Data Cross Cutting: Who

- Stephane Ploix
- Ricardo Brunino
- Giovanni Aliosio
- Sandre Fiore
- Modesto Orozoco
- Bryan Lawrence
  - (Rapportuer, bouquets for ideas to the others, brickbats for misrepresentation to me.)

Data, Context: New Paradigm

- New HPC usage more and more about direct numerical simulation of the real world, expect that to be more prevalent at exascale.
  - Complex assemblies of the real world, e.g. the brain, the climate, solid earth etc.
  - Large data problems; often I/O bound now, unable to use all the FLOPS available now.
  - Significant data analysis and visualisation problems.
- Data Analysis has become an HPC problem
  - Genomics: millions of individuals each with a 100 GB genome. Distirbuted storage. Distributed analysis (but still an HPC problem per se).
  - We can already envisage a sensible climate problem that can produce EB of data from a foreseeable computing system. Think of an exascale not as a “computer” (not HP*C*), more a data source: HPD! (high performance data)

Storage Problem

- Expected to be relatively tracticable to hold data
- The issues are the act of storage, and the act of retrieval.
- Bandwidth problems; from the machine to storage, between storage, and more widely.
  - Not all data can be stored?
- Expect a mult-tiered “cache” storage future, local, national, European, wider …
- All leads to distributed access and distributed processing
- Storage problem is already primarily a software problem, and will be worse at exascale.

Solutions? 1: Abstraction & Metadata

- Change the way we think about I/O
- New layer between solvers and post-processing, new “component” to manage I/O and “on-the-fly” processing.
  - Store less primary data; sub-sample ensemble PDFs on the fly and only save representative members, rather than all members.
  - “Whatever you can do on the fly is good for you”, but
  - Abstraction which irrevocably hides where the data is, may cause latency problems.
- Provide more metadata at run-time, to support more targeted I/O in analysis phase.
- Better workflow management (concentrating on data aspects).

Solutions, 2: Parallelisation

- Need scalable algorithms for statistical analysis.
- Stop talking about “post” processing: Post-processing is just processing which is I/O bound.
- Many analysis codes are single-threaded (even some which can be distributed).
- Invest in improving performance of analysis codes as much as primary solver codes.

Solutions, 3: Distribution

- Need better s/w tools for indexing, searching, and querying massive datasets.
  - An exascale computer is useless to many communities if the product of the calculation cannot be consumed!
- Expect large discipline specific cache archives on the European scale, analysis done close to those caches.
- Enable remote processing services, both by defined services, and by allowing “users” to move computations to the data.
  - (Software as a service, platform as a service, both adjacent to data).

Solutions, 4: Standardisation

- Much of the existing data handling is done by talented amateurs (from a s/w engineering perspective), who spend the minimum time possible, developing adhoc scripts.
  - Repetitive. Fragile. Inefficient. Possibly even wrong!
  - Real time pressures (improved data handling doesn't produce Nature papers).
  - But investment yields no returns in a non-standard environment.
- From too many formats to standardised environments?
  - Metadata to support search and query
  - Standard interchange formats.
  - Standard storage formats.

Solutions, 5: New Technologies

- Strongly urge investment in parallel file systems
  - Need to support massivel parallel I/O read and write that will scale to exascale.
  - Enhance with abiliity for WAN interfaces (cf IRODS)
- Active Storage
  - Computation “within” the storage, to avoid data movement (e.g. correlation?)
  - Active research area, but vendors are very “shy” about tellng us what it can/will do! An opportunity for “co-design” being missed?

Politics

- The exascale data handling problem is not just about the lack of s/w, it's also about sustained s/w investment.
- European initiatives are often not sustained long enough to be competitive with other (American, and probably soon, Chinese) offerings.
- Data life cycle is expected to be long. Analysis s/w will need to have longevity. 3-5 year funding life cycle is not representative of the reality of big data handling.
- Need to ask the question as to whether European funding can be better spent enhancing existing (foreign) s/w (which has sustained investment) rather than buidling products aimed at being competitive (but without sustained investmentment).
  - Invest in collaboration, rather than competition may yield better results?