Applications Breakout Report
Big Data and Extreme Computing
Barcelona
29-30 January 2015

James Ahrens
Jean-Claude André
Rosa Badia
Thierry Bidot
Mathis Bode
Fernando Cucchietti
Anshu Dubey
Françoise Genova
David Keyes
Jysoo Lee
Takemasa Miyoshi
Morris Riedel
Alejandro Ribés
Joel Saltz
Alex Szalay
Bill Tang
Akira Ukawa
Pier Luigi Vidale

Comparing “HPC” and “BDA”

- **What are the main differences and commonalities between the HPC and BDA requirements/ technologies/working-assumptions in this area?**
- **We prefer to reject the given labels, considering instead “numerically intensive” (NI) and “data intensive” (DI), both of which can avail themselves of HPC. (These are also imperfect; the right labels is a “work in progress.”)**
- **Both NI and DI approaches share the common challenge of gaining scientific insights, making prediction, and quantifying uncertainty – the former through first principles models and the latter through statistical models.**

“We shape our buildings; thereafter they shape us.”

Winston Churchill, 1943

Comparing “NI” and “DI”

- **How to distinguish?**
  - “Amdahl number” (Gray & Szalay), the ratio of bits through the CPU to cycle/s: O(10^-3) for “HPC” and O(1) “BDA”
  - Data access pattern is a useful distinguishing feature

Significant overlap in the spatio-temporal data realm

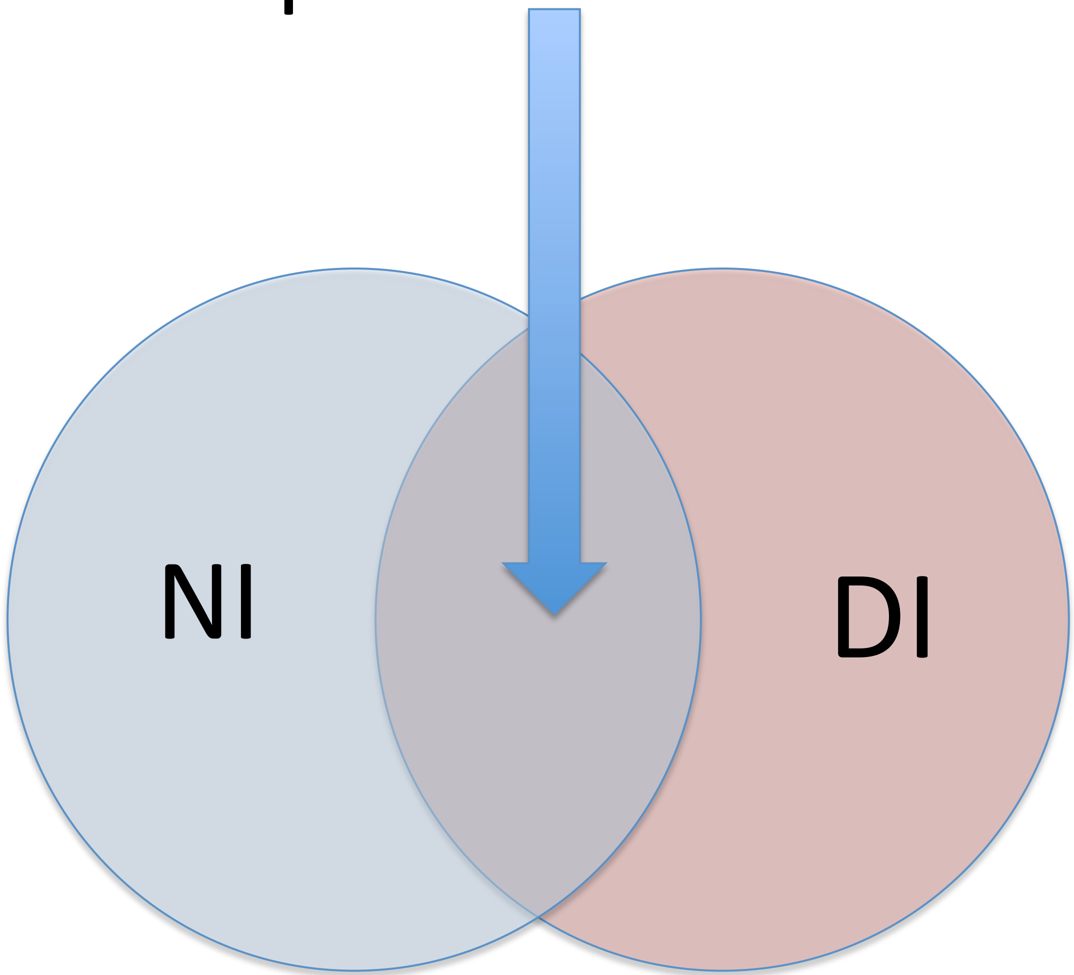

Comparing “NI” and “DI”

- **What are the main differences and commonalities between the HPC and BDA requirements/ technologies/working-assumptions in this area?**
- IDC classifications (2011) into which either may fall:
  - Traditional applications with growing data sets
  - “Needle in haystack” applications (low information density)
- HPC has always been oriented towards causality (dynamics, hypothesis-driven), whereas BDA has traditionally been oriented towards archiving data for future discovery
  - “The most important thing about big data is that you can go back to it.” (Amazon)

Comparing NI and DI

- **What are the main differences and commonalities between the NI and DI requirements/technologies/working-assumptions in this area?**
- NI data, arising from continuous models on meshes or swarms, tends to be highly structured and highly correlated, which may not be true of DI data
  - Example: correlation of neighbors' backyard temperature vs. Twitter behavior
- NI data tends to do lots of writes, whereas DI tends to write once and read many times.

Comparing NI and DI

- **What are the main differences and commonalities between the NI and DI requirements/technologies/working-assumptions in this area?**
- There is a growing overlap between NI and DI in applications. Many NI applications produce Big Data and DI is a growing consumer of HPC capabilities.
- There is a growing opportunity for common approaches, common software frameworks, and even reuse of data structures.

Comparing NI and DI

- **What are the main differences and commonalities between the NI and DI requirements/technologies/ working-assumptions in this area?**.
- DI has a tendency to use machines with larger memory spaces and I/O infrastructure more oriented towards random access, instead of the bulk streaming writes that dominate NI usage. Still, a growing degree of parallelism with asynchronous, distributed computation is a trend in applications from both sides.
- In NI applications the main action occurs in memory and I/O tends to be regular, whereas in DI the I/O stream is much more complex.

Axis	Sub-axis	Numerically Intensive	Data Intensive
Hardware	Nodes and Interconnect	High performance and power	Lower performance and power
	Storage	Separate, independent	Integrated
SW	Synchronization	Tightly coupled	Loosely coupled
	Reliability	Checkpoint restart	Replication
Workload	Number of Users	**Single per node**	**Multiple per node**
	Data	Dynamic, heterogeneous (unstructured grid)	Static, homogeneous (text, images)
	Algorithms	Global	Distributed
	User Interface	**Complex Application**	**Simple Web**
	Data Model	**Files**	**Database**
Workflow	Scheduling	Batch	Interactive
	Analysis	Offline post-processing	Online
	I/O	Bulk parallel writes	Streaming writes

c/o Ahrens, 2012

Commonalities

- **Are there common needs/problems/interfaces could serve as the basis (or as stepping stones) along a path to (some reasonable level of) infrastructure and application convergence?**
- To some degree, convergence has been dictated by the hardware trends and what vendors make available. There is less convergence in software.
- Under all hardware scenarios, data movement is becoming relatively increasingly expensive and analytics should computed *in situ,* or as close as possible to the data source. This may dictate changes in the store vs. recompute spectrum.

Testbeds for convergence

- **Are there interdomain testbeds that combine "BDA" and "HPC" workflows in ways that could help uncover pathways toward convergence?**
- Yes: a joint initiative between academia and industry to investigate applicability of high end "Google-type" platforms to scientific, medical, and engineering problems
  - e.g., Google Earth collaboration with European Geoscience Users (https://sites.google.com/a/earthoutreach.org/google-egu-2014/)

Game changers?

- **What is/are the technology or new research that may be a game changer?**
- Combination of embedded processing components, 3D memory technology, and out-of-core memory algorithms exploiting burst buffers have the potential of enhancing the capability of massive data processing while in transit and therefore enable many of the computations needed to extract knowledge in both domains.
- Systematic assessment is needed to determine in different domains what are the true requirements in terms of actionable information needed by users/ scientists together with quality compromises that can give rise to approximate computation that may be improved progressively.

Game changers?

- **What is/are the technology or new research that may be a game changer?**
- There is a need to investigate more deeply the potential for models that get away from classical batch processing and take more into account how the information is rendered accessible to decision makers (including e.g. scientists running multiple experiments on expensive equipment and therefore needing to make hard choices/difficult trade-offs). This includes measuring cost in all forms (time, memory, acquisition, power, user productivity, etc.).

Highest priority for “convergence”?

- **What action would be your number one priority to be taken rapidly to ensure success of the convergence of Extreme computing and Big Data infrastructures?**
- Build open data repositories and develop data challenges that highlight the unification of the techniques, ranging from data services to facilitate access, down to general data query and processing for investigation and exploration (ultimately knowledge discovery).
- We also need conferences and workshops that foster a unification of the research communities that are now highly fragmented, as evidenced by several presentations at this workshop.

Highest priority for “convergence”?

- **What action would be your number one priority to be taken rapidly to ensure success of the convergence of Extreme computing and Big Data infrastructures?**
- Allocation processes for major facilities are currently adapted to the needs of a typical simulation user (e.g., batch, temporary use of storage), and needs to become friendlier for data intensive workflows.

Highest priority for “convergence”?

- **What action would be your number one priority to be taken rapidly to ensure success of the convergence of Extreme computing and Big Data infrastructures?**
- “HPC” is driven by prioritizing efficient use of CPU cycles.
- “BDA” is driven by prioritizing the data flow and data parallelism.
- “Convergence” is driven by prioritizing the huge data sets that are emerging, structured or unstructured, and emphasizing use of commodity hardware and software in extracting the scientific benefits.

Highest priority for emergence?

- **What action would be your number one priority to be taken rapidly to ensure the emergence of efficient Extreme computing and Big Data applications?**
- Enable it(!) through the provision of large-scale resources that are not regimented into the scheduling algorithms of traditional HPC and allow the persistent storage of relevant data bases.
- Algorithmic research is needed to achieve scaling of parallel techniques especially on the Big Data front, which does not have a long tradition of parallel computing compared to HPC.

Highest priority for emergence?

- **What action would be your number one priority to be taken rapidly to ensure the emergence of efficient Extreme computing and Big Data applications?**
- Hierarchical representations of data that instead of storing all data at finest resolution stores representative sets at finest resolution, along with recursively space and time coarsened sets, and statistics.
- New data structures based, e.g., on space filling curves may be efficient in accessing data at adaptively increasing scales, with coarser scales useful for transmission and low-res visualization and finer scales useful for analysis once the phenomena of interest are identified.

Highest priority for emergence?

- Portability
  - Performance
  - Scheduling
- Usability
- *In situ* use of the data
  - On the fly extraction
- Means of tracking provenance through data transformations

Caveats

- We should not force a “shotgun” marriage of “convergence” when a love-based marriage is inevitable in the near future.
- The ultimate test of our efforts will be whether the emergence of NI+DI applications meets mission-critical needs in scientific discovery and engineering design.
- We need a follow-up meeting with additional input from extreme computing big data people.

Yesterday’s slides

Role of Applications Breakout

- Define the scientific customers for exascale
  - Identify the ready, eager communities
  - Raise expectations for computational scientists and engineers who are “satisfied” with current scale
  - Discover new applications with new capabilities
- Interface outward to stakeholders
  - Importance of predictive science, whether from first principles or statistical tools
- Interface inward to algorithm developers
- Ensure scientific accountability
  - Validation and verification improved at the new scale and integrated with observation
  - What did we do with the last nine orders of magnitude? (importance of real data)

What do we add to the discussion?

- Call attention beyond the Google-type problems already optimized by commercial applications
- In scientific applications, exploratory searches are followed by confirmatory simulations
  - Correlation **and** causation

Application types (1/2)

- Third paradigm
  - PDE-based models
  - Particle-based models
  - Linear algebra-based models (e.g., DFT)
  - Image processing
- Fourth paradigm
  - Archiving and retrieving from massive data sets
  - Clustering
  - Searching
  - Knowledge discovery

Application types (2/2)

- Combinations of Third and Fourth paradigms
  - Fourth informs Third
    - Inverse problems
    - Data assimilation
    - Visualization and computational steering
  - Third informs Fourth
    - Design of experiments
  - Both paradigms in a virtuous loop

Algorithmic Requirements

- Pipeline between simulation and analytics
  - Common data structures
  - *In situ* use of the data
- More efficient, scalable tools for the "inner" problems
  - Adaptive, scalable solvers
  - SVD, SVM
  - Search algorithms
  - Graph algorithms
- Efficient I/O and strategic checkpointing for ever-increasing data sets
- Combination of the usual "-omics" with spatial dimensions

Requirements on the SW/HW environment

- Portability (and performance portability)
- Usability
- *In situ* use of the data
  - On the fly extraction
- Adequacy of bulk network transfers for global research community acting on curated data sets
- Means of tracking provenance through data transformations

In/adequacy of existing national strategies and plans (1/3)

- Value-chain orientation focuses "big data" on the MapReduce-type problems
  - May miss some of the major scientific opportunities by being driven by the commercial opportunities
  - Commercial engineering drivers are growing, but not yet as exascale-intensive as the science drivers
- Big data and HPC are disconnected at the science policy and industry level in many countries
  - Big data is oriented towards social/business data
  - HPC is oriented towards scientific challenges
  - Potential of machine learning is underexploited in science and engineering

In/adequacy of existing national strategies and plans (2/3)

- Some calls are worded in ways that make scientists wary of bridging from HPC to big data
- Observational campaigns are sometimes inadequate (e.g., in resolution) to fully exploit or provide checks on simulational campaigns
  - And *vice versa*
  - Missed opportunities from lack of coordination or lack of synchronization in time or between nations

In/adequacy of existing national strategies and plans (3/3)

- Some calls are too narrowly posed to dynamically respond to scientific opportunities
- Funding for sustainability – code development and data curation – is rare compared to funding for research
- What lessons from the co-design centers for HPC can influence the big data campaigns, e.g., the EU Centers of Excellence

Comments to kick back to plenary

- Want to save full-resolution simulations, since the newest phenomena are in the smallest scales
- Cannot current cache the data stream

Questions to kick back to plenary

- Will the exascale machine look just like the petascale machines?
- Will the "killer apps" be something completely different from what we are planning on today?

Possible Case Studies emphasizing time to solution

- Complex simulations
- Merging of simulation with sensors
- Real-time aspects
- Examples
  - Weather forecasts
  - Wildfires
  - Earthquakes
  - Hazardous release evacuation
  - Reactor failure
  - Radiation treatment planning

Profile of the group

- Nations represented
  - EU: France, Germany, Iceland, Spain, UK
  - Asia: Japan, Korea
  - USA
- Disciplines represented
  - PDE-based: combustion, climate/meteorology, fusion
  - Observation-based: astronomy, biomedical informatics
  - Tools-based: PDE frameworks, runtime systems

EOF

Questions to address

- What are the national strategies and plans?
- What is missing from the respective plans?
- What recommendations can we make for an internationally coordinated converged plan for HPC and Big Data