Exascale Vision for Mesh Adaptivity

Joaquim Peiró, David Moxey, Spencer Sherwin, James Slaughter
Imperial College London and University of Exeter

With contributions from Michael Park (NASA), John Chawner (Cadence), and Bob Haimes (MIT)

ELEMENT

Current State of the Art

Increase/decrease resolution:

- ***h***
  - Mesh modification operators: Node Addition, Edge/Face/Cell Splitting, Collapsing & Edge-Face Swaps
  - Contained in Fluidity, LibMesh, Deal.II
  - Used for anisotropic geometric or physics related adaptation
- ***r***
  - Clustering of current nodes
  - Local feature capturing – crack propagation, shocks, boundary /shear layers
- ***p***
  - Increase/decrease order of polynomial expansion basis functions
  - Increase accuracy of smooth solutions
- ***Re-meshing***
  - Regenerate selected regions of the domain

Workshop Learnings – Exascale Bottlenecks

- *Meshing workflow will be linked to hardware choices at Exascale*
- *Disk I/O and storage limitations*
- *Adaptation algorithms and dynamic load balancing.*
- *CAD retention and query operations*
- *Error estimation: local/global (interpolation/adjoint)*

ELEMENT

Meshing Workflow Linked to Exascale Hardware

- ***Algorithm choices driven by hardware selection of exascale computers***
  - Partly by scalability concerns across heterogeneous architectures – eg Task-based execution for GPU/FPGAs
- ***Base Mesh Generation***
  - Too large for workstation hardware – may necessitate parallel meshing even for base mesh
- ***Mesh Quality***
  - What is a good mesh? – Are sizing functions 'quality aware'?. Anisotropy?
  - Mesh enhancement raises communications concerns yet needed to ensure mesh quality
- ***Robustness, Automation, and Fault Tolerance***
  - Simulation scale will prohibit manual interaction with processes

Disk I/O and Storage

- ***Space requirements and limitations for adapted meshes***
  - Disk I/O needs to be limited and size of these meshes will be restrictive
    - Could only retain meta-data and reconstruct for post-processing and visualization though this presents robustness challenges
    - Exacerbated with highly dynamic/moving meshes

ELEMENT

Exascale Scalability of Meshing Algorithms

- ***Initial Graph Partitioning***
  - If partitioning using the full heterogeneous architecture issues with:
    - Memory write synchronization over thousands of GPU cores
    - SIMD GPU architectures and anisotropy of initial graph cause bottlenecks
    - Smaller VRAM vs System Memory causing access and transfer/communications considerations
    - Use of both SIMD and MIMD architectures
- ***Adaptation Algorithms***
  - Redesign of existing algorithms to reduce parallel communication
    - Focus on locality
- ***Dynamic Load Balancing***
  - Likely to consume a large portion of time at exascale
    - User-specified load-balancing frequency?

ELEMENT

CAD Retention and Query Operations

- ***Retention of Geometry definition***
  - Surface values and gradients necessary for certain error estimation techniques.
- ***Need for Geometry Representations designed for Parallel Distribution***
  - With notable exceptions, most current representations are not.

ELEMENT

Short-term Research Agenda

- ***Identification of current algorithms for both meshing and partitioning/dynamic load balancing for use on heterogeneous exascale architectures***
    - Identification of promising existing algorithms and techniques suited for chosen exascale hardware

  Eg, multi-level partitioning methods hard to scale and spectral methods (eigenvalue based) being computationally expensive.

- ***Improvement of CAD APIs for Exascale***

ELEMENT

Medium-term Research Agenda

- ***Improvement of Identified Algorithms***
  - Improvement of previously identified space-time (4D) adaptation algorithms for exascale workflows – i.e. task-based execution for GPU/FPGA hardware.

- ***Lightweight Geometry Definitions and APIs***
  - Reduction in size of CAD APIs for exascale use.
    - CAD systems could be very large – yet only a few functions are used in geometrical queries
  - Lightweight Geometry definition to ensure easy storage and distribution.

Long-term Research Agenda

- ***Methods of Integrating Combined hrp 4D Adaptation***
  - CPU time vs efficiency considerations when considering application of each individual technique.
  - Combination of multiple strategies likely needed for complex problems at the exascale.
  - Techniques to synergize* the application of adaptation strategies are needed.

* *The whole is more than the sum of its parts.*