      97               2  0
    CON            1       1       1       1      65      97       0
     4           257       1       1     257      65      97               1  2
    CON          257       1       1     257      65      97       0
     4             1       1       1       1      65      97               1  2
    CON            1      65       1     129      65      97       0
     4           257      65       1     129      65      97               2  0
    CON          129      65       1     257      65      97       0
     4           129      65       1       1      65      97               2  0
    WAL            1       1       1     257      65       1       1
    FAR            1       1      97     257      65      97       0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    5      65     129      97        7
```

Figure 1: Topology file

Topology Refinement and Graph Structure

```
# BLOCKS
     5
=====Block   Imax    Jmax    Kmax    Total surfaces================
    1      65        65         97         6
  ---BCtype       Ista     Jsta     Ksta     Iend     Jend     Kend     icomp   iiconnect---
     CON            1       65        1       65       65       97        0
      2            97        1        1      161        1       97                2  0
     SYM            1        1        1        1       65       97        0
     SYM           65        1        1       65       65       97        0
     SYM            1        1        1       65        1       97        0
     WAL            1        1        1       65       65        1      100
     FAR            1        1       97       65       65       97        0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    2     257        65         97        10
  ---BCtype       Ista     Jsta     Ksta     Iend     Jend     Kend     icomp   iiconnect---
     CON           97        1        1      161        1       97        0
      1             1       65        1       65       65       97                2  0
     CON            1        1        1        1       65       97        0
      2           257        1        1      257       65       97                1  2
     CON          257        1        1      257       65       97        0
      2             1        1        1        1       65       97                1  2
     CON            1        1        1       33        1       97        0
      5            33        1        1        1        1       97                2  0
     CON          225        1        1      257        1       97        0
      5            65        1        1       33        1       97                2  0
     CON            1       65        1      257       65       97        0
      3             1        1        1      257        1       97                2  0
     SYM           33        1        1       97        1       97        0
     SYM          161        1        1      225        1       97        0
     WAL            1        1        1      257       65        1      100
     FAR            1        1       97      257       65       97        0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    3     257        33         97         6
  ---BCtype       Ista     Jsta     Ksta     Iend     Jend     Kend     icomp   iiconnect---
     CON            1        1        1      257        1       97        0
      2             1       65        1      257       65       97                2  0
     CON            1        1        1        1       33       97        0
      3           257        1        1      257       33       97                1  2
     CON          257        1        1      257       33       97        0
      3             1        1        1        1       33       97                1  2
     CON            1       33        1      257       33       97        0
      4             1        1        1      257        1       97                2  0
     WAL            1        1        1      257       33        1        1
     FAR            1        1       97      257       33       97        0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    4     257        65         97         7
  ---BCtype       Ista     Jsta     Ksta     Iend     Jend     Kend     icomp   iiconnect---
     CON            1        1        1      257        1       97        0
      3             1       33        1      257       33       97                2  0
     CON            1        1        1        1       65       97        0
      4           257        1        1      257       65       97                1  2
     CON          257        1        1      257       65       97        0
      4             1        1        1        1       65       97                1  2
     CON            1       65        1      129       65       97        0
      4           257       65        1      129       65       97                2  0
     CON          129       65        1      257       65       97        0
      4           129       65        1        1       65       97                2  0
     WAL            1        1        1      257       65        1        1
     FAR            1        1       97      257       65       97        0
=====Block   Imax    Jmax    Kmax    Total surfaces================
    5      65       129         97         7
```

Figure 1: Topology file

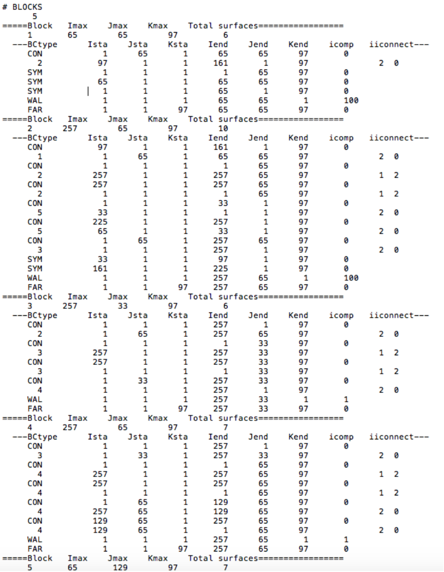

```
# BLOCKS
    5
=====Block   Imax     Jmax     Kmax     Total surfaces================
    1       65       65       97       6
  ---BCtype      Ista     Jsta     Ksta     Iend     Jend     Kend     icomp    iiconnect---
     CON          1        65       1        65       65       97       0
      2           97       1        1        161      1        97                2  0
     SYM          1        1        1        1        65       97       0
     SYM          65       1        1        65       65       97       0
     SYM          1        1        1        65       1        97       0
     WAL          1        1        1        65       65       1        100
     FAR          1        1        97       65       65       97       0
=====Block   Imax     Jmax     Kmax     Total surfaces================
    2       257      65       97       10
  ---BCtype      Ista     Jsta     Ksta     Iend     Jend     Kend     icomp    iiconnect---
     CON          97       1        1        161      1        97       0
      1           1        65       1        65       65       97                2  0
     CON          1        1        1        1        65       97       0
      2           257      1        1        257      65       97                1  2
     CON          257      1        1        257      65       97       0
      2           1        1        1        1        65       97                1  2
     CON          1        1        1        33       1        97       0
      5           33       1        1        1        1        97                2  0
     CON          225      1        1        257      1        97       0
      5           65       1        1        33       1        97                2  0
     CON          1        65       1        257      65       97       0
      3           1        1        1        257      1        97                2  0
     SYM          33       1        1        97       1        97       0
     SYM          161      1        1        225      1        97       0
     WAL          1        1        1        257      65       1        100
     FAR          1        1        97       257      65       97       0
=====Block   Imax     Jmax     Kmax     Total surfaces================
    3       257      33       97       6
  ---BCtype      Ista     Jsta     Ksta     Iend     Jend     Kend     icomp    iiconnect---
     CON          1        1        1        257      1        97       0
      2           1        65       1        257      65       97                2  0
     CON          1        1        1        1        33       97       0
      3           257      1        1        257      33       97                1  2
     CON          257      1        1        257      33       97       0
      3           1        1        1        1        33       97                1  2
     CON          1        33       1        257      33       97       0
      4           1        1        1        257      1        97                2  0
     WAL          1        1        1        257      33       1        1
     FAR          1        1        97       257      33       97       0
=====Block   Imax     Jmax     Kmax     Total surfaces================
    4       257      65       97       7
  ---BCtype      Ista     Jsta     Ksta     Iend     Jend     Kend     icomp    iiconnect---
     CON          1        1        1        257      1        97       0
      3           1        33       1        257      33       97                2  0
     CON          1        1        1        1        65       97       0
      4           257      1        1        257      65       97                1  2
     CON          257      1        1        257      65       97       0
      4           1        1        1        1        65       97                1  2
     CON          1        65       1        129      65       97       0
      4           257      65       1        129      65       97                2  0
     CON          129      65       1        257      65       97       0
      4           129      65       1        1        65       97                2  0
     WAL          1        1        1        257      65       1        1
     FAR          1        1        97       257      65       97       0
=====Block   Imax     Jmax     Kmax     Total surfaces================
    5       65       129      97       7
```

Figure 1: Topology file

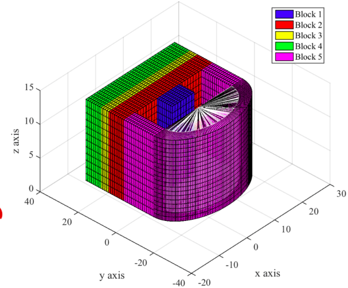

Figure 2: Three dimensional plot of the surface provided

Grid Examples Used For The Partitioning Model

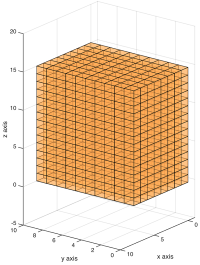

Figure 3: $8 \times 10 \times 15$ structured grid

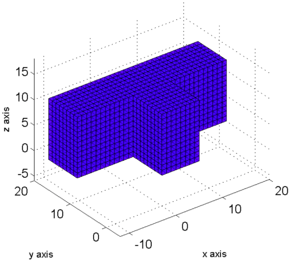

Figure 4: T shape object formed by 2 structured grids

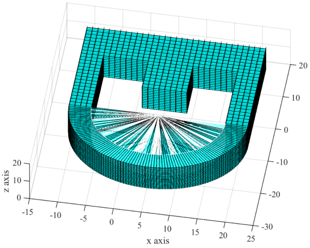

Figure 5: Shape object formed by 3 structured grids

Partitioning Problem Of The Block Graphs (1)

Idea: Turn our problem into a balanced graph clustering problem.

Input: Graph $G = (V, E)$, node weights $w_n : V \mapsto \mathbb{N}_+$, edge weights $w_e : E \mapsto \mathbb{N}_+$, number of colours k.

Objective: Find a colouring of nodes $c : V \mapsto P$ where $P = \{1, \ldots, k\}$ such that $T = c_1 \max_{p \in P} V_p + c_2 \max_{p \in P} B_p$ is minimized. Here the volume of a partition is $V_p = \sum_{u \in V} w_n(u)\delta(c(u), p)$ and the boundary of a partition is $B_p = \sum_{(u,v) \in E} w_e((u, v))[1 - \delta(c(u), c(v))]\delta(c(u), p)$, where δ denotes the indicator function.

Methods for balanced graph clustering:

- Heuristic based on label propagation
- Integer programming formulation
- Freely available programs: KaHIP and METIS (slightly different objective functions)

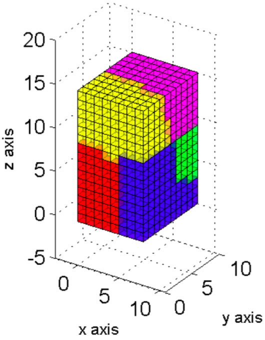

Figure 6: $8 \times 10 \times 15$ structured grid

KaHIP Applied To Structured Grids

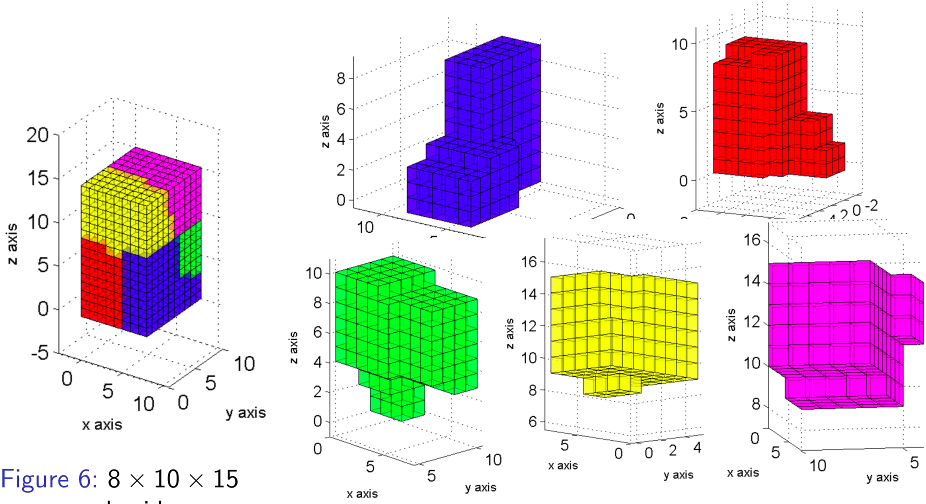

Figure 6: $8 \times 10 \times 15$ structured grid

Heuristic Partitioning Method Applied To Structured Grid

0. Initialize
    - Set nodes to random partitions
    - Use previous partition as a starting point
1. For each node $u \in V$ (in random order):
    - Find colours in the neighborhood of u: $C_{\text{candidates}} = \{c(v) : (u, v) \in E\}$
    - For each $p \in C_{\text{candidates}}$ find change in the objective function $\Delta_{u,p}T$ if $c(u)$ is set to p.
    - Select p^* that minimizes $\Delta_{u,p}T$. If $\Delta_{u,p^*}T < 0$, set $c(u)$ to p^*.
2. Repeat point 1, if at least one colour was changed in last run of point 1.

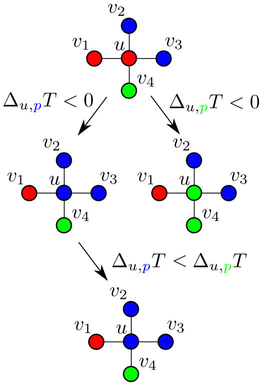

Heuristic Partitioning Method Applied To Structured Grid

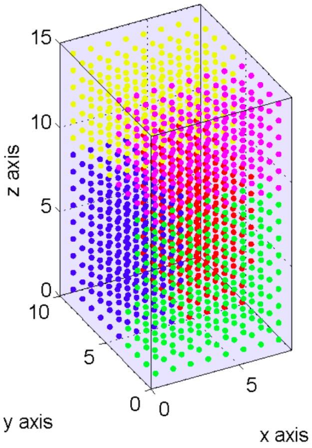

Figure 7: $8 \times 10 \times 15$ Structured Grid

Heuristic Partitioning Method Applied To Structured Grids

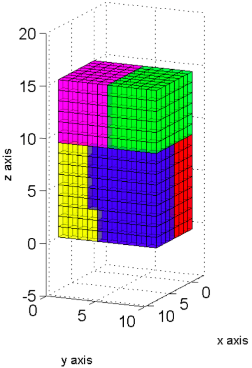

Figure 8: $8 \times 10 \times 15$ structured grid partitioned by our algorithm

Heuristic Partitioning Method Applied To Structured Grids

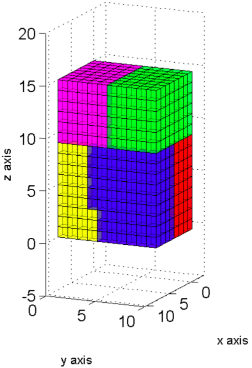

Figure 8: $8 \times 10 \times 15$ structured grid partitioned by our algorithm

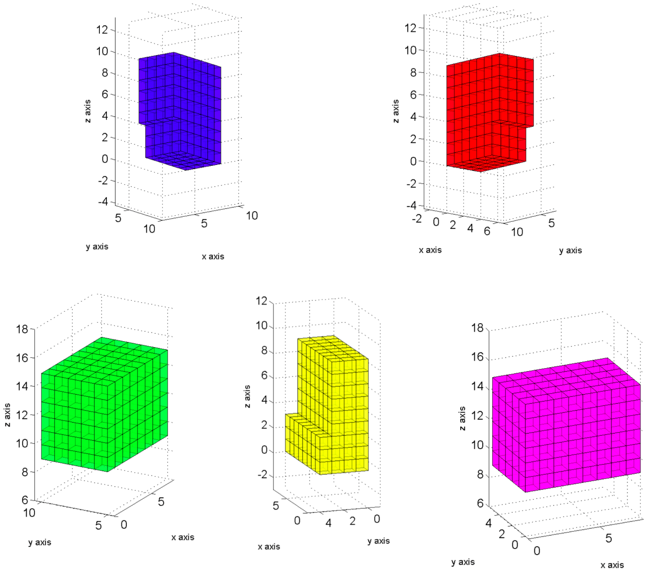

Figure 9: Partitions assigned to 5 processors

Heuristic Partitioning Method Applied To T-shape Grid

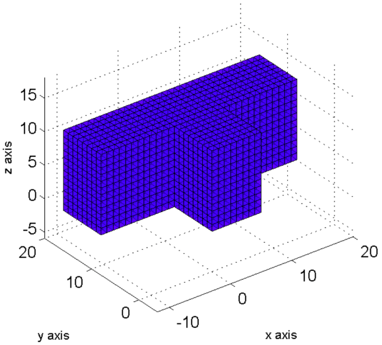

Heuristic Partitioning Method Applied To T-shape Grid

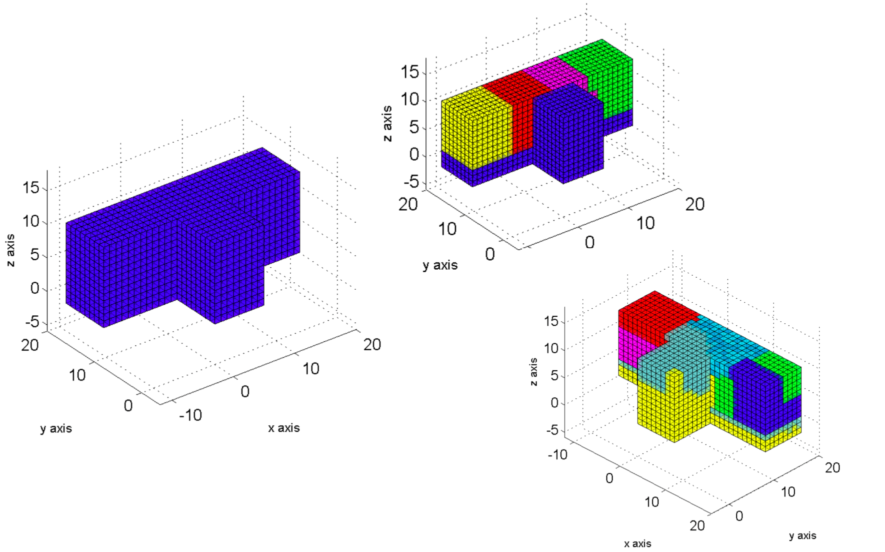

Figure 10: Structure partitioned to 5 and to 8 processors

Partitioning Problem Of The Block Graphs (2)

- Assign the work units to processors in order to minimize the sum of the maximum load and the maximum communication cost
- Mathematical formulation

minimize

$$\left(C_1 \max_{p \in P} \left\{ \sum_{v \in V} w_v y_{vp} \right\} + C_2 \max_p \left\{ \sum_{p^\star \neq p} \sum_{(v,u) \in E} w_{(v,u)} y_{vp} y_{up^\star} \right\} \right)$$

such that

1. $\sum_{p \in P} y_{vp} = 1, \forall v \in V$
2. $y_{vp} \in \{0, 1\}, \forall p \in P, v \in V$

Partitioning Problem Of The Block Graphs (2)

Here is an equivalent formulation.

$$\text{minimize } (C_1 T_1 + C_2 T_2)$$

such that

1. $\sum_{v \in V} w_v y_{vp} \leq T_1, \forall p \in P$
2. $\sum_{p^\star \neq p} \sum_{(v,u) \in E} w_{(v,u)} y_{vp} y_{up^\star} \leq T_2, \forall p \in P$
3. $\sum_{p \in P} y_{vp} = 1, \forall v \in V$
4. $y_{vp} \in \{0, 1\}, \forall p \in P, v \in V$

Partitioning Problem Of The Block Graphs (2)

We now linearize the term $y_{vp}y_{up^\star}$.

1. $$x_{vu}^{pp^\star} = \begin{cases} 1 & \text{if node } v \text{ (resp. } u\text{) belongs to cluster } p \text{ (resp. } p^\star\text{)} \\ 0 & \text{otherwise} \end{cases}$$

2. $$\begin{cases} x_{vu}^{pp^\star} \geq y_{vp} + y_{up^\star} - 1 \\ x_{vu}^{pp^\star} \leq y_{vp} \\ x_{vu}^{pp^\star} \leq y_{up^\star} \end{cases} \qquad \forall (u,v) \in E; p, p^\star \in P : p^\star \neq p$$

3. The minimization structure of the model implies that x will automatically be binary and the last two constraints are not necessary. We redefine the variables.

$$x_e^p = \begin{cases} 1 & \text{if e is an outgoing arc from cluster } p \\ 0 & \text{otherwise} \end{cases}$$

4. $x_{(v,u)}^p \geq y_{vp} + y_{up^\star} - 1, \forall (u,v) \in E; p, p^\star \in P, p^\star \neq p$
5. $x_{(v,u)}^p \geq 0, \forall (u,v) \in E, p \in P$

Partitioning Problem Of The Block Graphs (2)

Here is the final formulation, assuming that all vertex weights are identical and all edge weights are also identical.

$$\text{minimize } C_1 w_1 T_1 + C_2 w_2 T_2$$

such that

1. $\sum_{v \in V} y_{vp} \leq T_1, \forall p \in P$ and $\sum_{p \in P} y_{vp} = 1, \forall v \in V$
2. $\sum_{e \in E} x_e^p \leq T_2, \forall p \in P$
3. $x_{(v,u)}^p \geq y_{vp} + y_{up^\star} - 1, \forall (u, v) \in E; p, p^\star \in P, p^\star \neq p$
4. $y_{vp} \in \{0, 1\}, \forall p \in P, v \in V$
5. $x_{(v,u)}^p \geq 0, \forall (v, u) \in E; p \in P$

Partitioning Problem Of The Block Graphs (2)

Conclusions after implementing our algorithm.

- Our problem is a very challenging combinatorial problem.
  - It is highly degenerate.
  - The LP is very weak (the two halves of each node are in different clusters).
  - The network size is massive.
- We tried the following approaches.
  - Cplex
  - Benders decomposition
  - MIP-based heuristic (I, II)

If we know the values of the y variables, the communication cost is almost trivial to compute.

Pseudocode of MIP-based heuristic (II)

1. Find the degree of each node.
2. For each cluster (associated to a given processor), execute the following steps.
   1. Find the node with the smallest degree (i.e., v).
   2. While the load limit is not reached and some neighbour exists, do
      1. Find all neighbours of v and add them to the current list;
      2. Update v by selecting one of its neighbours for exploration.
   3. Fix the binary variables associated with each node in the cluster list.
3. Solve the "restricted MIP."

Partitioning Problem Of The Block Graphs (2)

Results (the time limit was set to 30 minutes)

Approach	Instance I	Instance II
Cplex	149951	1966080
Benders	150480	N/A
Heuristic I	162240	369280
Heuristic II	130880	291136

Reorganizing the Blocks Assigned to a Given Processor

We have considered two approaches.

1. Treat the union of all the blocks as a single structure and try to decompose it into “cuboids”, i.e., grid-like regions.
2. Attempt to merge the blocks themselves by applying some algorithm (probably a greedy algorithm).

The first approach requires very complicated and time-consuming algorithms and the second one does not guarantee that the solution produced is optimal.

General Algorithm To Build The Final Blocks

Algorithm 1 General Algorithm

```
mark all the blocks with the labels of the original block from which they were extracted
if all the blocks belong to the same original block then
    call function to get structured blocks
else if the original blocks are compatible (Call Routine to check for degeneracy) then
    merge all the blocks
    call function to build structured blocks
else
    divide the blocks in 3D representable sets (TODO)
    for all the 3D representable sets do
        call function to build structured blocks
    end for
end if

function BUILD STRUCTURED BLOCKS(3D representable blocks)
    call routine to detect concave edges (implemented)
    if there are no concave edges then
        all the blocks assigned to the processor form a structured block
        return the blocks structure
    else
        call routine to build structured blocks (TODO)
        return the blocks structure
    end if
end function
```

General Algorithm To Build The Final Blocks

Comparison between the original set of blocks and the final result where the concave edges are detected

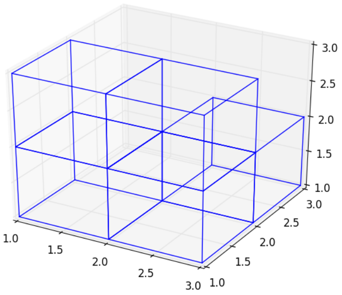

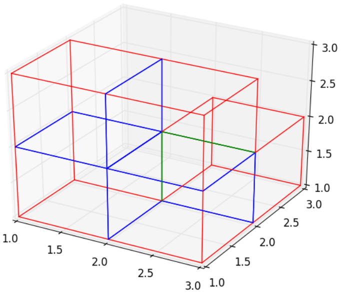

Let L denote the original list of cuboids. Repeat the following steps as long as the list contains at least one mergeable cuboid.

1. Choose a mergeable cuboid (let it call A).
2. Among the cuboids that can be merged with A, select a cuboid B.
3. Merge cuboids A and B, yielding a new cuboid C.
4. Replace cuboids A and B by cuboid C in the list L.

Future Directions

- Implement refinement and heuristic algorithm for a more complicated structure.
- Compare with a deterministic optimization technique.
- Study various clustering methods for parallel optimization.

Thank You For Your Attention!