

WG 4.3 activity

WG 4.3: Numerical libraries, solvers, and algorithms

- Iain Duff, Chair
- Andreas Grothey, Vice-Chair

Composition of Working Group 4.3 of EESI

Iain Duff	STFC/CERFACS	UK	Sparse Linear Algebra
Andreas Grothey	University of Edinburgh	UK	Cont & Stoch Optimization
Patrick Amestoy	ENSEEIH-IRIT, Toulouse	FR	Sparse Direct Methods, Solvers
Peter Arbenz	ETH Zürich	CH	Eigenvalues, HPC
Jack Dongarra	Tennessee/Manchester	UK/US	HPC, Numerical LA
Salvatore Filippone	Università di Roma	IT	Numerical Software
Mike Giles	University of Oxford	UK	GPU, CFD/Finance
Luc Giraud	INRIA Bordeaux	FR	Iterative & Hybrid Methods
Thorsten Koch	Zuse-Institut Berlin	DE	Combinatorial Optimization
Bo Kågström	Umeå University	SE	HPC, Dense Linear Algebra
Karl Meerbergen	K.U. Leuven	BE	Preconditioners, ExaScience Lab
Volker Mehrmann	TU Berlin	DE	Linear Algebra, HP Applications
Gerard Meurant	ex-CEA	FR	HPC, PDE solution
François Pellegrini	Université de Bordeaux & INRIA	FR	Partitioning
Julius Žilinskas	Vilnius University	LT	Global Opt, Meta-heuristics



General comments on area covered

- An **enabling technology**
- Motivated by **needs of applications** and critical to their success
- Much interdependence between areas
- Hierarchically structured
- Importance of **also working at Peta and Terascale level**
- Will illustrate main issues and some solutions
- Also highlight **European involvement** and support required



Main areas covered in our report

- Dense linear algebra
- Graph and hypergraph partitioning
- Sparse direct methods
- Iterative methods for sparse matrices
- Eigenvalue problems, model reduction
- Optimization
- Control of complex systems
- Structured and unstructured grids



Software stack

The use of libraries will be **hierarchical**, for example, from bottom up, we may have:

- BLAS (eg dense matrix-matrix multiplication)
- Dense linear algebra
- Sparse solver
- Hybrid solver
- Optimization/Eigenvalues/Control

Another issue is the relationship of the solver to the application. How **specific is it to the application** and are we concerned by **special purpose hardware**.



Algorithmic issues

- Wide range of algorithms depending on functionality, problem, **and** problem structure
- Hierarchical**/Multilevel with several levels of parallelism
- Further levels of parallelism through the software stack
- Reducing synchronisation, say through use of directed acyclic graphs
- Blocking and remapping to **avoid data movement**
- Communication avoiding/hiding algorithms
- Development of **energy aware** algorithms
- Dynamic scheduling and load balancing
- Trade-off between speed and accuracy
- Provision of backward error estimates



Software issues

- **Interoperability.** Library software to be efficient over multiple domains, across machines and over time
- **Support action.** Need for credible support structure, constant adaptation of libraries to hardware (r)evolution and user needs
- **Several decades lifetime.** Needed for credibility of software and confidence of users
- **Open source with possible commercial exploitation**
 - Widest dissemination of the software libraries and tools
 - Leads to further collaborations and enhancement,
 - Preserve possible routes for exploitation (eg. LGPL licence).
- Ease of use
- More complex usage. **Inverse problems, stochastic programming**



Generic issues

- **Floating-point issues**
 - Accuracy
 - Reproducibility
 - Mixed-precision arithmetic

- **Programmability**
 - Higher level interface for applications people.
 - Support for multilevel parallelism. MPI/OpenMP ++
 - Further developments of MPI
 - Hardware support for mixed/quad precision?
 - UPC, co-array Fortran, and similar languages (**PGAS**)

- **Performance optimisation**
 - Autotuning
 - Software tools for code generation
 - Hardware monitoring and control



Fault tolerance

- A major concern
- Many algorithms have some inherent fault tolerance
- Difference between **MTBF** and **MTTI**
- Need as much help as possible from low level OS/software/hware
- Need a good FT-MPI
- Checkpointing** using FLASH memory
- Recomputation and redundant computation**
- Effect on accuracy
- Effect on performance (for example, cost of checkpointing)



Uncertainty quantification

- Occurs at several levels
 - Input data
 - Modelling of physical phenomena
 - Uncertainty in observed data
 - Approximation of continuous by discrete model
 - Solution of resulting equations
 - The effect of finite-precision arithmetic

- Algorithmic tools
 - Mixed-precision arithmetic and refinement
 - Backward error and software for assessing accuracy



Training

- For the **skills necessary** to develop the underlying mathematics and numerical algorithms and software
- Training in using programming models
- Training in using basic tools eg partitioning
- To alert users to **benefits** of using library software
- In the use of tools and library software that will be developed
- Relation to DEISA, PRACE, and HPC-Europa



European Expertise

- Many existing libraries either developed in Europe or **significant European input** to US-led/international projects
- Often few **isolated** centres within Europe, more coordinated efforts would be good
- Good **involvement in collaborations** with international/US researchers



Main concerns

- **Complexity** of moving to Exascale domain in terms of parallelism, heterogeneity, complexity of hierarchical structure
- Maintenance of **critical mass**
- Relatively **poor networking** of groups. Fragmentation of effort.
- **Lack of long-term funding** to support the maintenance of software libraries, including porting them to new hardware platforms



Trilinos

Contrast current European support with US projects like **Trilinos**.

- it is a project based on a framework that **federates multiple independent subprojects** covering many subject areas
- it pays close attention to **interoperability**
- it attempts to provide maintenance for **long term sustainability**



European Support Needed

- Funding model that secures **long term funding** for projects to ensure acceptance by user base, maintenance, porting to new architectures
- PhD, Postdoc funding to enable significant development/porting work
- Funding for specific research challenges (**targeted calls**)
- Provision of Exascale hardware as development platform
- Support for **networking and training**
- Support for **Exascale centre**



Quantification of support

What we need

- Around 2000 person years of effort
- a little over 10 million euros per year for the 2012-2030 timeframe
- A co-design centre at a cost of roughly 2 million euros a year for the centre and the same for an extensive visitor programme



Putting funding in context

Note that this funding request is **small** relative to costs of Exascale hardware; perhaps \$200 million with around \$20 million per year in electricity costs!

Cost of not doing it is **high**



THANK YOU FOR YOUR ATTENTION

