The EFDA HPC project

Roman Hatzky and Sibylle Günter
Max-Planck-Institut für Plasmaphysik

HPC-FF

EFDA ITM-TF Meeting, Lisbon, 13th - 15th September, 2010

High-performance computing in fusion needed for …

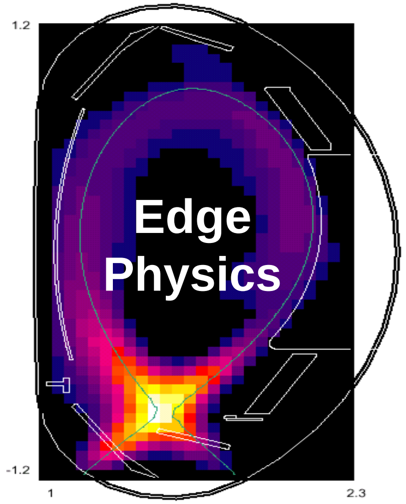

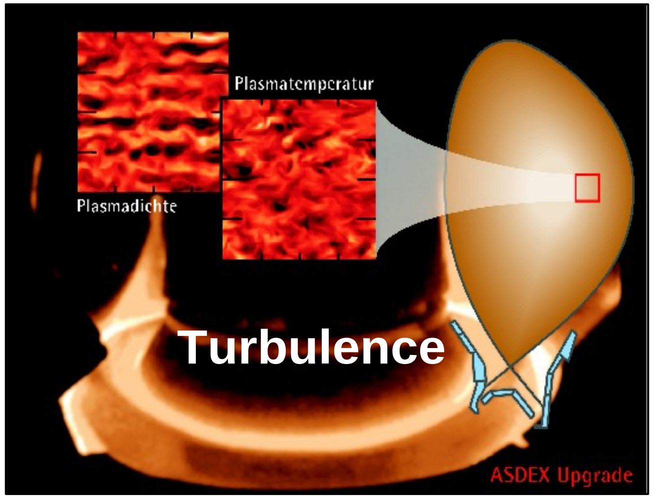

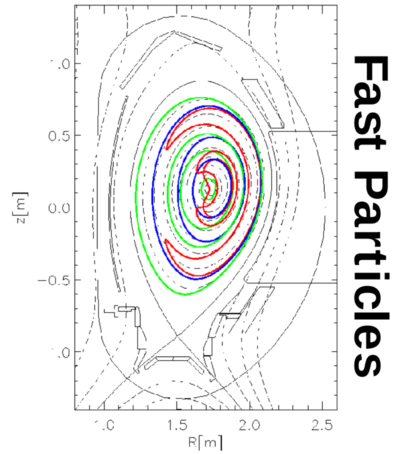

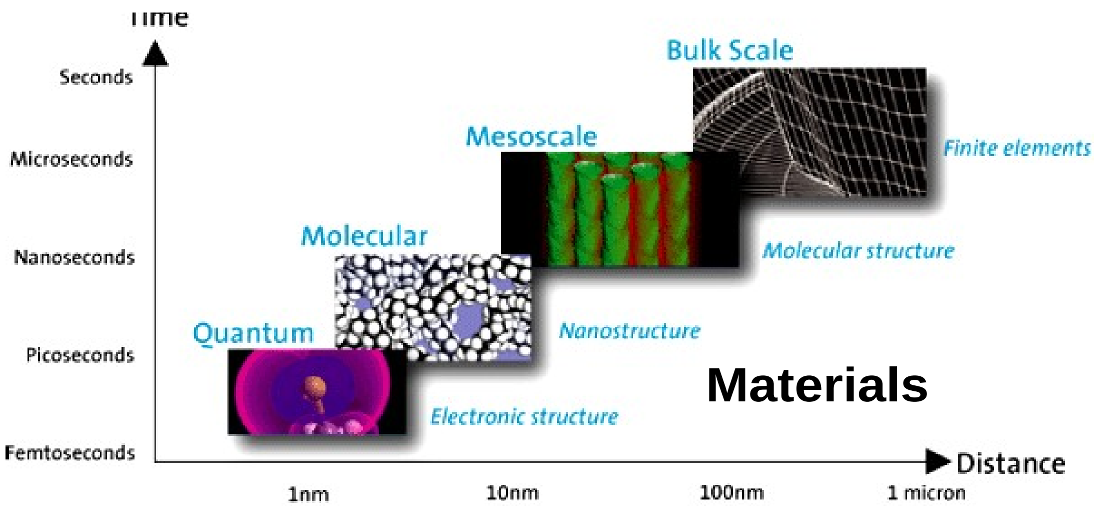

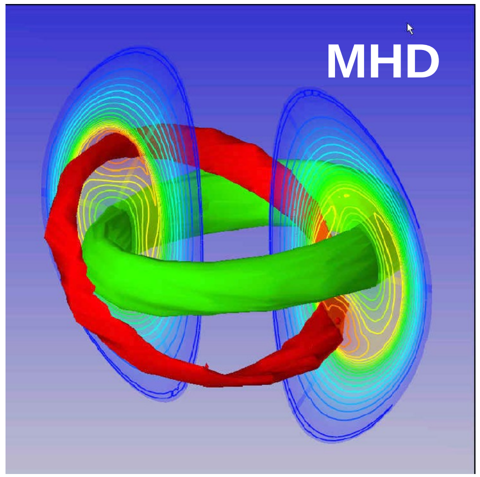

The EFDA HPC project

HPC-FF

- 1080 nodes @ 8 cores, peak performance 101 TFlop/s
- Start of production: August 2009

- numerical algorithms
- efficient parallelization strategies
- visualization
- education

core team (5 persons)

+ 4 ppy in other Associations

HPC Board

- Chooses HLST team members
- Monitors operation and exploitation of facility
- Allocation of resources (CPU time and high level support)
- User representation
- Annual work program for STAC approval
- ….

The board consists of nine members including:
S. Günter (IPP; chair person), R. Zagórski (EFDA; secretary)

HPC-FF at FZJ

Jülich's General Purpose Supercomputer

- 2208 nodes @ 8 cores
- 24 GB per node
- Intel Nehalem-EP
- Network: Infiniband QDR
- Peak Performance 207 TFlop/s

Supercomputer for
Fusion Science HPC-FF

- 1080 nodes @ 8 cores
  of same architecture
- Peak Performance 101 TFlop/s

Flexible swapping of resources

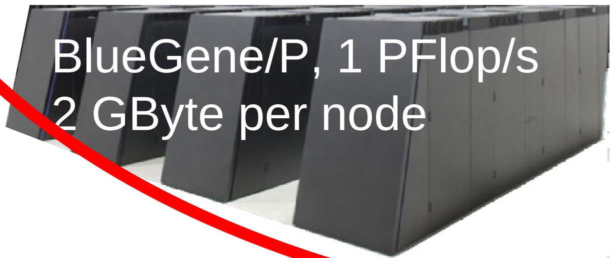

Storage
Environment

JuRoPA at FZJ

Whole computer Juropa-JSC/HPC-FF after installation:

No. 2 in Europe; No. 10 in Top 500 list: 308 TFlop/s Peak

High Level Support Team

HLST core team at IPP Garching

- Roman Hatzky - core team leader
- Matthieu Haefele (visualization specialist)
- Nicolay J. Hammer (astrophysicist with HPC experience)
- Nitya Hariharan (HPC specialist)
- Kab Seok Kang (numerical mathematician)

≈30 applications were evaluated to select the members of the core team

Quite difficult to find candidates with extensive
HPC & physics background

High Level Support Team

HLST staff members

- Antonio Gómez Iglesias (CIEMAT, Spain)
- Salomon Janhunen (TEKES, Finland)
- Peter Knight (CCFE, Great Britain)
- Huw Leggate (DCU, Ireland)
- Christophe Ortiz (CIEMAT, Spain)
- Florent Surbier (CEA, France)
- Trach-Minh Tran (CRPP, Switzerland)

Staff members are located at fusion relevant sites all over Europe

High Level Support Team

Management of HLST

- The work is coordinated by the HLST coordinator, Roman Zagórski.
- The core team leader, Roman Hatzky, guides the daily work.
- The whole HLST meets personally twice a year in Garching.

The HLST has its own web site URL: www.efda-hlst.eu for dissemination purpose. An internal domain offers special services for HLST members:

- A wiki to manage the knowledge base of HLST
- A version control system as repository for the codes of the projects

High Level Support Team

Main tasks for HLST

The HLST team is a support unit to ensure optimal exploitation of HPC-FF, i.e. it is not focused on its own academic research.

Support for code development

- Single processor performance optimization
- Parallelization & optimization of codes for massively parallel computers
- Improvement of the parallel scalability of existing codes already ported to parallel platforms
- Implementation of algorithms and mathematical library routines to improve the efficiency of codes
- Visualization of large data sets

Single processor performance optimization

Try to come as close as possible to the peak performance of the processor.

Performance = Frequency × Work/Instruction × Instructions/Cycle

- Frequency: further doubling the clock frequency can cause power consumption to increase by a factor of six or more (→ power crisis)
- Work/Instruction is given by the CPU instruction set; some compiler & developer influence through choice of algorithm and instructions
- Instructions/Cycles is architecture dependent; strong compiler & developer influence through optimization and usage of hardware vendor optimized numerical libraries, e.g. BLAS, LAPACK, Intel MKL, FFTW etc.

High Level Support Team

Multiple processor performance optimization

- Parallelizing codes with MPI and/or OpenMP
- Optimization of scaling properties on large numbers of CPUs
- Usage of highly optimized parallel libraries, e.g. PBLAS, ScaLAPACK, PETSc, IBM WSMP etc.

To develop an efficient parallelization concept it is already necessary to have a deep insight into the code structure.

→ Such work will be usually on scales of many months and its outcome is not always precisely predictable.

Improvement of the efficiency of algorithms

Try to reduce as much as possible the number of FLOPS needed to solve a given problem → usage of highly efficient parallel algorithms

Example:
Using multigrid instead of the conjugated gradient method to solve a discretized PDE.

- The multigrid method is very efficient but complex.
- There are no general purpose numerical libraries.
- A multigrid solver has to be adapted by hand
  → know how is mandatory

Note: Some algorithms become inefficient/efficient for massively parallel usage.

High Level Support Team

Education and Training

- Optimal flow of information within HLST core and staff (permanent internal training)

- Providing links about HPC issues on HLST web site, e.g. scripts about numerical algorithms triggered by HLST projects:
  - Discontinuous Galerkin Finite Element Methods (DG-FEM)
  - Parallelization of the Multigrid Method for HPC
  - Comparison of Different Methods for Performing Parallel I/O

- Collaboration with EUFORIA and GoTiT:
  - Exchange between EUFORIA and HLST on HPC projects
  - GoTiT e-seminar with HLST contribution

High Level Support Team

Cooperation with HLST

- Close collaboration with the developers is mandatory and should be established by personal meetings, video conferences and e-mail
  → project coordinators have to be prepared and have to be accessible
- Changes and improvements of the codes can be done only in agreement and with the support of the developers
- The provided support is flexible and problem-oriented within the framework of the submitted HLST proposal
  → flexible adaptation to problems which may occur
- Providing consultancy to further HPC specialists in and outside the fusion Associations

High Level Support Team

- The team should not be misused for doing the job of the developers:
  - low level programming work, e.g. clean-up work
  - code refactoring
  - implementation of new physics
- The HLST should be informed about other collaborators which work/worked on the performance improvement of the code
- Important contributions of HLST members to certain projects should be credited by co-authorship
- The code developers keep the responsibility for their codes:
  → The code changes have to be finally accepted by the developers!

High Level Support Team

HLST call

The call is launched once a year

- Addressed to scientists from the EFDA associates
- Improvement of existing codes, i.e. efficiency and/or scalability
- Development of numerical libraries/tools, e.g. visualization of large data sets
- Maximal allocatable resources of 12 ppm (exceptions possible)
- Special requests, e.g. performing benchmark suite for IFERC procurement on HPC-FF

High Level Support Team

Projects 2009

Call launched on 12th December 2008: 10 proposals → 5 selected

Project	Project Coordinator	Institution	Alloc. Resources (ppm)	Status
BEUPACK	J. David	CEA	4	finished
GYGLES	P. Helander	IPP	6	finished
OPTGS2	W. Arter	UKAEA	6	finished
ORBIS	A. Bottino	IPP, CRPP	8	finished
JOREK-HR	G. Huysmans	CEA	8	in progress
MGEDGE	B.D. Scott	IPP	6	finished

High Level Support Team

Projects 2010

Call launched on 28th October 2009:
13 proposals received and all have been approved

→ total request for 97 ppm with an average request of 7 ppm

The participating institutions are located in:
Austria, Finland, France, Germany, Great Britain, Spain and Switzerland

Start date: April 2010

High Level Support Team

Project	Project Coordinator	Institution	Alloc. Resources (ppm)	Status
ASCOT-10	T. Kurki-Suonio	Aalto University	9	in progress
ELMU	J. Heikkinen	VTT	6	in progress
EUTERPE	E. Sánchez	CIEMANT/ IPP	6	in progress
GENEOPT	F. Jenko	IPP/CRPP	6	in progress
GYGLES	A. Mishchenko	IPP/CRPP	6	in progress
GYNVIZ	M. Haefele	HLST	24	in progress

High Level Support Team

Project	Project Coordinator	Institution	Alloc. Resources (ppm)	Status
IMPGS2	W. Arter	CCFE	7	in progress
ITM-EU4IA	D. Coster	ITM	3	finished
JOREK-HR	G. Huysmans	CEA Cadarache/ Univ. Bordeaux I	8	in progress
KINSOL2D	T. Tskhakaya	Univ. of Innsbruck	6	in progress
PARAMAR	R. Vila	CIEMAT	6	in progress
SPKMC	A. Ibarra	CIEMAT	6	in progress
ZOFLIN	K. Hallatscheck	IPP	6	scheduled

High Level Support Team

Final goal of HLST evolution

The more insight into the physics/work flow of a code, the better are the chances to significantly improve the efficiency of a code.

→ The HLST members have to enhance their HPC skills to become specialists in algorithms used in certain fields of plasma physics simulations.

To mature such experts will take years but otherwise we will not reach future petaflop computing (IFERC-CSC in 2012) for a significant number of plasma physics codes.

HLST comes just in time!