BDEC Math, Algorithms and Libraries

MICHAEL A. HEROUX, SANDIA NATIONAL LABS
ALISON KENNEDY, EPCC
CONTRIBUTIONS FROM MANY OTHERS

BIG DATA AND EXTREME-SCALE COMPUTING

Bridging HPC-HPDA Computing Environment Gaps

HPC and HPDA have separate computing environment heritages.

- Data: R, Python, Hadoop, MAHOUT, MLLIB, SPARK
- HPC: Fortran, C, C++, BLAS, LAPACK, HSL, PETSc, Trilinos.

Determine capabilities, requirements (application, system, user), opportunities and gaps for:

- Leveraging HPC library capabilities in HPDA (e.g., scalable solvers).
- Providing algorithms in native HPDA environments.
- Providing HPC apps, libraries as appliances (containers aaS).

Note: HPDA = BD, HPC = EC

Refactoring & leveraging of HPC Capabilities for HPDA

Sparse computations:

- HPC: low, consistent degree graphs.
- HPDA: highly variable degree, "power law" graphs.
- Requires different partitioning, parallel strategies.

Dense LA for some machine learning.

High performance communication libraries (MPI).

- Global collectives for machine learning (dense).
- Point-to-point for graphs.

New Math & Algorithms

Math & Algorithms for Intrinsically Discrete Data (Ie.g., light sources)

- Model extraction.
- Surrogate development.
- Inverse problems.
- In general: Converting observations to models.
- Mature in HPC (e.g., Oil & Gas), but new areas: e.g., sensors.

Factorizations, spectral algorithms, other NA for tensors.

Algorithms based on random sampling.

- Stochastic Gradient Descent algorithms from sampling.
  - Already being done, but reconsider from HPC perspective.
- Better methods than gradient descent?

Streaming algorithms, “online” algorithms.

Complexity reduction: Decrease from n^2 to $n \log n$ or n.

- Similar to multi-pole expansion, FMM.

Low-rank representations: e.g. H-matrix approaches.

General: Revisit HPDA problems with mindset of “HPC is in your toolbox.”

New Libraries

HPC-HPDA libraries are needed.

- Scalable. Not trivial for many reasons.
- Support virtual resources (e.g. virtual clusters).
- Agreed upon abstractions.
    - Graph, KV, pixel ?
    - File formats (HDF5, FITS): Reconcile common data/file formats with big data.
- Usability, accessibility: “Bring to the HPDA community”
    - Address multiple situations from long tail to big science.
- Conceptual software stack.
    - Low-level services to high-level knowledge.

Requirements for other BDEC Focus areas

A well defined infrastructure (virtual cluster concept):

- Important for providing libraries.
- It's a good model in general.
- Must be high performance.

High performance virtual network APIs.

- Infiniband is fast, need virtual, fast API.

Programming model & communication layers:

- Bring together the best of HPC and HPDA.
- Examples: MPI+Hadoop/Spark, Load balancing + Giraph/Pregel

Support for workflow, data fusion.

- E.g., Drawing from multiple data sources.

Containerized Libraries

HPDA seems to require container approaches.

HPC can provide, and can benefit. Containers becoming new tarfile format.

Eliminates complicated build process.

- The dozens of 3rd party libraries can be pre-built, installed.

Full-featured installation, custom installations.

- Example: Trilinos has 57 packages.
- Most people use “aspects” of Trilinos.

Portable debugging environments.

Many other opportunities.

Lots to be learned…

NERSC Cori Docker effort: “Shifter”.

Opportunities & Challenges

Funding opportunity announcements: Numerical Analysis for HPDA.

- Appears to be rich exploration space for expansion to tensors.

New math models for discrete data sources.

Leverage of HPC sparse technologies in HPDA.

Containerization technologies from HPDA for HPDA & HPC.

- Turnkey of common computations, e.g., SVD.

A lot of domain knowledge building.

- HPDA and HPC conversations.
- Translation of terminology.
- Domain model understanding.