Plasma Physics: Scientific and Computational Challenges

Fusion, EFDA, ITM and EUFORIA

David Coster

Contributors to EUFORIA

Contributors to EFDA-TF-ITM

Grid Computing: a new tool for Science and Innovation

2009-08

Kulturni Dom of Veli Lošinj

EUFORIA FP7-INFRASTRUCTURES-2007-1, Grant 211804

Outline

- Fusion
  - Background
- EFDA (European Fusion Development Agreement)
  - Organization of fusion within Europe
- ITM
  - EFDA Task Force on Integrated Tokamak Modelling
- EUFORIA
  - EU for Iter Applications

Fusion

- Energy source for the sun and other stars
- Provides a potential source of base load energy production
- Been working on this for more than 50 years
- Has turned out to be a very difficult problem

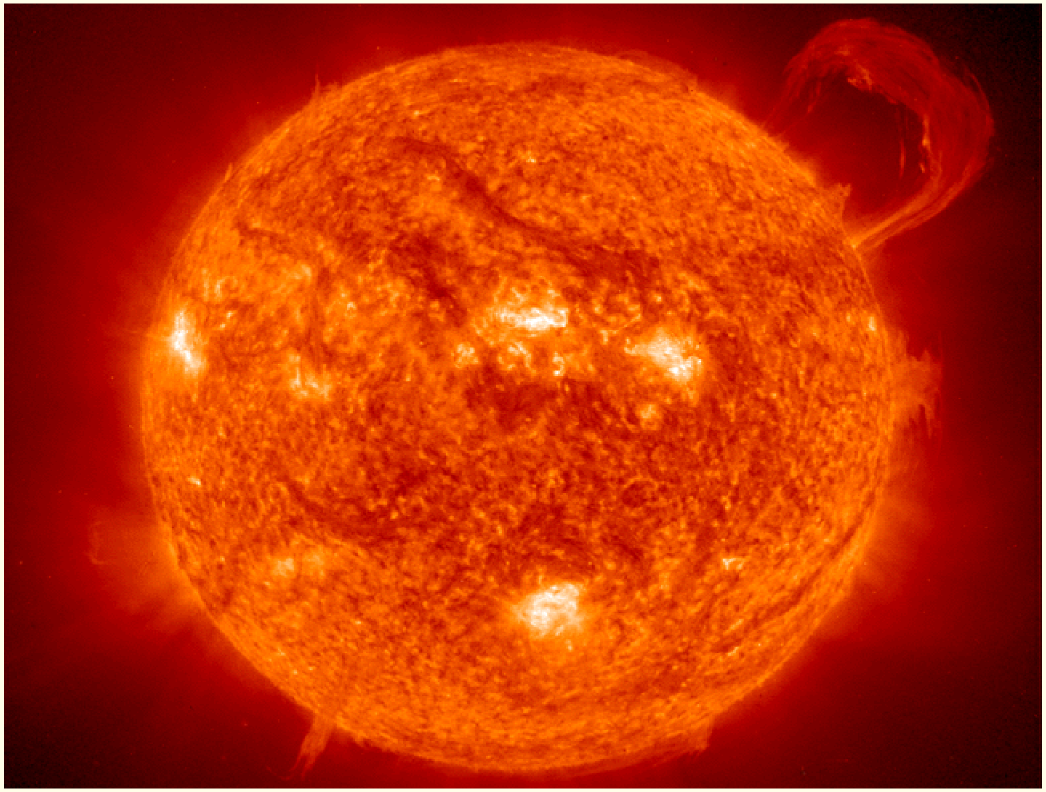

Fusion

D + T → ^{4}He (3.561) + n(14.029)
D + D → T(1.011) + p(3.022) (50%)
D + D → ^{3}He(0.820) + n(2.449) (50%)
p + T → ^{3}He + n – 0.764
T + T →^{4}He + 2n + 11.332
D + ^{3}He →^{4}He(3.712) + p(14.641)

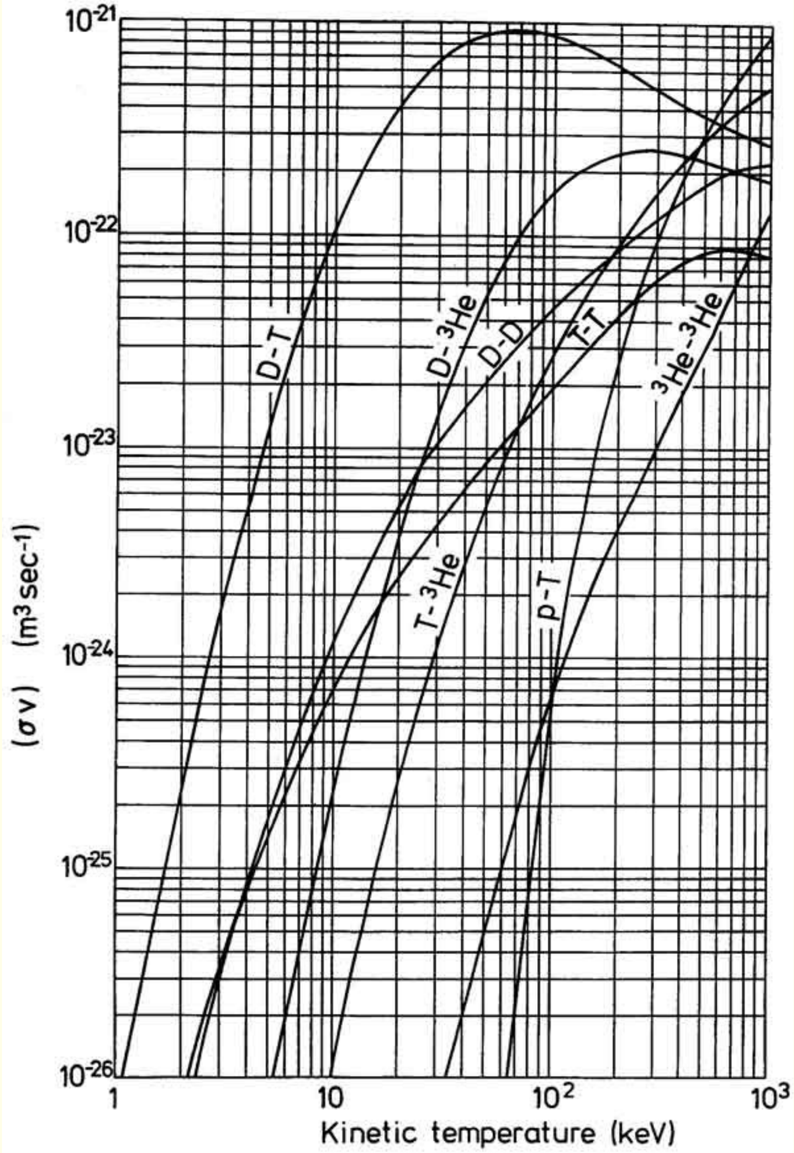

Fusion

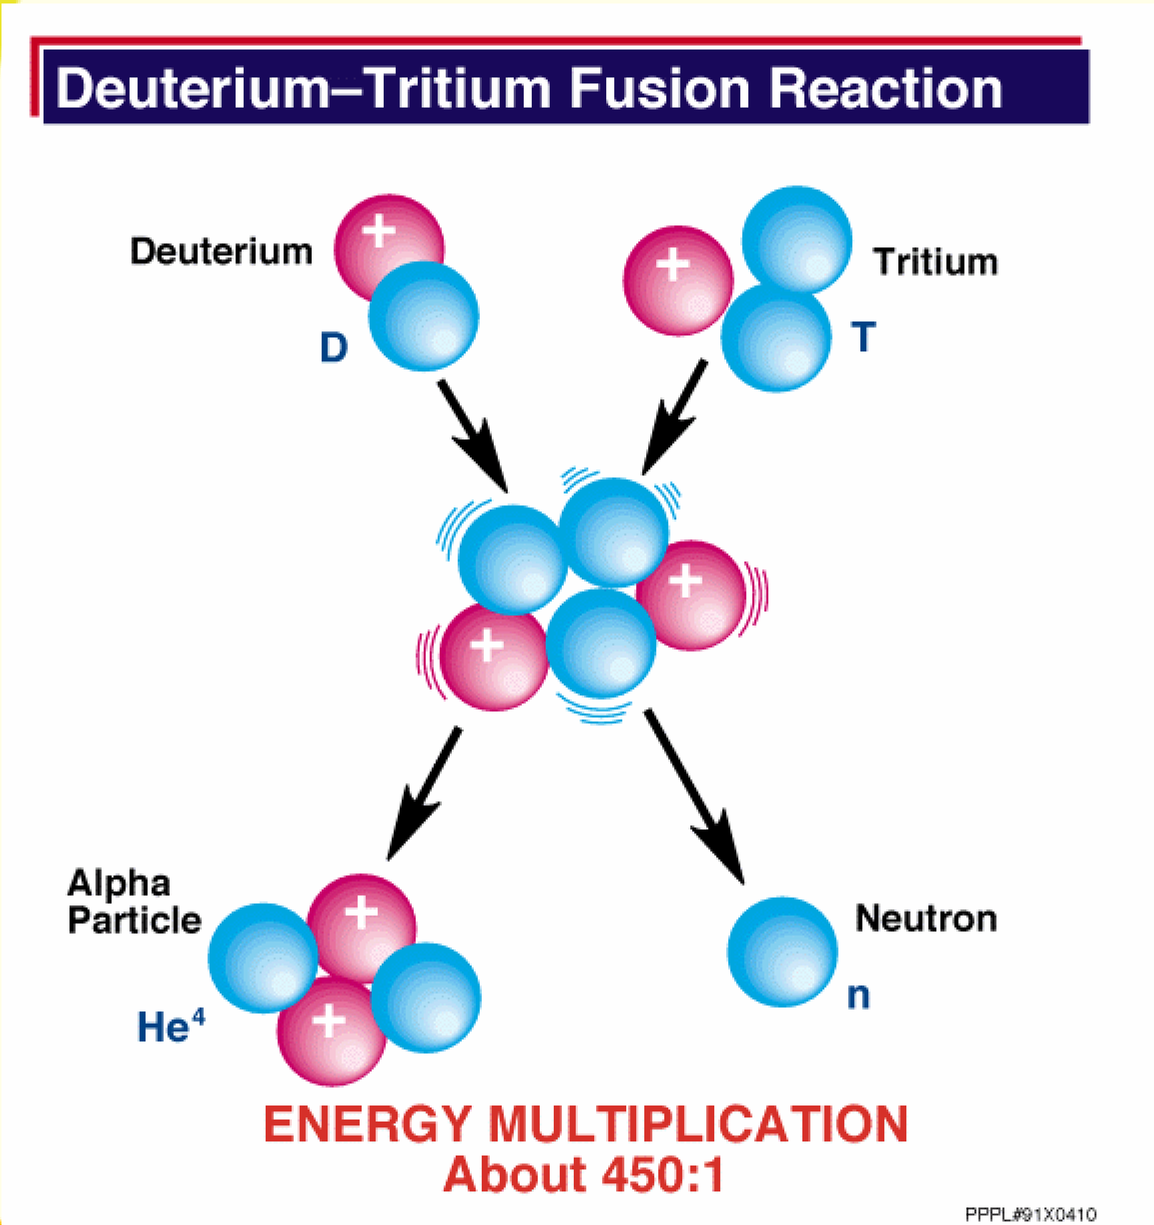
Deuterium–Tritium Fusion Reaction
Deuterium
D
Tritium
T
Alpha Particle
He4
Neutron
n
ENERGY MULTIPLICATION
About 450:1
PPPL#91X0410

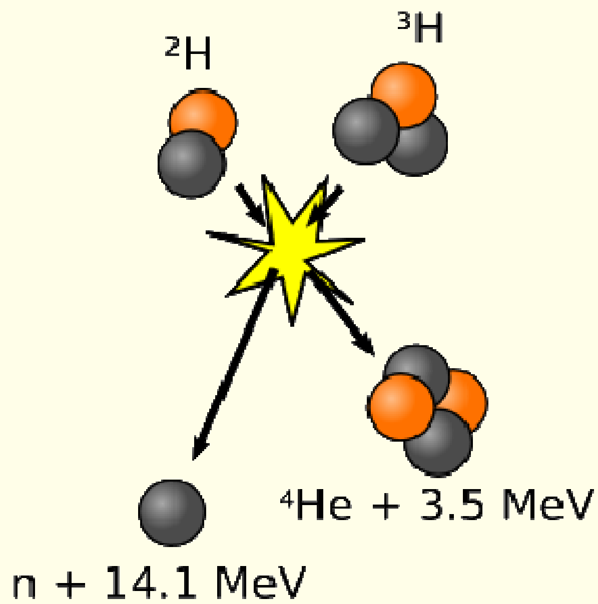
²H
³H
⁴He + 3.5 MeV
n + 14.1 MeV

Fusion

- Two main lines of research
  - Inertial confinement
    - Implosion of small pellets
    - NIF at LLNL
  - Magnetic confinement
    - Two main lines of research at the moment
      - Stellarator – W7X
        - Currently under construction in Greifswald in Germany
      - Tokamak – ITER
        - To be constructed in Cadarache in France

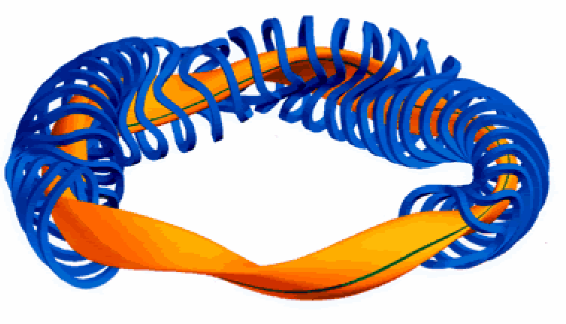

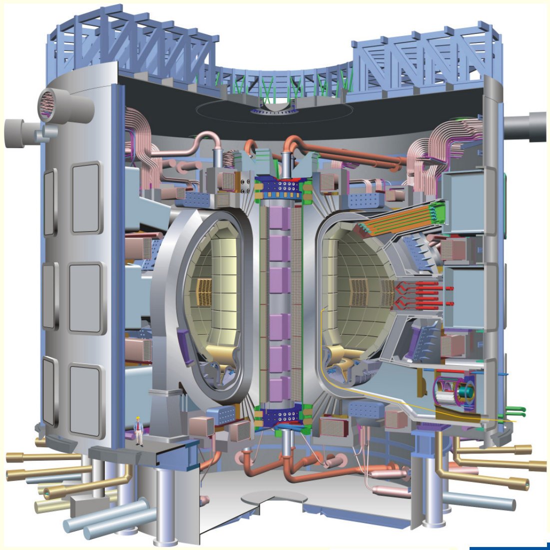

NIF at LLNL

W7X

ITER

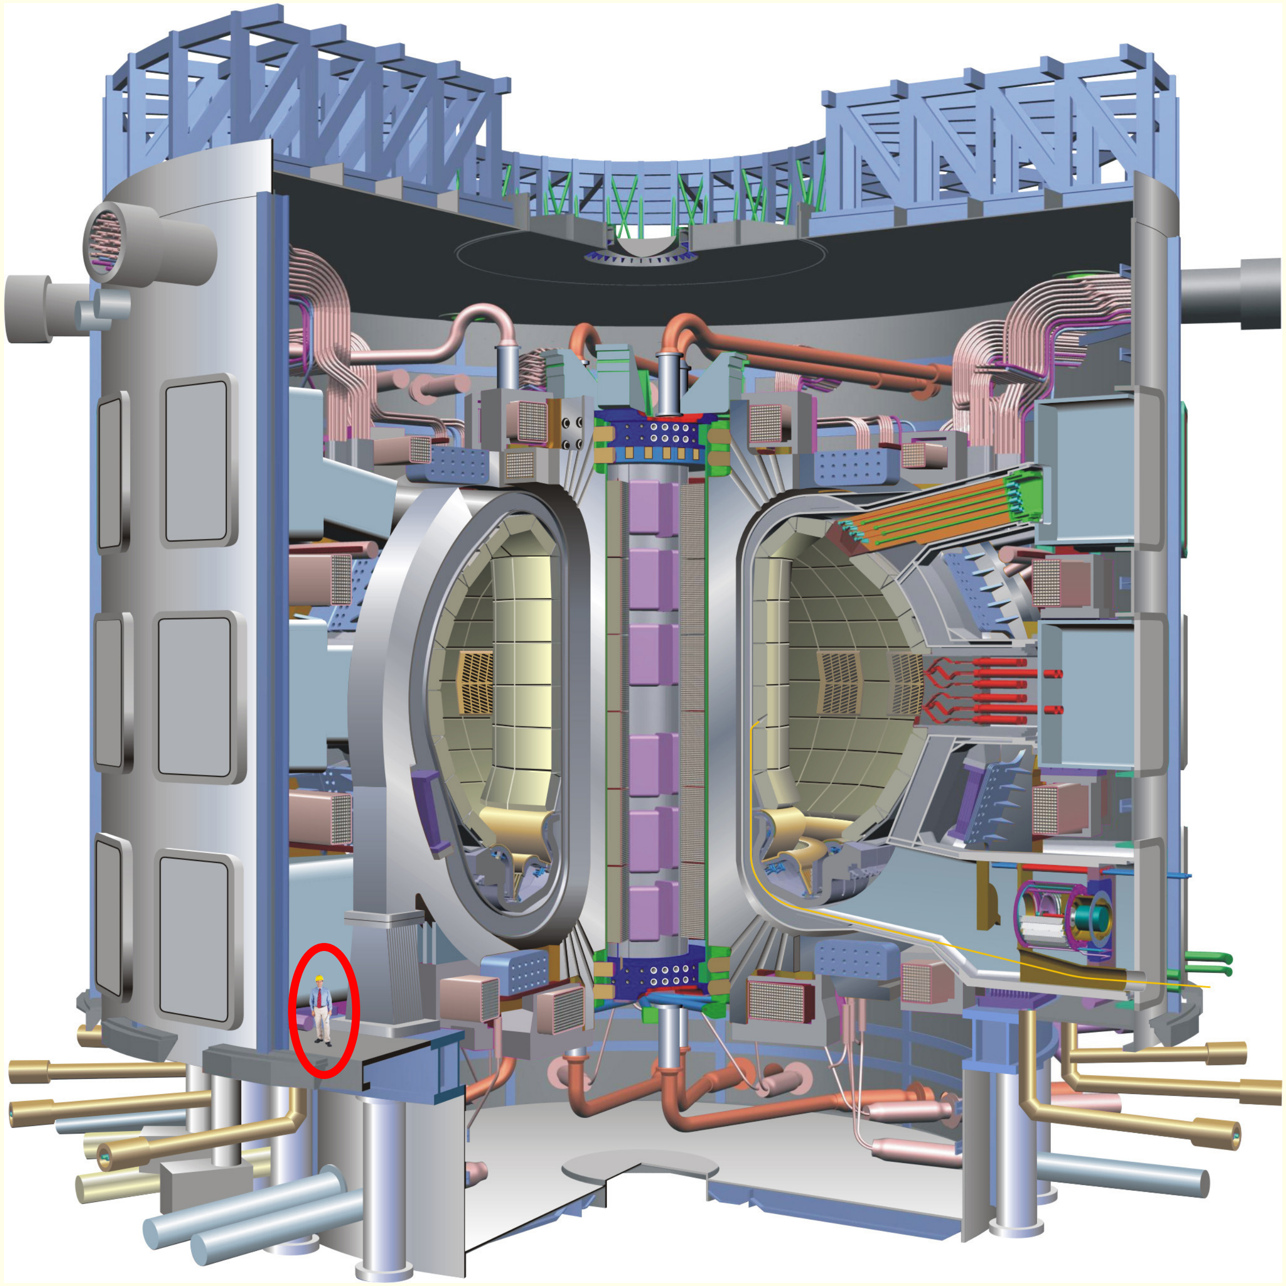

Fusion

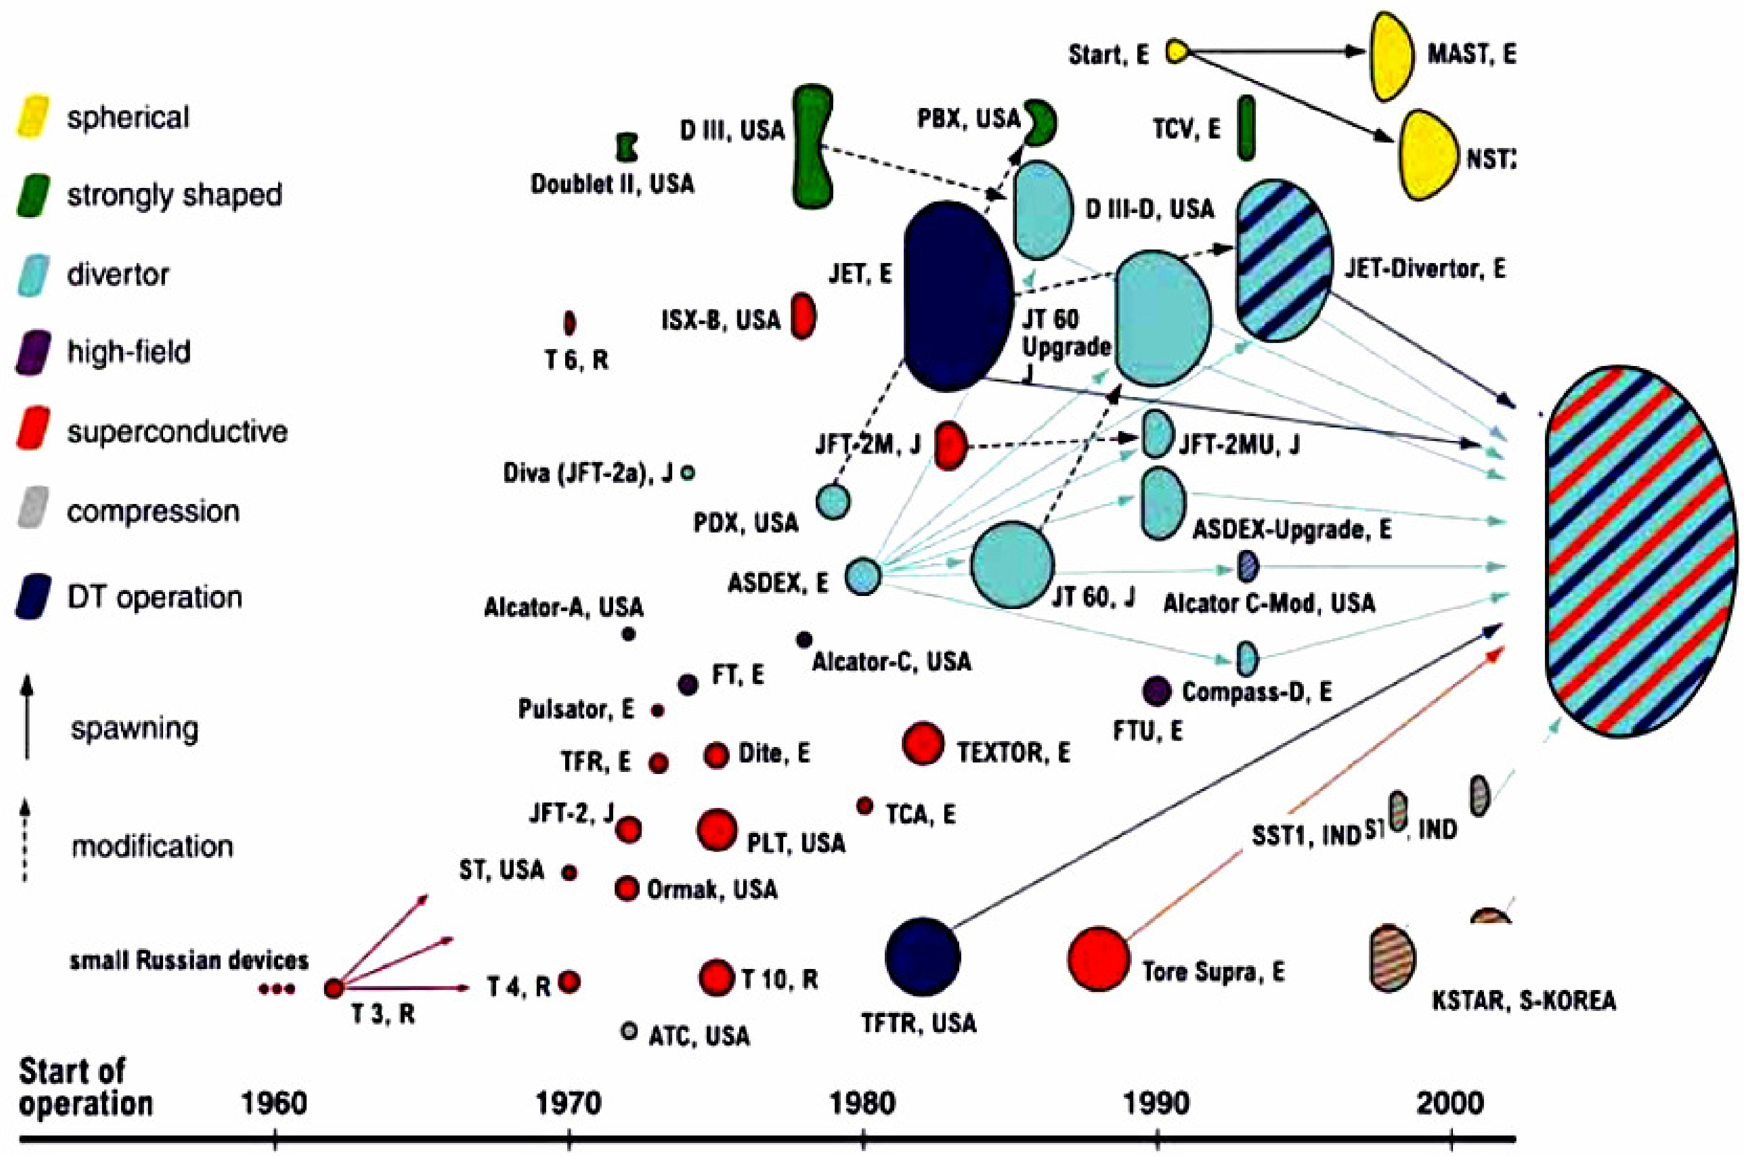
spherical
strongly shaped
divertor
high-field
superconductive
compression
DT operation
spawning
modification
Start, E
MAST, E
NST:
TCV, E
PBX, USA
D III, USA
Doublet II, USA
D III-D, USA
JET, E
JET-Divertor, E
JT 60 Upgrade J
ISX-B, USA
T 6, R
JFT-2M, J
JFT-2MU, J
Diva (JFT-2a), J
PDX, USA
ASDEX, E
ASDEX-Upgrade, E
JT 60, J
Alcator C-Mod, USA
Alcator-A, USA
Alcator-C, USA
FT, E
Compass-D, E
Pulsator, E
FTU, E
TFR, E
Dite, E
TEXTOR, E
JFT-2, J
PLT, USA
TCA, E
SST1, IND
IND
ST, USA
Ormak, USA
small Russian devices
T 3, R
T 4, R
T 10, R
ATC, USA
TFTR, USA
Tore Supra, E
KSTAR, S-KOREA
Start of operation
1960
1970
1980
1990
2000

Fusion developing faster than Moore's Law

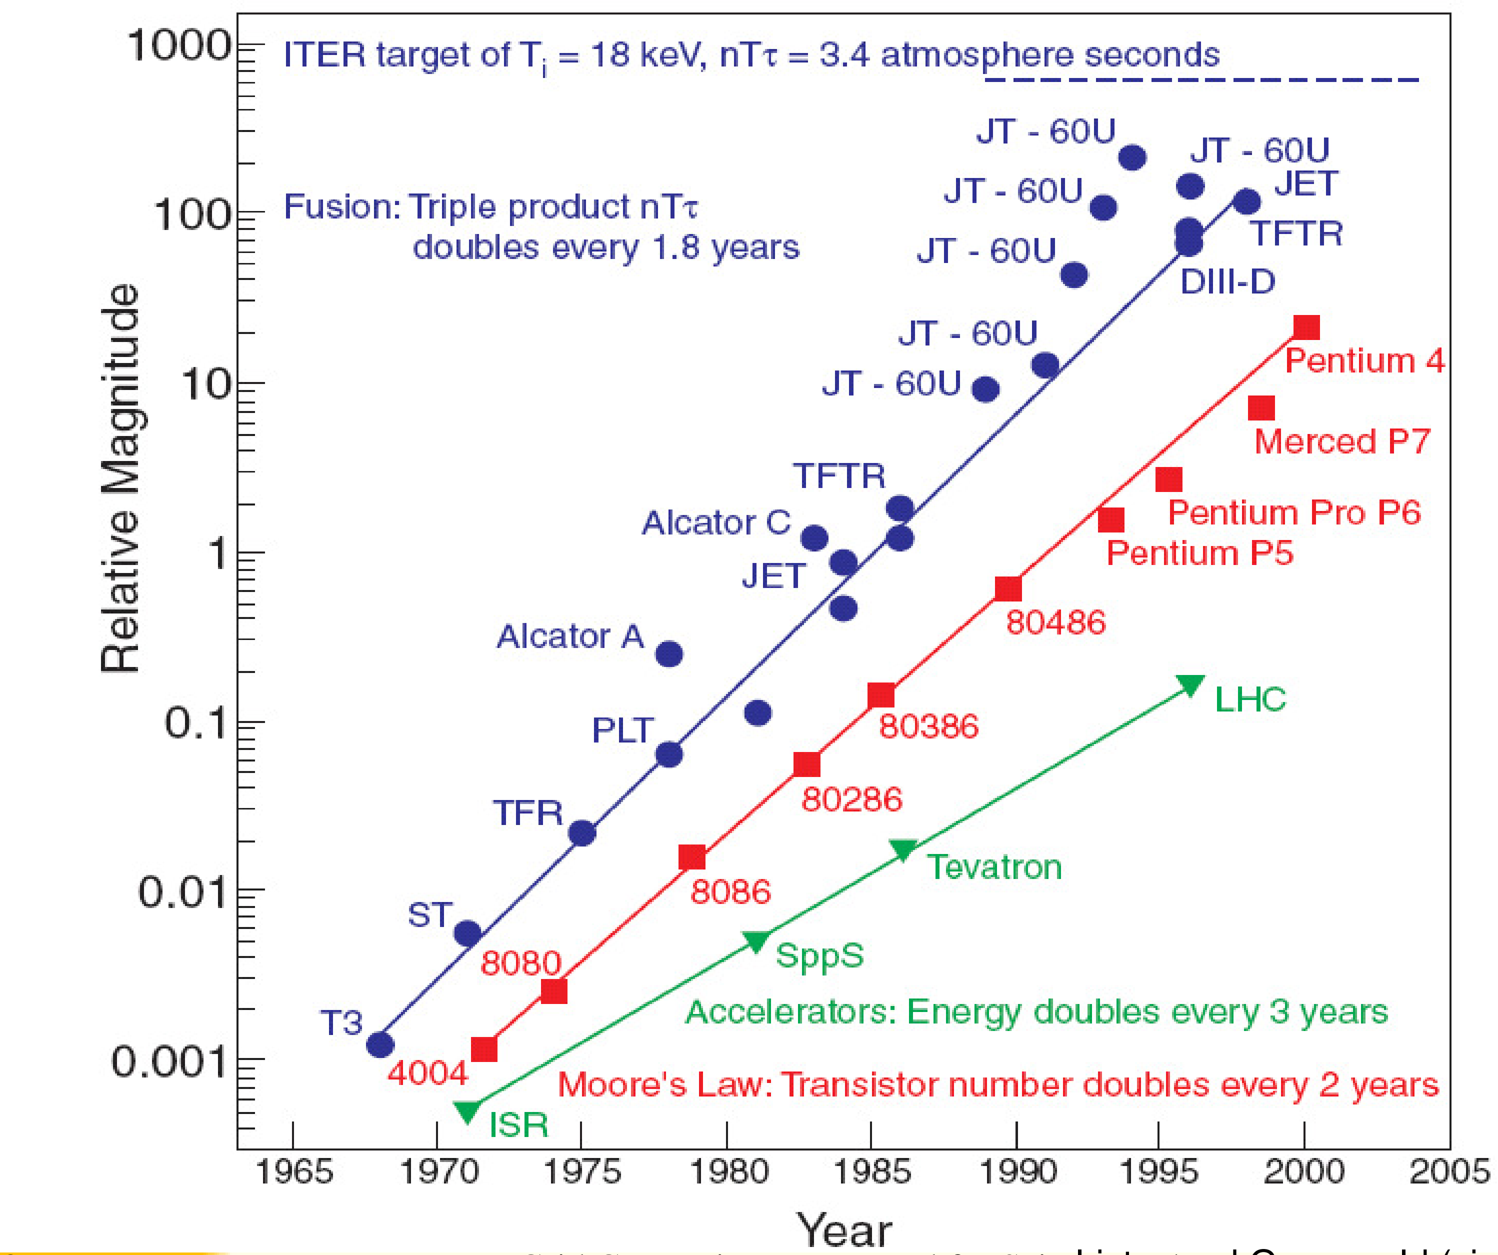

Lister and Greenwald (via Webster)

ITER

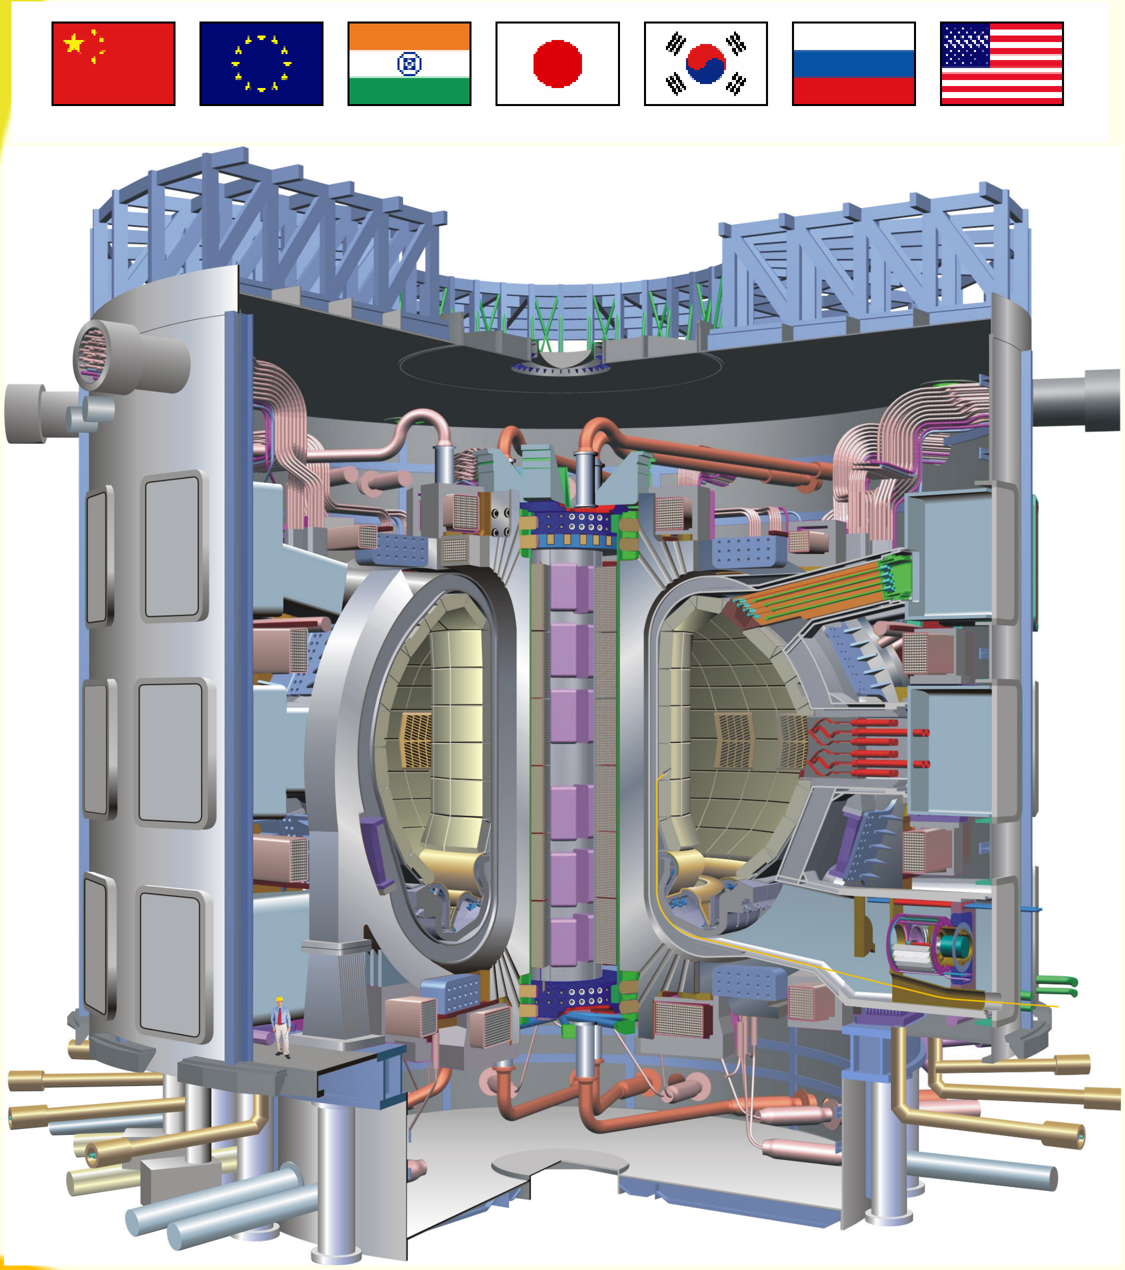

Involves 7 partners representing more than 50% world population
Costs > 10 G$
Under construction in Cadarache, France
Key element on the path to fusion energy production

ITER

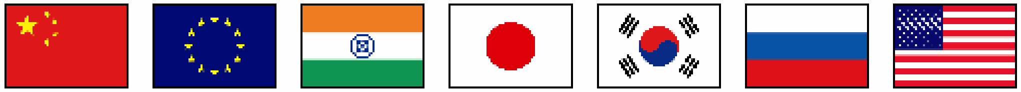

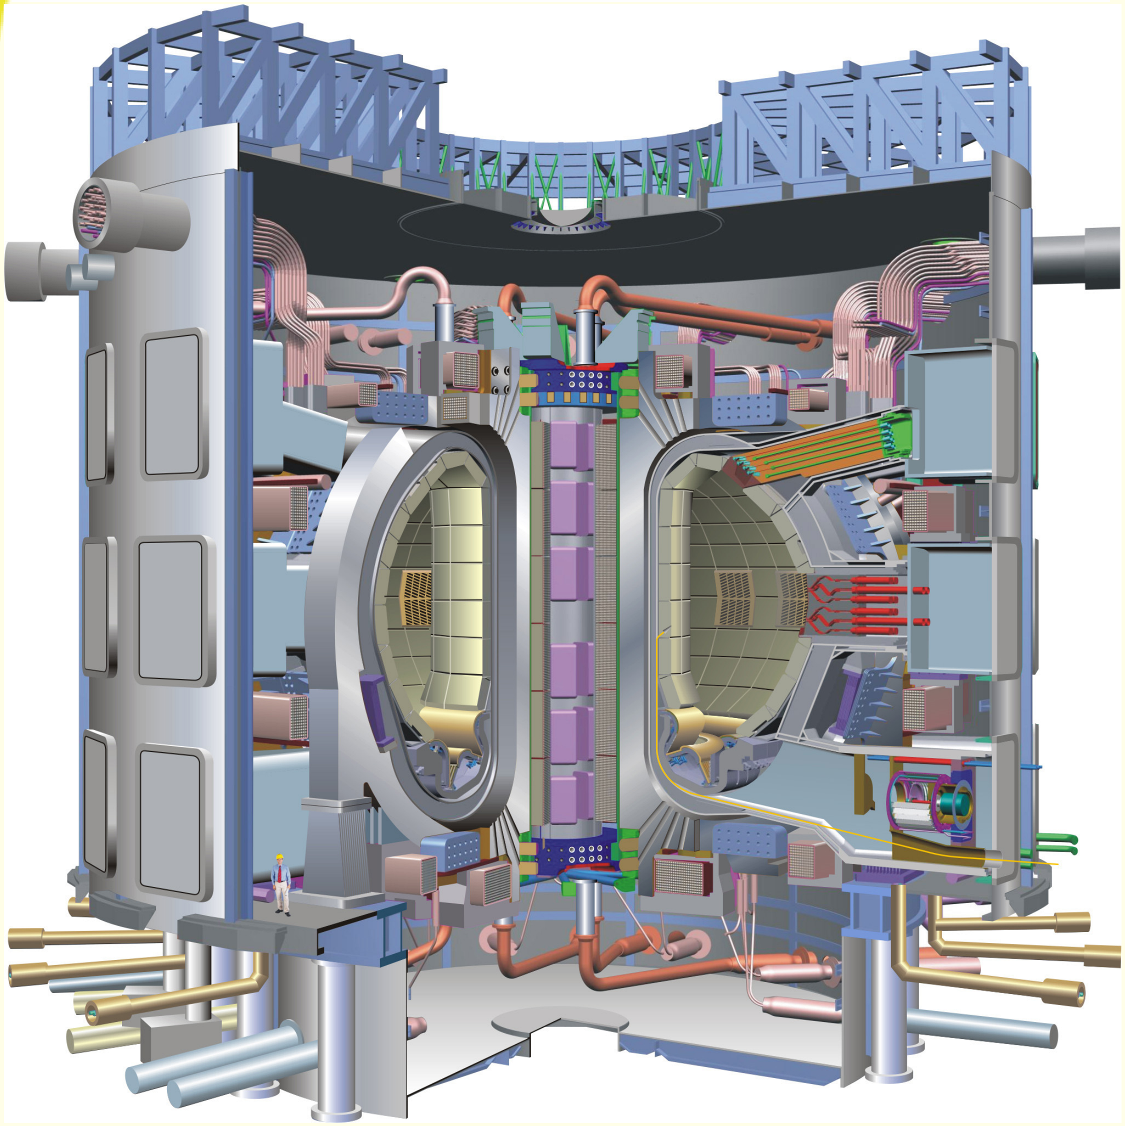

		Units
Plasma Major Radius	6.2	m
Plasma Minor Radius	2.0	m
Plasma Volume	840	m3
Plasma Current	15.0	MA
Toroidal Field on Axis	5.3	T
Fusion Power	500	MW
Burn Flat Top	>400	s
Power Amplification	>10	

EFDA

European Fusion Development Agreement

All EU Laboratories / Institutions working on Fusion are parties to EFDA

Defined under EURATOM under "Contract of Associations"

EURATOM : KEY ACTION FUSION Associated Laboratories, parties to EFDA

Euratom - Belgian State
(Brussels) - (Mol)

Euratom - CEA
TORE SUPRA (Cadarache)

Euratom - CIEMAT
TJ-II (Madrid)

Euratom - Conf. Suisse
TCV - SULTAN
(Lausanne) - (Villigen)

Euratom - DCU
(Dublin) - (Cork)

Euratom - ENEA
FTU - RFX
(Frascati) - (Milan) - (Padua)

Euratom - FOM
(Petten) - (Nieuwegein)

Euratom - FZJ
TEXTOR (Julich)

Euratom - FZK
TOSKA (Karlsruhe)

Euratom - Greece
(Athens) - (Heraklion) - (Ioannina)

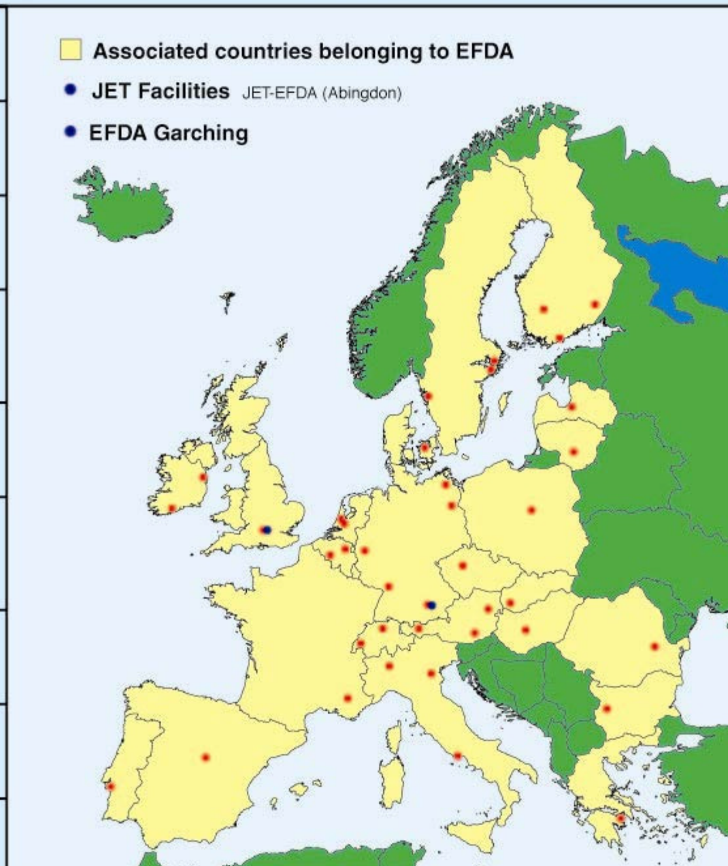

Euratom - INRNE
(Sofia)

Euratom - LEI
(Kaunas)

Euratom - CU
TOSKA (Bratislava)

Euratom - HAS
(Budapest)

Euratom - IPP
Asdex Upgrade - Wendelstein 7-AS
Wendelstein 7-X
(Garching) - (Greifswald) - (Berlin)

Euratom - IPP.CR
CASTOR (Prague)

Euratom - IST
ISTTOK (Lisbon)

Euratom - Latvia
(Riga)

Euratom - MEC
(Bucharest)

Euratom - ÖAW
(Vienna) - (Graz) - (Innsbruck)

Euratom - RISØ
(Roskilde)

Euratom - TEKES
(Helsinki) - (Tampere) - (Lappeenranta)

Euratom - UKAEA
MAST - JET (Culham)

Euratom - VR
EXTRAP T2R (Stockholm) - (Lund)
(Gothenburg) - (Studsvik) - (Uppsala)

JG03.241-6c

EFDA

- **EURATOM Budget (for the period 2007-2011)**
  - Nuclear Security
  - Fusion energy research* €1947 m
    - *Within the amount foreseen for fusion energy research, at least €900 million will be reserved for activities other than the construction of the fusion energy source ITER.
  - Nuclear Fission and radiation protection €287 m
  - Nuclear Activities of the Joint Research Centre €517 m

Multi-facetted physics

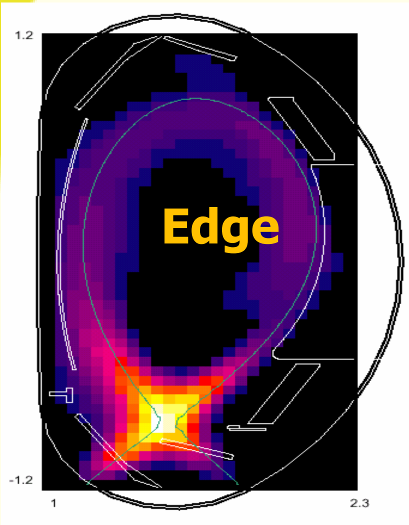

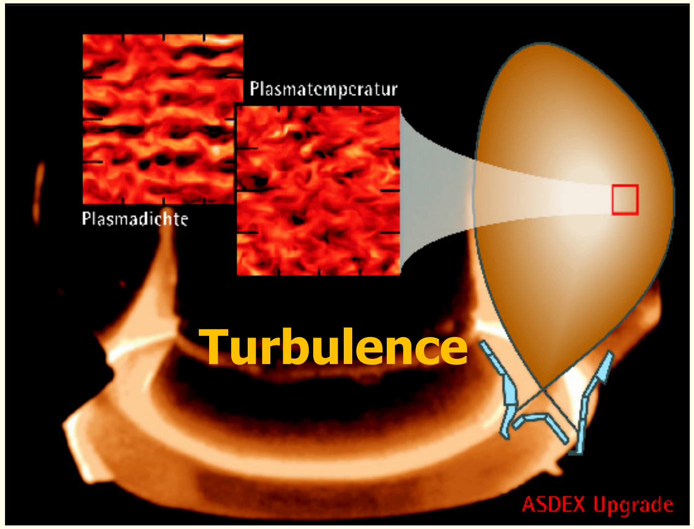

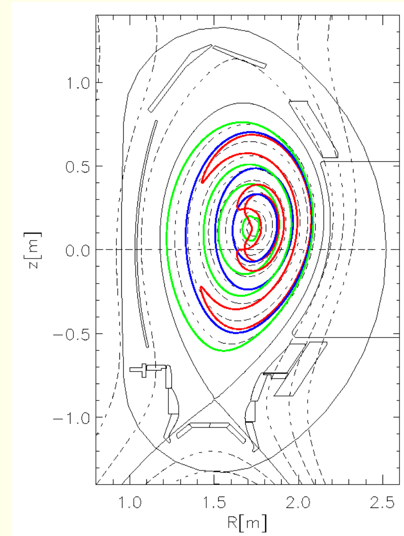

Energetic particles

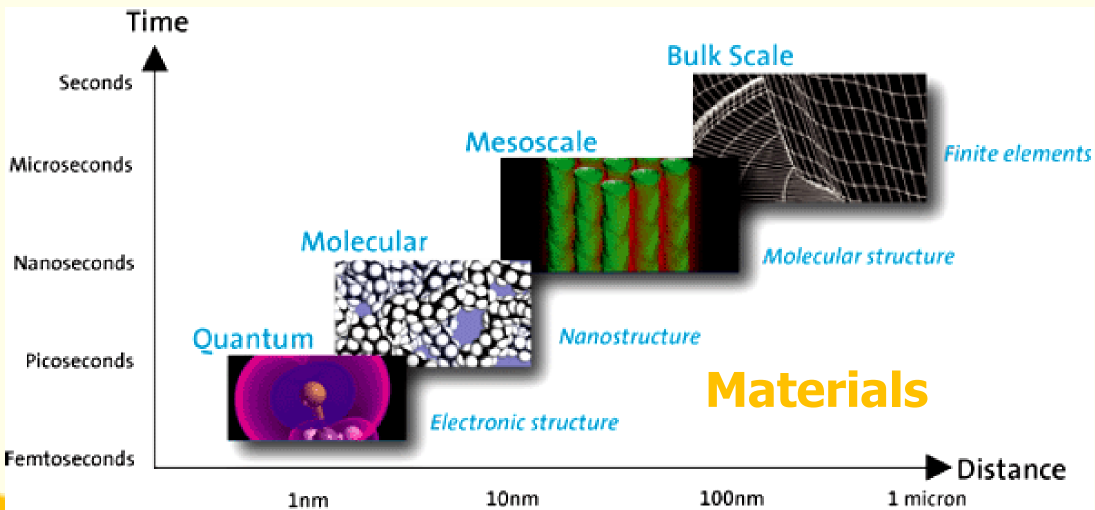

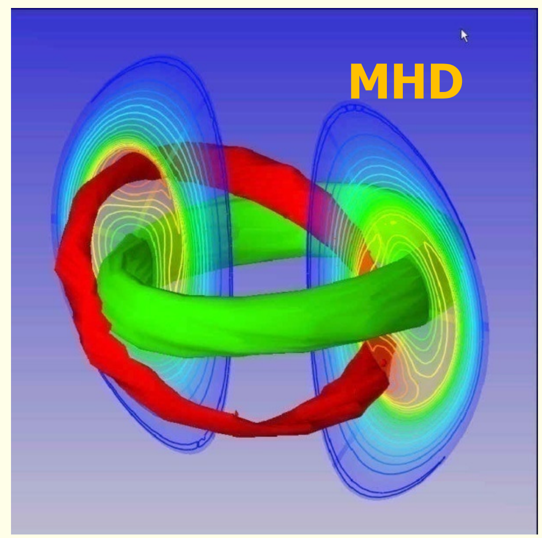

Large Computation Requirement

Have requirements for

- Experimental data processing
- Experimental scenario development
- Theory
  - Turbulence
  - Large scale MHD
  - Materials

Experimental Data

- JET (per shot)
  - ~ 10 Gb of raw data (JPF)
  - ~ 0.4 Gb of processed data (PPF) from “Chain1”
  - Chain1 takes around 4 minutes to run on 6 processors
  - ~ 0.2 Gb of additional processed data per pulse by “Chain2” and other analysis codes
  - A typical “Chain2” run (they are quite variable as the ICRH and NBI FP codes take most of the time) takes an hour
- AUG
  - ~ 9 GB Raw data
  - ~ 0.6 GB Processed data

Shot cycle: computing with a deadline

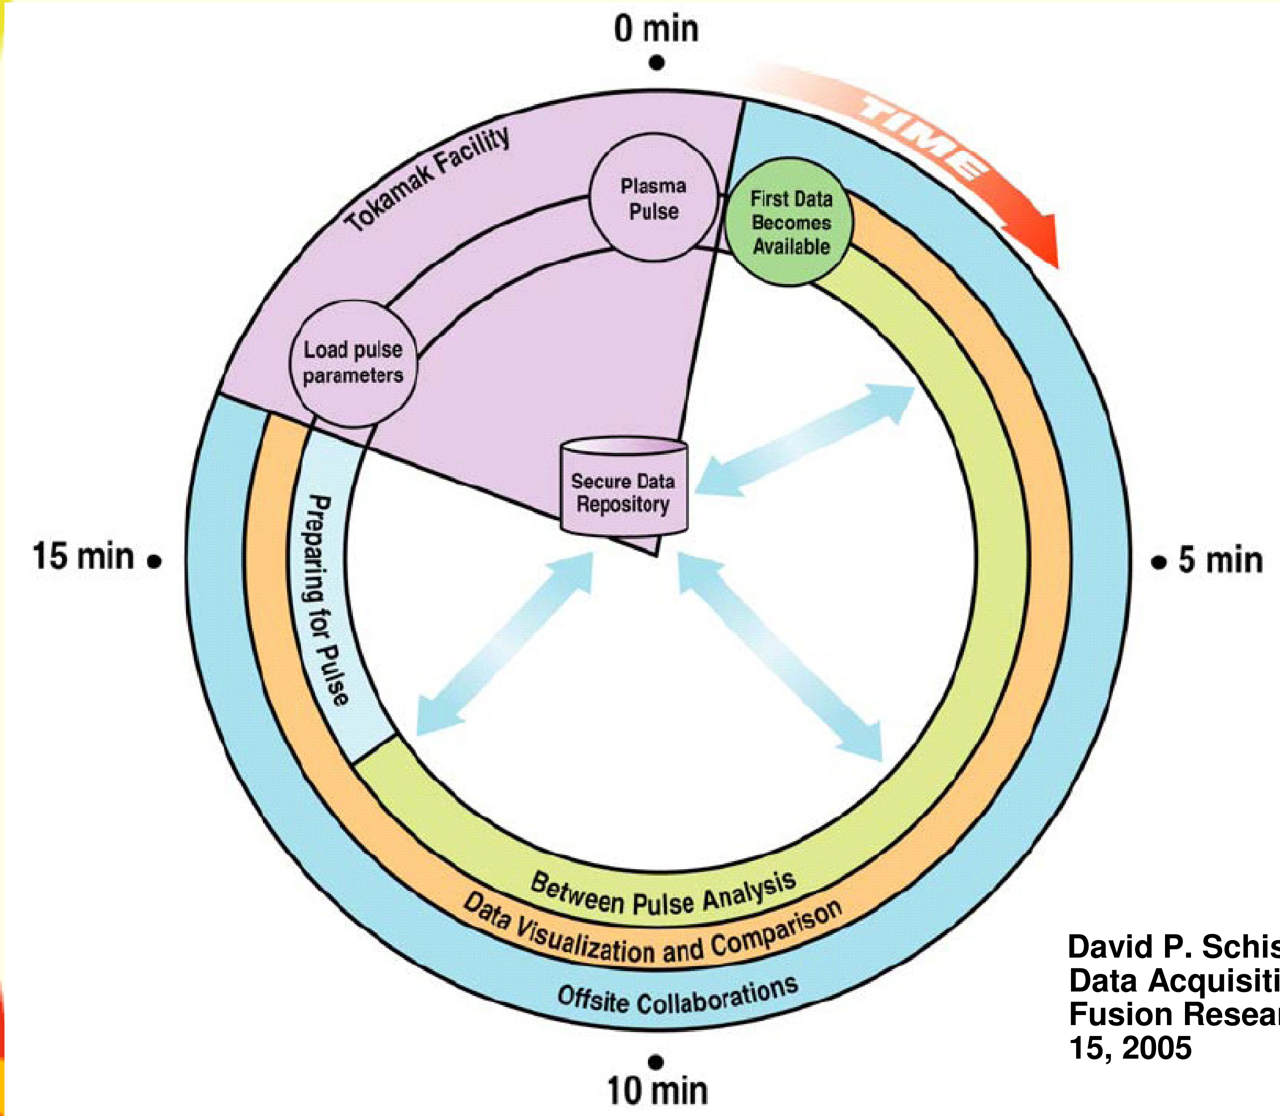

David P. Schissel, 5th IAEA TM on Control, Data Acquisition, and Remote Participation for Fusion Research, Budapest, Hungary, July 12–15, 2005

Shot cycle: JET

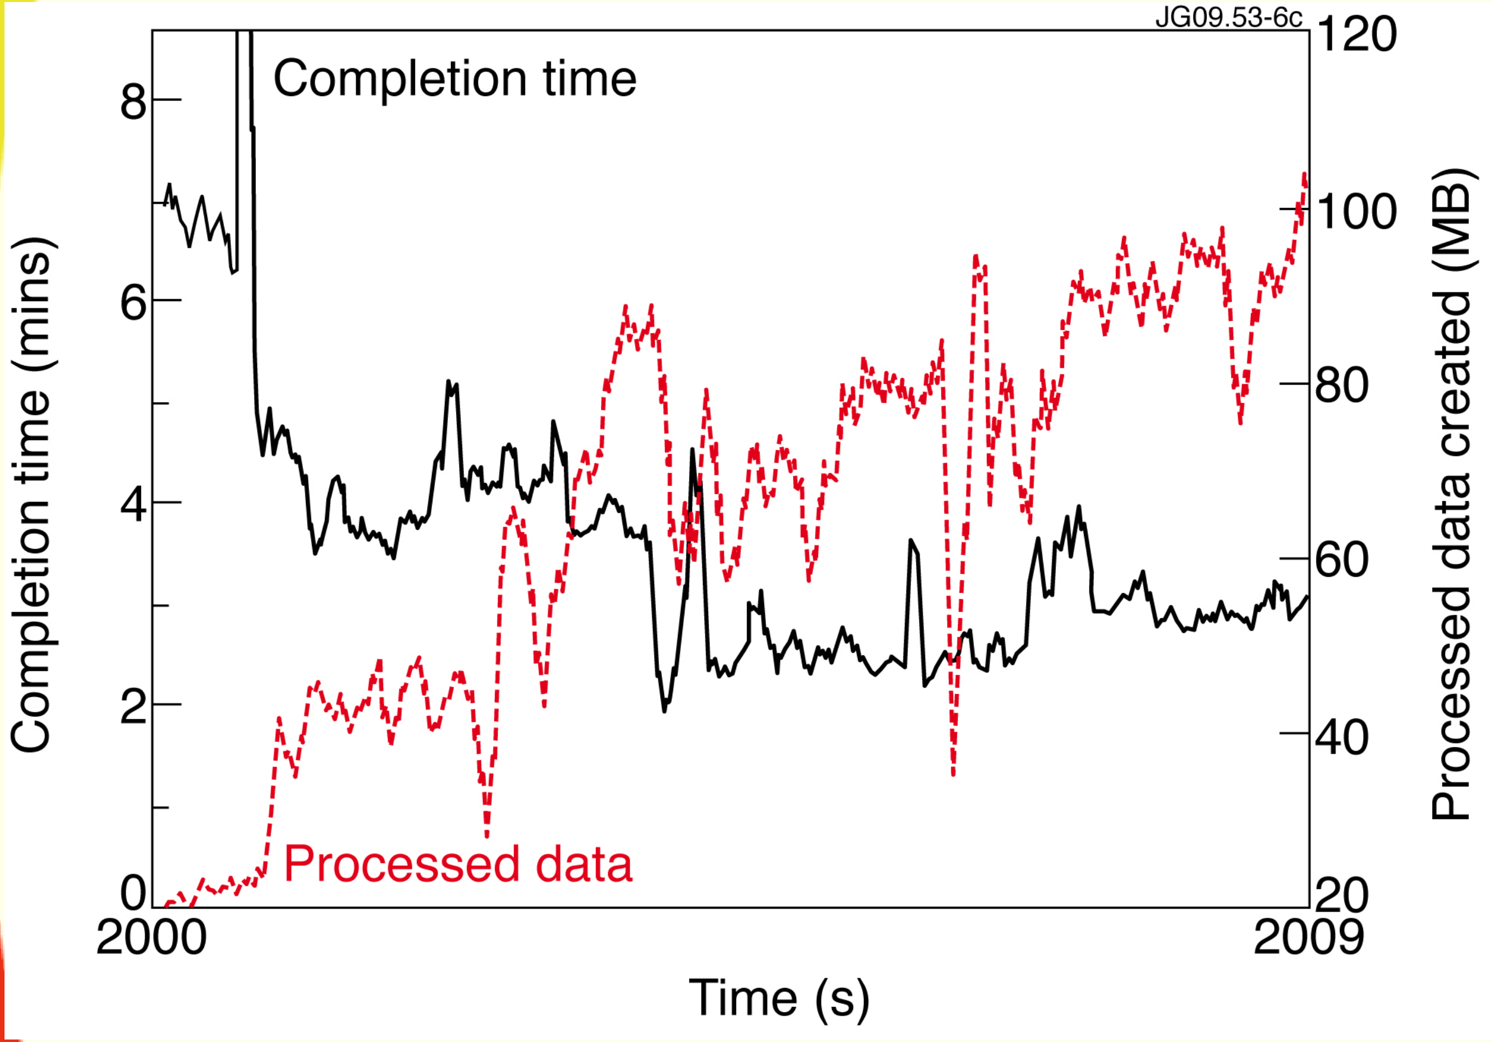

Richard Layne, Euratom-UKAEA Fusion Association/ Tessella plc, 7th IAEA TM on Control, Data Acquisition, and Remote Participation, 15/06/09

Integrated Data Analysis

- IDA: (ne, Te)
  - 1 time slice for Te and ne: approx. 1 min
  - 10 s shot / 50 μs resolution -> 140 days on a single CPU
  - (standard is 1 ms resolution only due to limited computation capabilities (power and memory), 50 μs are calculated only for selected shots)
- IDZ: (Zeff)
  - 1 shot with 50 ms resolution approx. 1 hour on a single CPU
- Er
  - more than days
- The plan is to introduce into IDA also the equilibrium solver (equil) and the corresponding magnetic signals. I expect that this will increase computation time considerably.

Theory Challenges

Grand challenge “single physics” simulations

- Turbulence simulations
- Materials simulations
- Wave simulations

Combined physics simulations

- Need to combine physics at very different time and space scales
- Very different computational demands for the individual pieces
- Modules (codes) have been developed over a number of years using a variety of computer languages

Large Range of Scales

Time

- $10^{-11} << 10^{-8} << 10^{-4} << 10^{-2} << 10^{+1}$
- (plasma freq; ion cyclotron freq; electron collision time; ion collision time; confinement time)

Space

- $10^{-4} << 10^{-3} << 10^{+0} << 10^{+4}$
- (Debye; ion Larmor radius; a; electron mfp)

Eliminate some processes by choosing the physics model

But still need to bridge

- 10^{+5} in time
- $(10^{+3})^2$ in space

Major Theoretical & Algorithmic Speedups

relative to simplest brute force, fully resolved, algorithm, for ITER $1/\rho_* = a/\rho \sim 700$

Nonlinear gyrokinetic equation	
• ion polarization shielding eliminates plasma freq. $\omega_{pe}/\Omega_{ci} \sim m_i/m_e$	$x10^3$
• ion polarization eliminates ρ_e & Debye scales $(\rho_i/\rho_e)^3$	$x10^5$
• average over fast ion gyration, $\Omega_{ci} / \omega_* \sim 1/\rho_*$	$x10^3$
Continuum or δf PIC, reduces noise, $(f_0/\delta f)^2 \sim 1/\rho_*^2$	$x10^6$
Field-aligned coordinates (nonlinear extension of ballooning coord.) $\Delta_{\|\|} / (\Delta_\perp\, q\, R / a) \sim a / (q\, R\, \rho_*)$	x70
Flux-tube / Toroidal annulus wedge, ↓ simulation volume	
$k_\theta \rho_i = 0, 0.05, 0.1, \ldots, 1.0$ $n = 0, 15, 30, \ldots, 300$ (i.e., 1/15 of toroidal direction)	x15
$L_r \sim a/5 \sim 140\, \rho \sim 10$ correlation lengths	x5
High-order / spectral algorithms in 5-D, 2^5 x 2	x64
Implicit electrons	x5-50
Total combined speedup of all algorithms	**$x10^{23}$**
Massively parallel computers (Moore's law 1982-2007)	**$x10^5$**

Hammett, APS DPP, Orlando, 2007

Current demands

A present day large, routine run may take 0.1 to 1.0 MCPUhrs (note that 1 MCPUhrs = 114 cpu years)

Large fusion projects in the USA might get 50 MCPUhrs

Awards in Europe are currently in the 1 to 5 MCPUhrs range for fusion projects

Some physics runs are being postponed because we cannot afford to run the cases!

Grid Computing: a new tool for Science and Innovation

Kulturni Dom of Veli Lošinj

EFDA

EUROPEAN FUSION DEVELOPMENT AGREEMENT

Task Force
INTEGRATED TOKAMAK MODELLING

ITM

TF Leader : P. Strand,
Deputies: L-G. Eriksson, R. Coelho

EFDA CSU Contact Person: D. Kalupin

EFDA ITM-TF

ITM-TF charge

- *Co-ordinate the development of a coherent set of validated simulation tools for ITER exploitation*
- *Benchmark these tools on existing tokamak experiments*
- *Provide a comprehensive simulation package for ITER and DEMO plasmas.*
- *Coordinate the European Software developments with the aim to increase quality and reduce parallel efforts. (Streamline the code base)*

ITM-TF Remit

- *Development of the necessary standardized software tools for*
  - *interfacing code modules and*
  - *accessing experimental data.*

Medium term activities

- *Support the development of ITER-relevant scenarios in current experiments.*

EFDA SC (03)-21/4.9.2 (June 24th, 2003)

EFDA ITM-TF

- Staffed by people from the associations under “Task Agreements”
- 240 people involved
- 60 ppy (1/3 under “Preferential” or “Priority” support ➔ 40% from Brussels (remaining paid 20% by Brussels))
- ~ 400 people in Europe involved in fusion theory and modelling

ITM-TF Milestones

2008

- **Start of the Gateway for IM and deployment of code platform (Real start of ITM collaborative activities)**
  - Implies a complete set of data structures and associated tools
  - A fully operational portal/workflow configuration
  - Major code releases from all of Integrated Modelling Projects

2009

- **Extended set of platform tools forming a predictive core physics capacity for ITER**
  - Production activities – local clusters and grid, HPC-FF

2011

- **Whole device modelling capability including comprehensive core-edge coupling and first principles elements**
  - Aiming towards Broader Approach IFERC level computations
  - HPC-FF & Prace

ITM

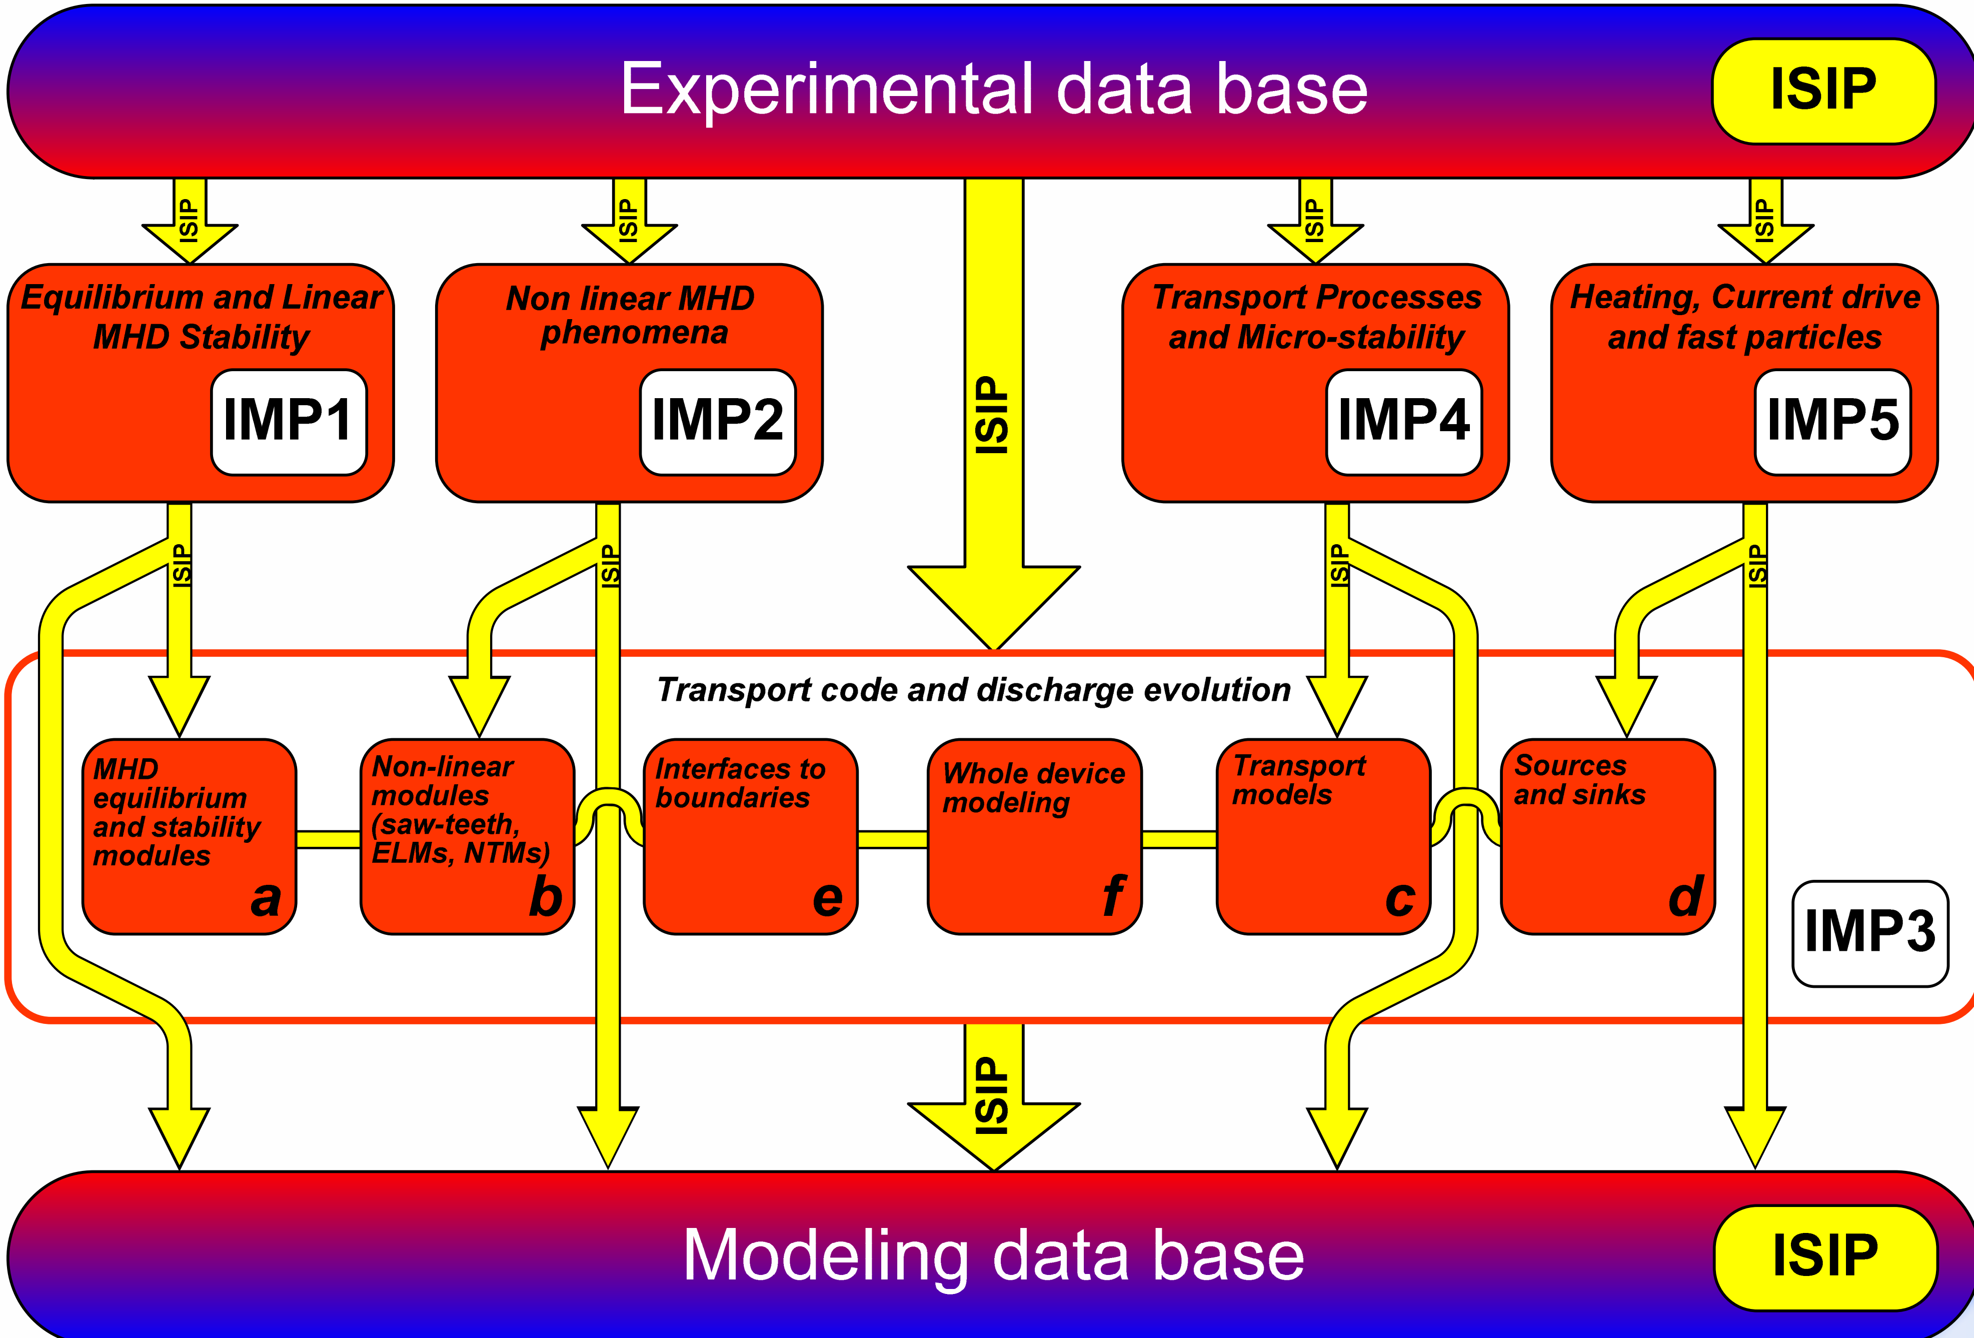
Experimental data base
ISIP
ISIP
ISIP
ISIP
ISIP
ISIP
Equilibrium and Linear MHD Stability
IMP1
Non linear MHD phenomena
IMP2
Transport Processes and Micro-stability
IMP4
Heating, Current drive and fast particles
IMP5
ISIP
ISIP
ISIP
ISIP
Transport code and discharge evolution
MHD equilibrium and stability modules
a
Non-linear modules (saw-teeth, ELMs, NTMs)
b
Interfaces to boundaries
e
Whole device modeling
f
Transport models
c
Sources and sinks
d
IMP3
ISIP
Modeling data base
ISIP

ITM-TF Application Structure

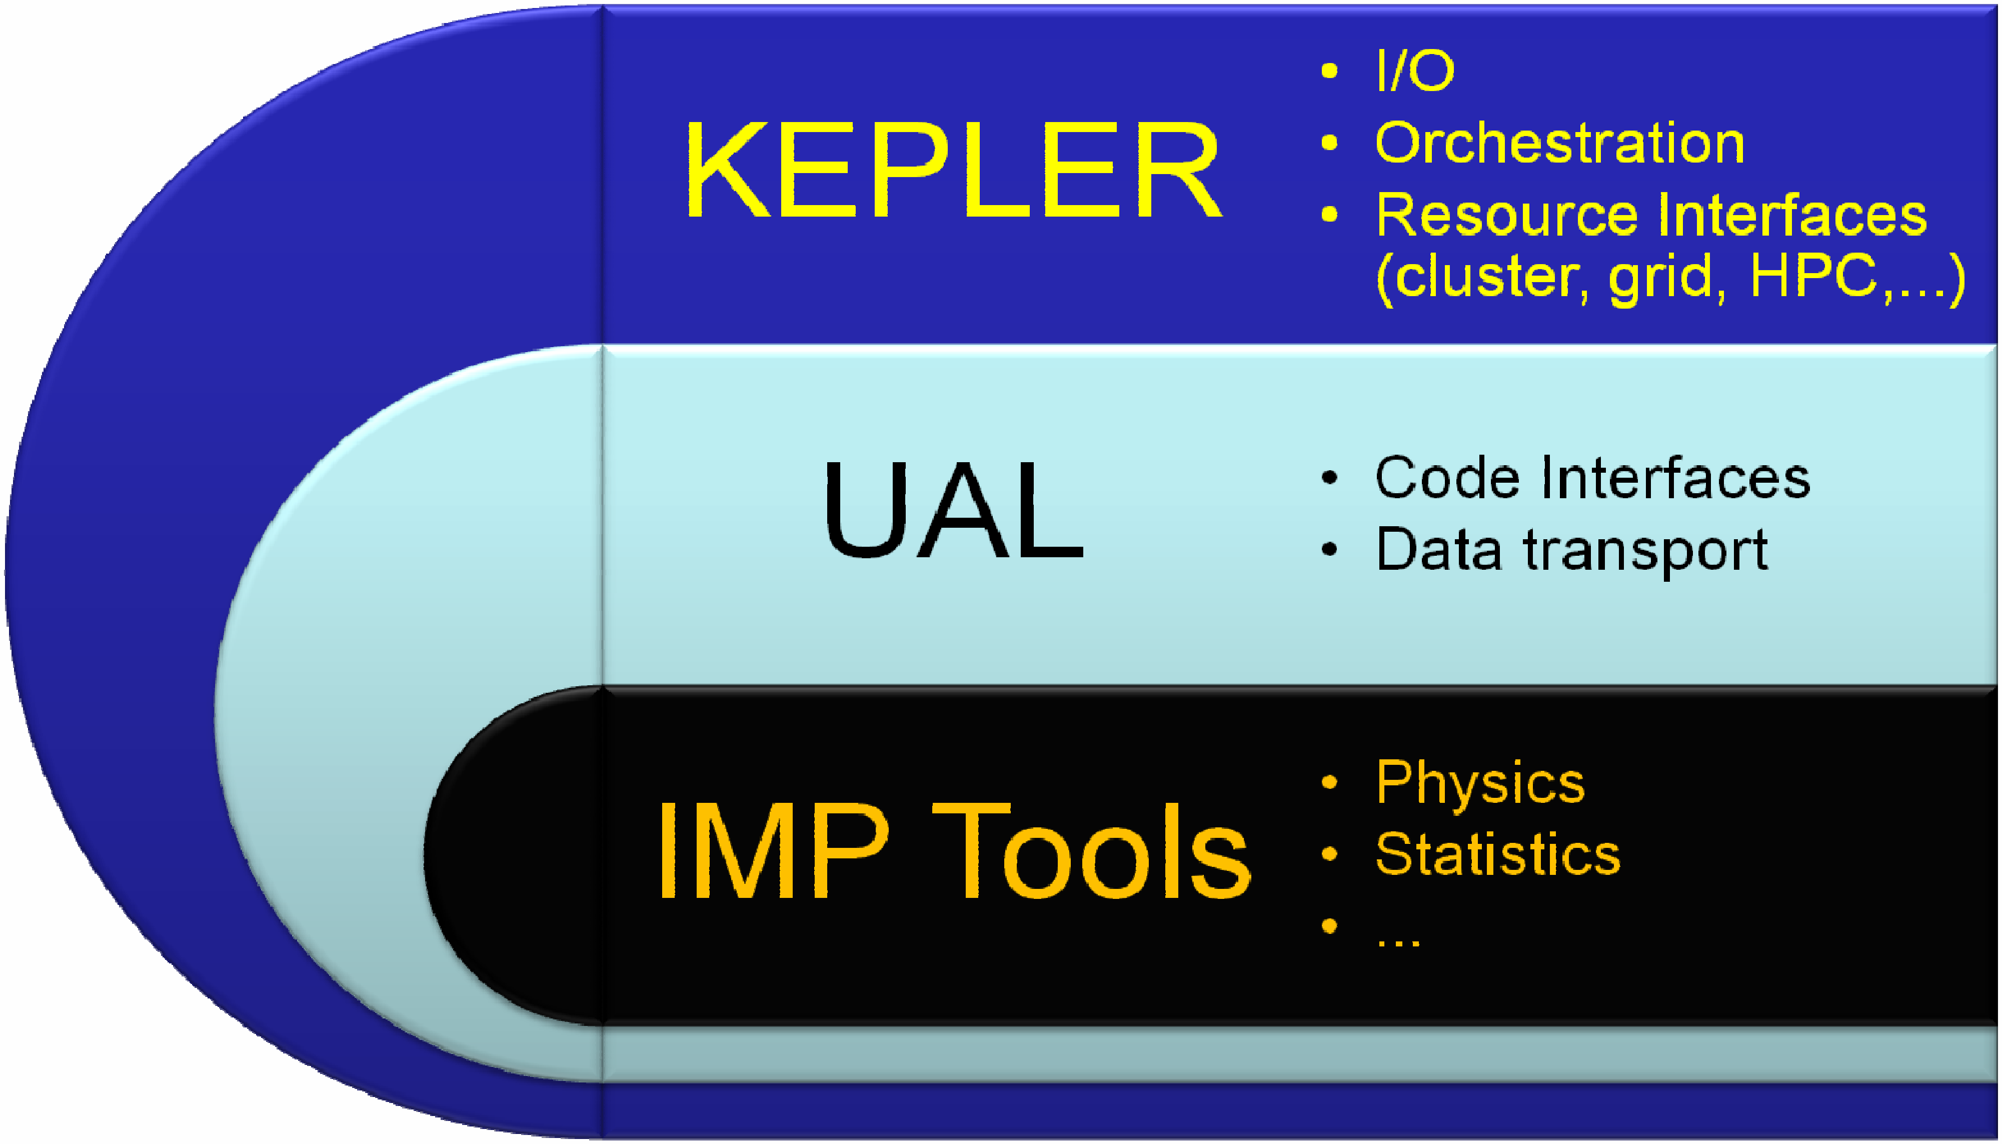

ITM-TF Application Structure

Framework (Kepler)

Calls the wrapper, specifying the present time of the simulation, orchestrates run.

Wrapper (FC2K, WS2K)

Calls UAL to GET the CPOin and CPOout at the requested time slices

Physics code

Receives CPOin, CPOout,

Physics calculations

Updates CPOout

Updates data management nodes

Calls UAL to PUT the CPOout

ETS Workflow

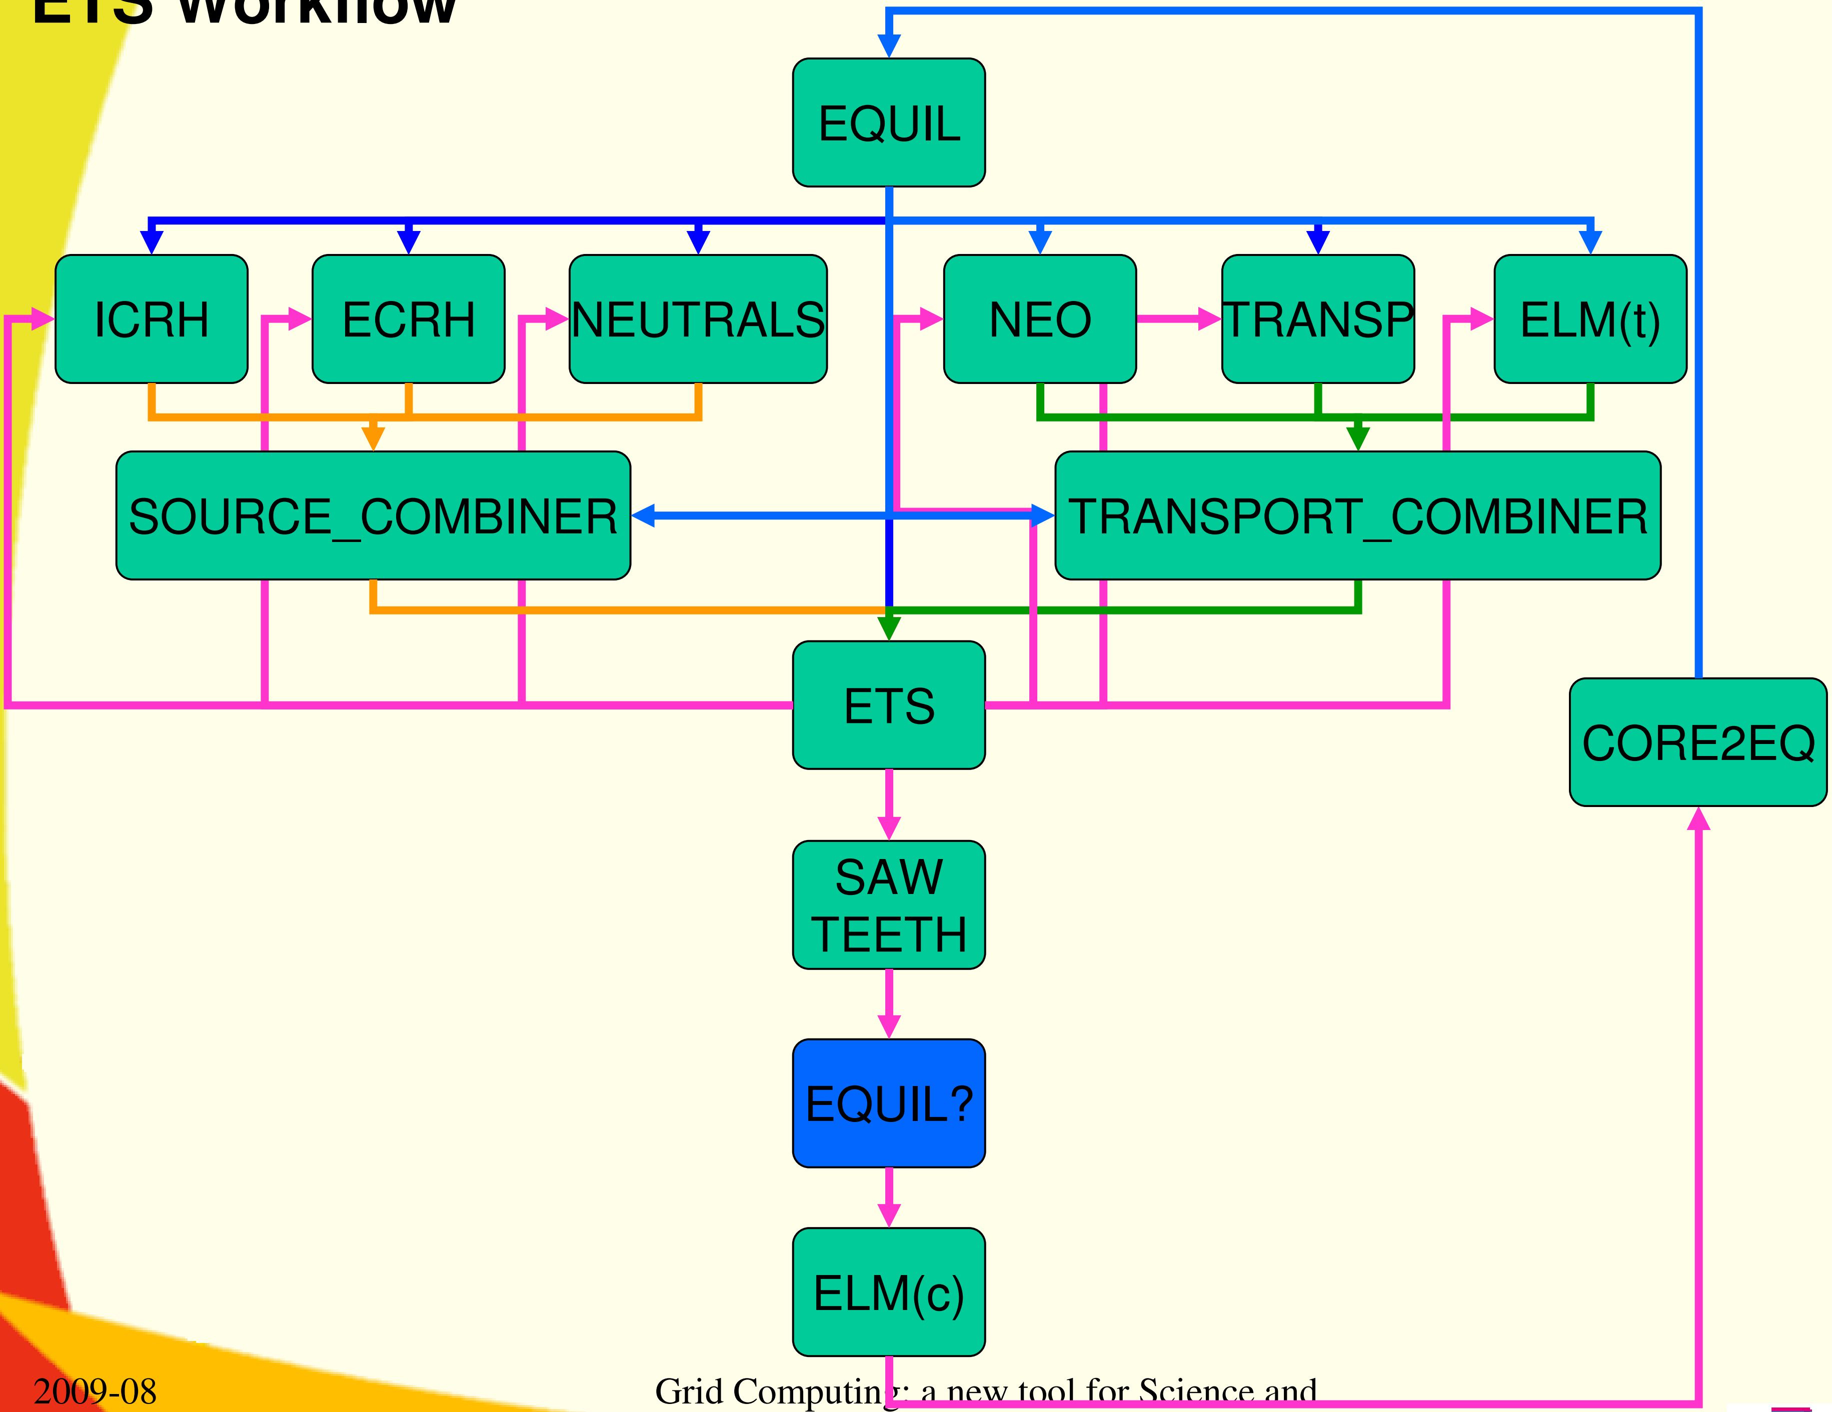

KEPLER Transport Workflow

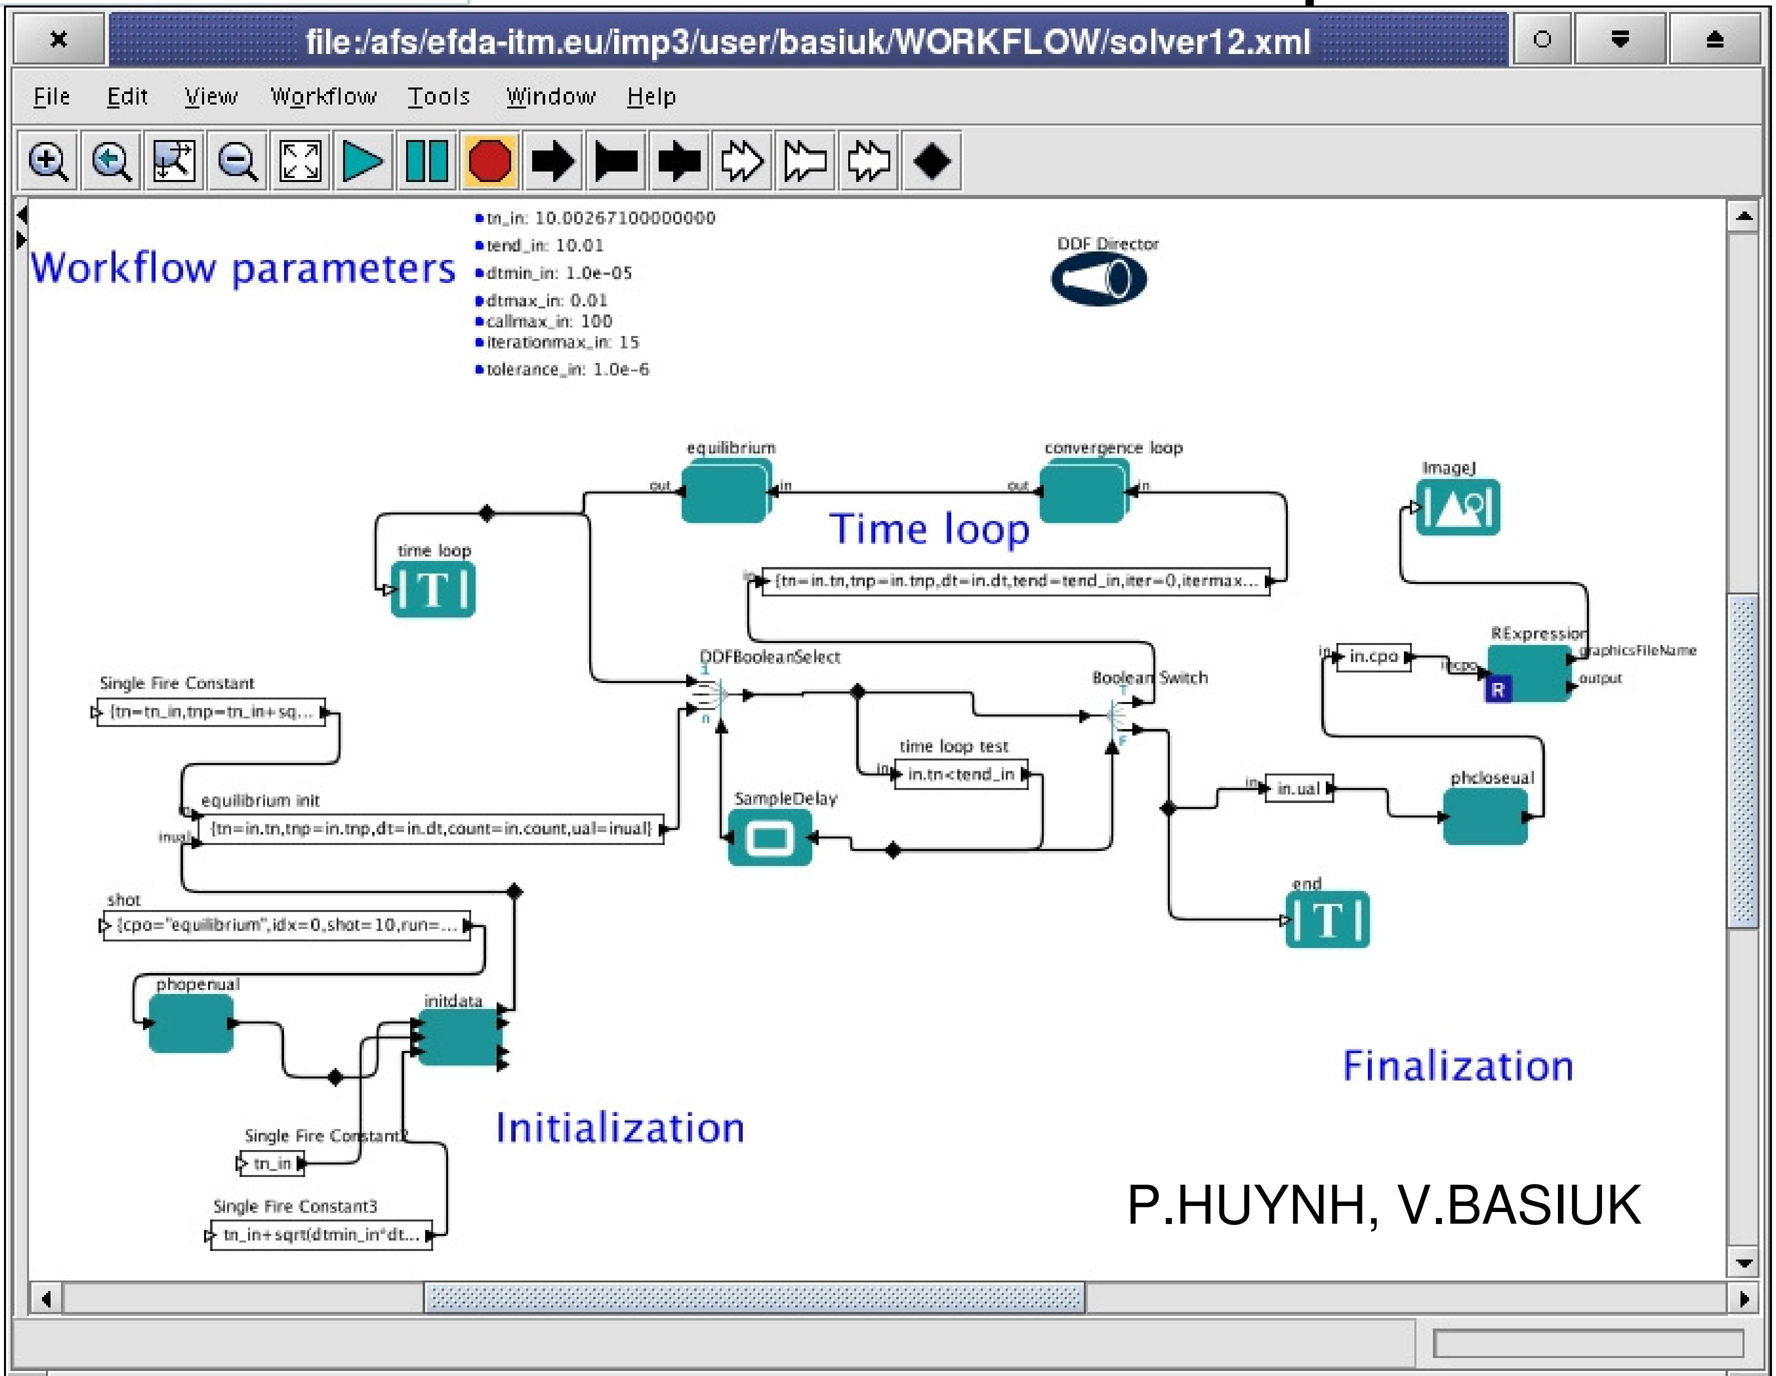

EUFORIA

14 member Institutes
3.65M€ over 36 months

522pms covering

- Management
- Training
- Dissemination
- Grid and HPC infra-structure & support
- Code adaptation & optimization
- Workflows
- Visualization

Consortium Members

Country	Institute	Capabilities
SWEDEN:	*CHALMERS University of Technology (coordinating)*	*Fusion, Grid, (CS)*
FINLAND:	*CSC - Tieteellinen laskenta Oy*	*HPC, (Grid),*
	Åbo Akademi University	*Code Optimization, CS*
FRANCE:	*CEA - Commissariat à l'énergie atomique – Cadarache*	*Workflow, Fusion, CS*
	Université Louis Pasteur	*Visualization, Applied Math*
GERMANY:	*Forschungszentrum Karlsruhe GmbH -FZK*	*Grid, Code parallelisation*
	Max-Planck-Institut für Plasmaphysik - IPP	*Fusion, (HPC, Grid)*
ITALY:	ENEA	Fusion, Grid, HPC, **GATEWAY**
SLOVENIA:	*University of Ljubljana -LECAD*	*Visualization, CS*
POLAND:	Poznan Supercomputing and Networking Centre	Grid, Migrating Desktop, CS
SPAIN:	*Barcelona Supercomputing Center – Centro Nacional de Supercomputación -BSC*	*HPC, Code optimization*
	Centro de Investigaciones Energéticas Medio Ambientales y Tecnológicas -CIEMAT	*Grid, Code parallelization, Fusion, Grid, NA*
	Consejo Superior de Investigaciones Cientificas -CSIC	*Grid, CS, (NA activities)*
UNITED KINGDOM:	*The University of Edinburgh - EPCC*	*HPC, Code Optimization, NA, User support, (GRID)*

Project objectives

Bringing a comprehensive framework and e-infrastructure to the fusion modelling community oriented to the development of ITER physics needs with particular emphasis on Grid and HPC activities

Focus area: Edge and core transport and turbulence

- Deployment of Grid and HPC infrastructure
- Adaptation and Optimization of Fusion Codes
  - Platform oriented Grid and/or HPC
- Development of advanced tools for
  - Workflow management
  - Visualization tools
  - Data mining

In addition, there is a number of outreach and dissemination activities planned to introduce the fusion community at large to the developed infrastructure and make contact with other infrastructure and research projects with similar or associated orientations.

Work plan outline

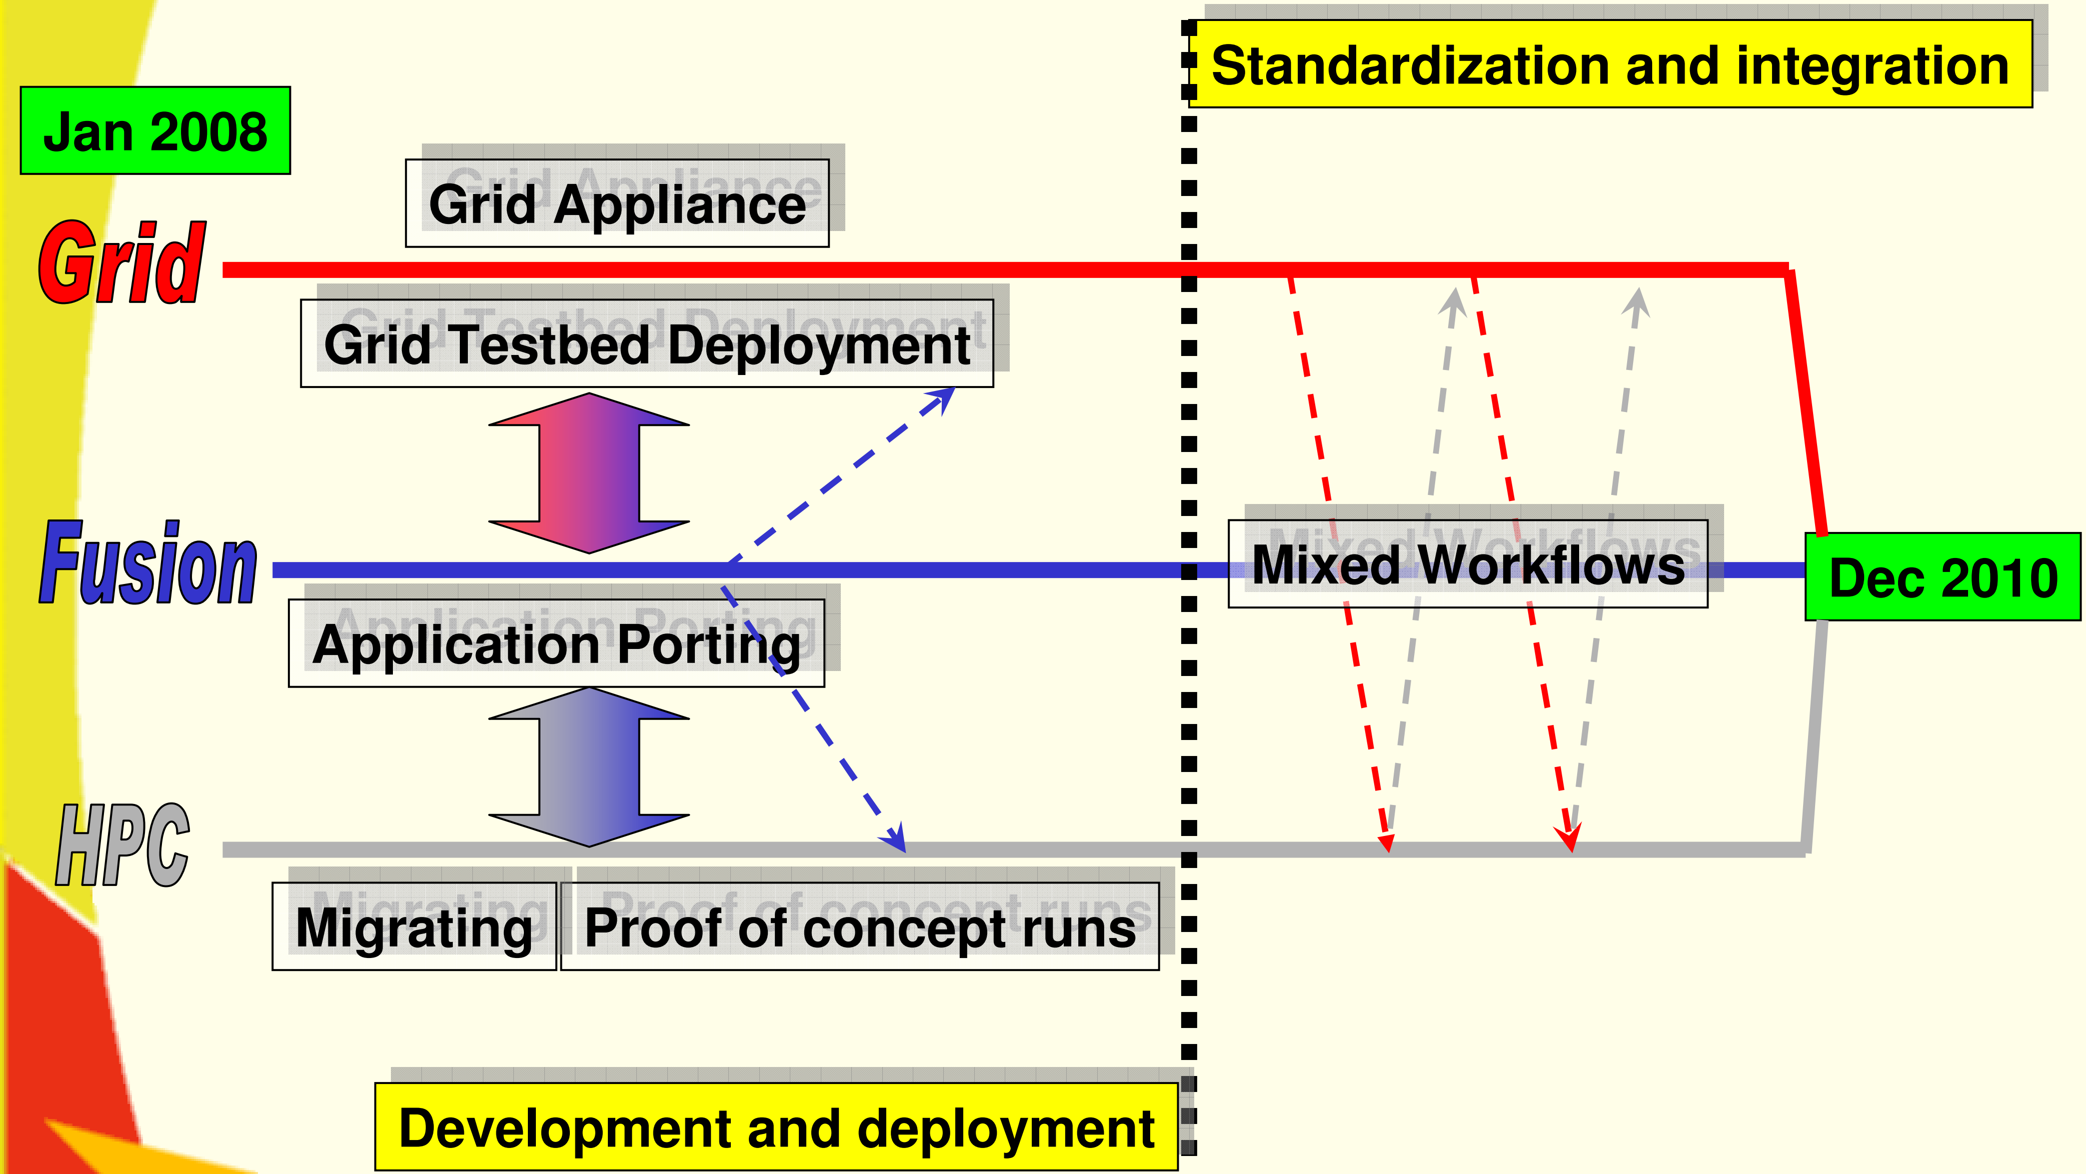

Team - Work Package Managers

The EUFORIA programme is organized under three headings: Networking (NA), Service activities (SA) and Joint Research Activities (JRA) and 10 WP's.

	WP #	# of pm's allocated	Work package title	Lead participant short name	WP managers
Networking	NA1	16	Management	Chalmers	Pär Strand (PSG)
	NA2	25	User Documentation and Training	UEDIN	Lorna Smith (PSG)
	NA3	24	Dissemination	CIEMAT	Miguel Cardenas
Service Activities	SA1	115	Grid Infrastructure Deployment and Operation	CSIC	Isabel Campos (PSG)
	SA2	25	HPC Infrastructure	BSC	Jose Maria Cela (PSG)
	SA3	31	HPC and Grid Support	UEDIN	Lorna Smith
Joint Research Activities	JRA1	72	Adaptation of codes for Grid Infrastructure	CIEMAT	Francisco Castejon
	JRA2	110	Adaptation of codes for HPC infrastructure	Åbo	Jan Westerholm (PSG)
	JRA3	68	Workflow orchestration - physics integration support	CEA	Bernard Guillerminet (PSG)
	JRA4	36	Visualization	STRASB	Eric Sonnendrucker

Workflows

- The objective is to schedule jobs on the GRID and HPC infrastructures together with jobs running on other computing facilities.
  - launch and control jobs in a transparent manner for the users.
  - data communication: data transfers and visualizations are required at run time and to access the experimental and simulated data.
  - middleware connectivity for Kepler Workflow (java Library for gLite, HPC)
  - scheduler notification tools
  - Integration with ITM tools (UAL – Universal Access Layer)

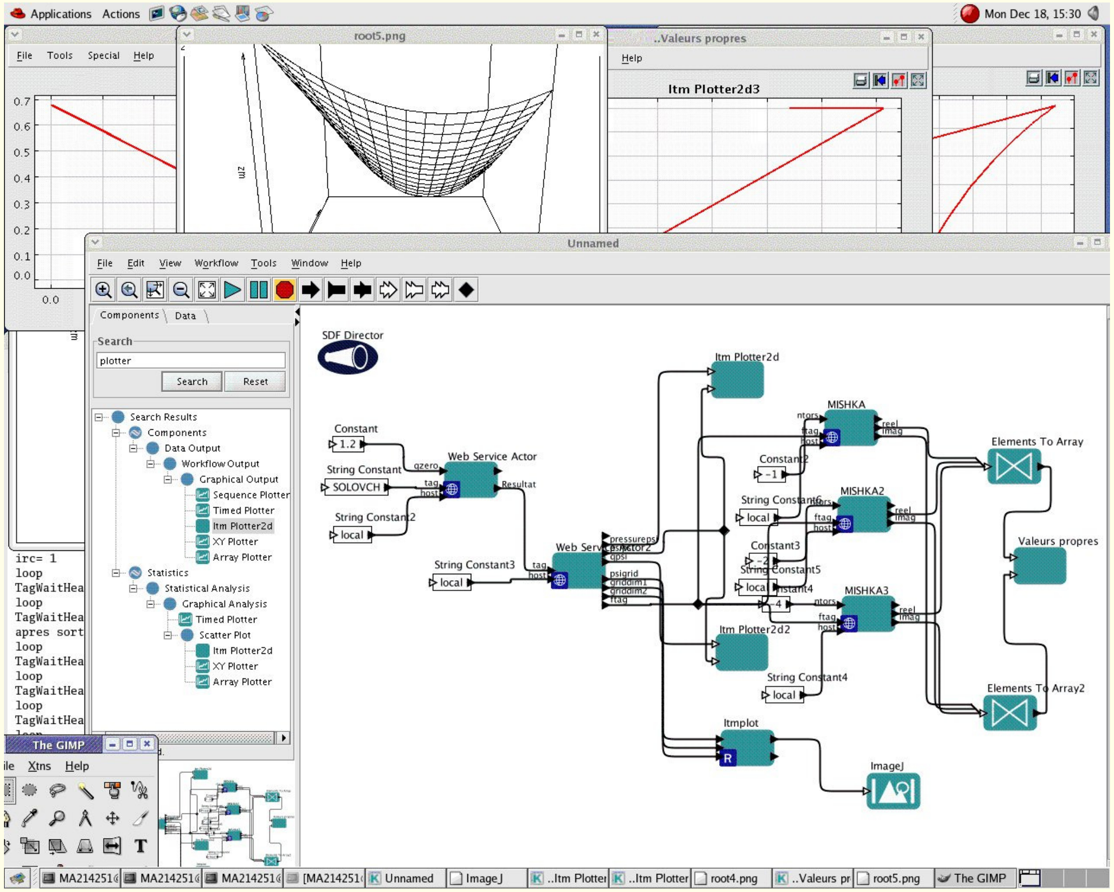

Visualization

- Development of post-processing data analysis and visualization tools based on Python, NumPy, and Matplotlib packages.
- Development of post-processing data analysis and visualization tools based on VisIt.
- Integration of Python, NumPy, Matplotlib into the Kepler workflow.
- Integration of VisIt or VTK into the Kepler workflow.
- Development of a lossy wavelet based compressed data format.
- Develop visualization techniques for 4D and 5D particle distributions.

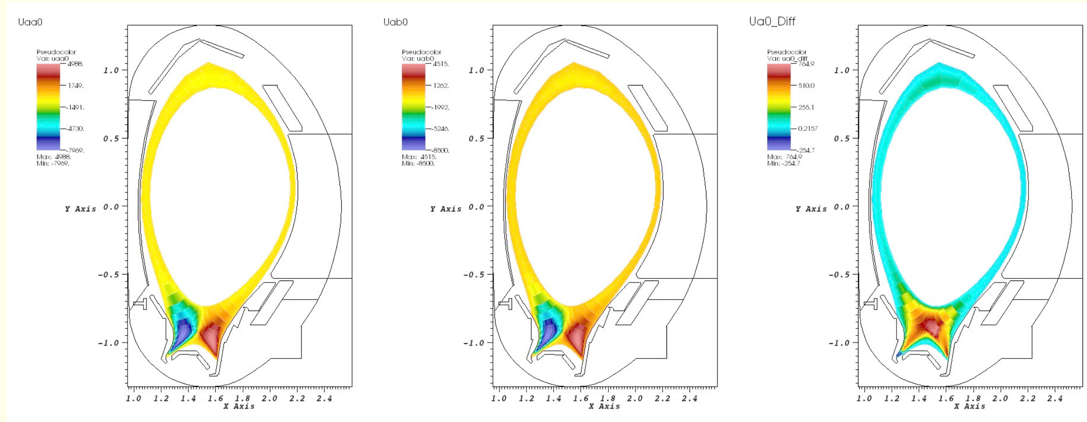

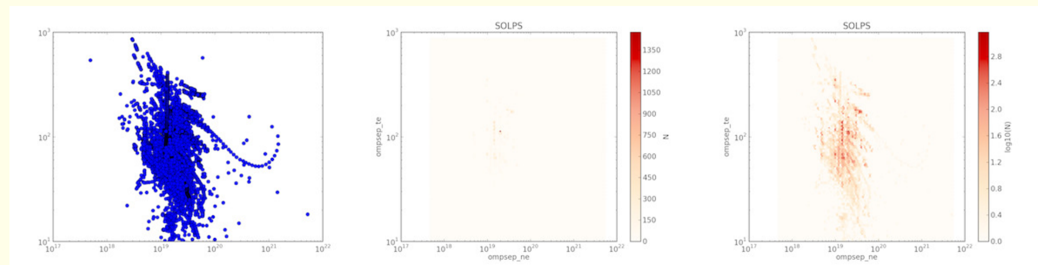

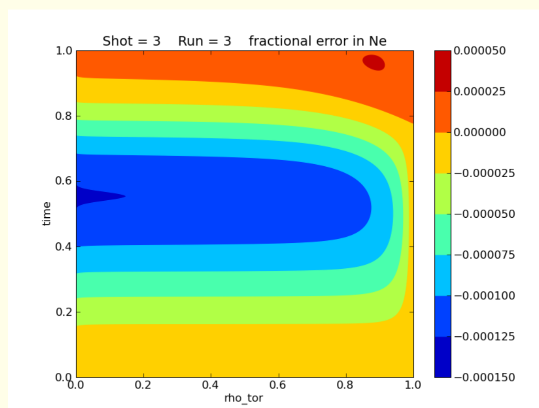

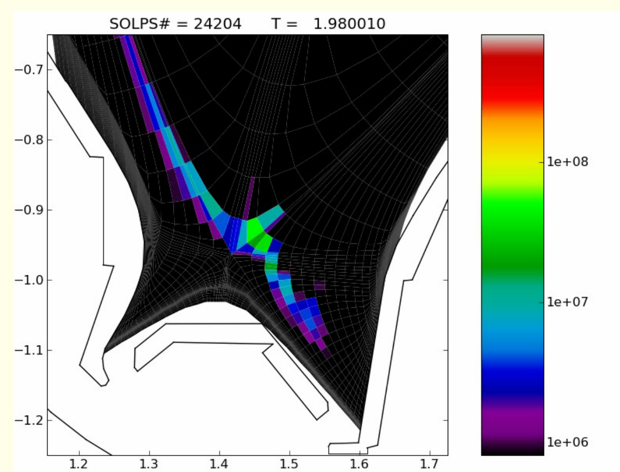

EU Fusion "Ontology"

- Abstracted data structures
  - Description of a "Complete" set of data for describing plasma operation and simulations
    - Abstracted through XML schemas
    - Unambiguous description with agreed sign and other conventions
    - SI-units (with eV)
    - Consistent Physical Objects (CPO)
      - Groupings of related data - basis for code interfaces
    - Machine descriptions
      - Allowing for device independent codes!
  - Serializations
    - Transformations providing
      - Language specific implementations of CPOs (f90, C++,..)
      - Database structures
  - Used in workflow tool to connect modules.
    - Automagic creation of Kepler Actors

EU Fusion “Ontology” (II)

- Access and Storage
  - Universal Access Layer (UAL) providing invariant API based on CPOs
  - Plug-in API for backend
    - MDS+ implemented - implications on data structures
    - HDF5 next step
- Code Coupling Mechanism
  - Used in workflow tool to connect modules.
- NB: The tools and structures are generic, the EUFORIA toolset is portable to other modelling areas

EUFORIA in brief

- 3.65M€, 36 months, 522 pm's
- Coordinated by Chalmers (IPP deputy)
  - ITM-TF connectivity to guarantee broad access
  - Intended to boost ITM-TF activities towards HPC and grid computing
- "Edge and core turbulence and transport"
- Deployment of Grid and HPC infrastructure
- Adaptation and Optimization of Fusion Codes
  - Platform oriented Grid and/or HPC
  - First profiling and priorisation of codes done.
- Development of advanced tools for
  - Workflow management
  - Visualization tools
  - Data mining
  - Range of outreach and training activities

Developing a new paradigm for fusion computing

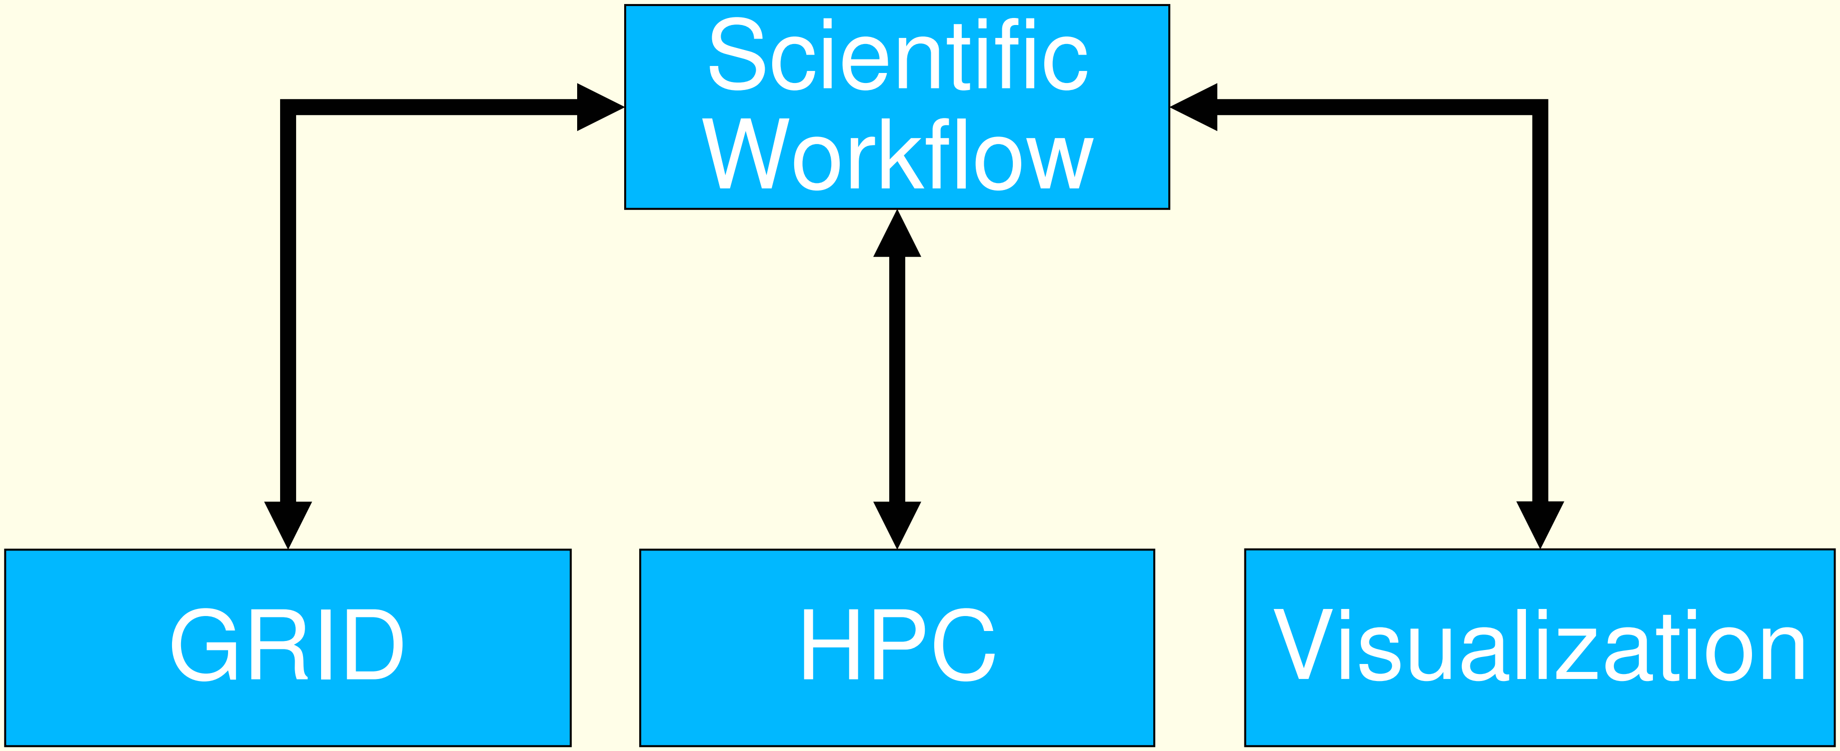

How is EUFORIA Working towards this Goal?

Scientific Workflow	• JRA3
GRID	• JRA1, SA1 • EGEE, EGI
HPC	• JRA2, SA2 • DEISA
Visualization	• JRA4

But we also need to "sell" this paradigm to the fusion community

•NA2, NA3

Batch, GRID and HPC

	Batch	GRID	HPC
Current usage within the fusion community	Extensive • JET (339 cpus) • Most associations (IPP-TOK 400 cpus) • Gateway (128 cpus)	• Only at CIEMAT • Starting to see some usage at Juelich • Plans for usage at IPP	• Extensive use of national facilities • Extra-national usage via DEISA
Main issue preventing more usage	• Lack of more local cpus	• Lack of knowledge • Difficulties getting certificates	• Difficult for Associations wothout strong national facilities
Currently being addressed by		EUFORIA • has provided training • also developing facilities to launch jobs from the Gateway	EUFORIA • has provided training • has worked on codes HPC-FF starting August 1st • HLST
Longer term		?	IFERC computer (2012)

Developing a new paradigm for fusion computing

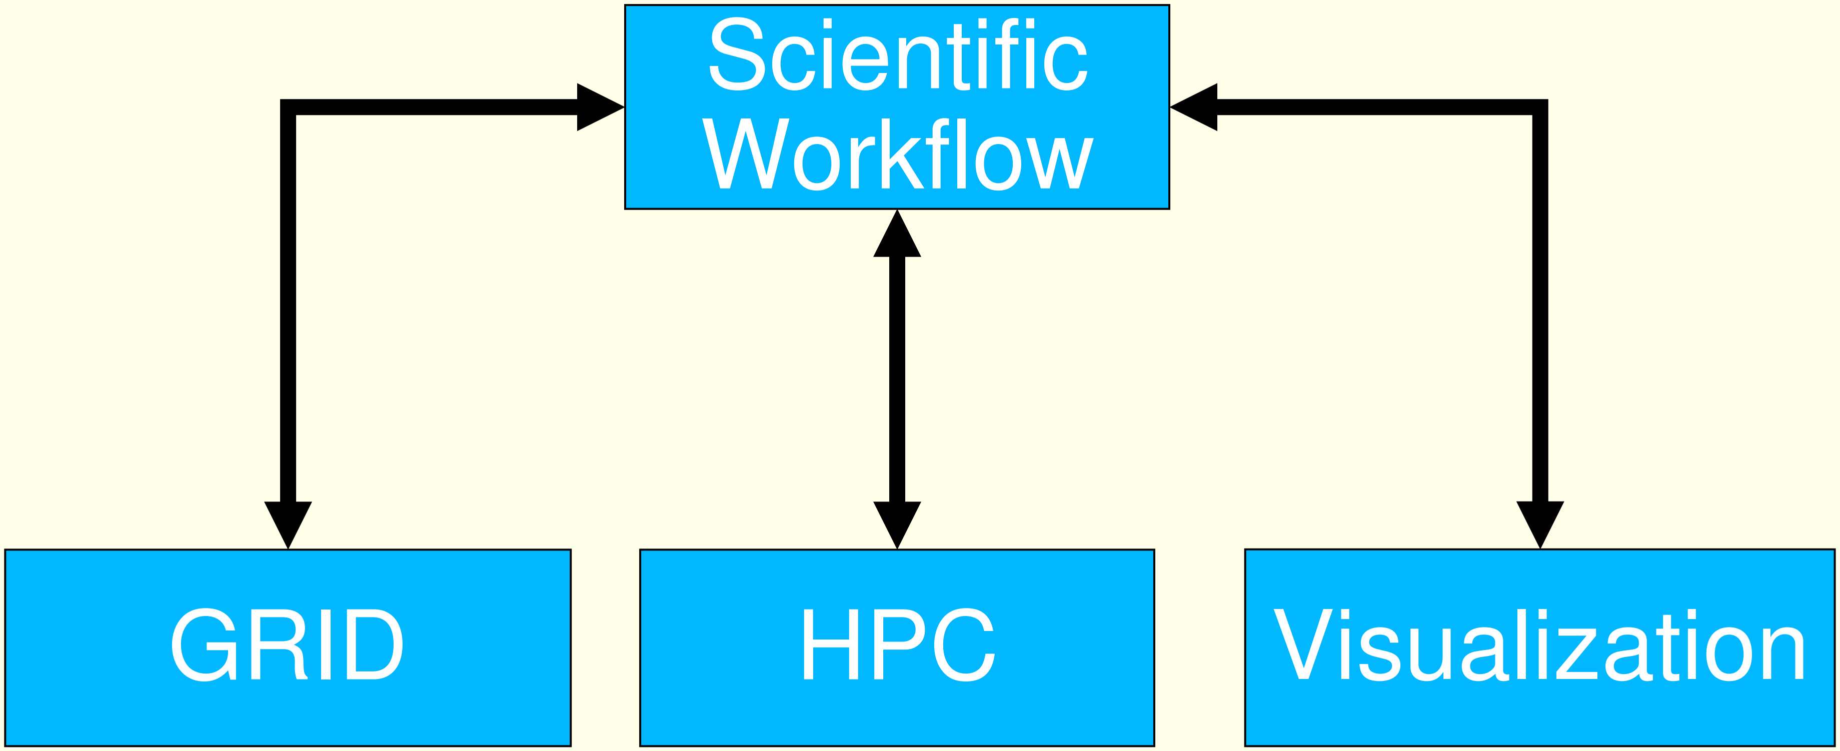

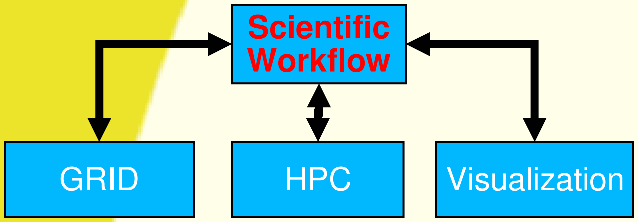

Scientific Workflows

- Need to make a convincing case that scientific workflows will
  - Enhance productivity
  - Allow for new approaches problems to be solved
  - Allow for traceability, reproducibility, …
  - Allow for a better use of resources

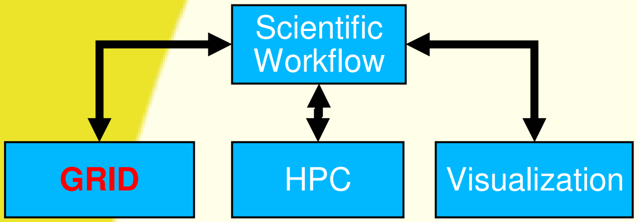

Grid Computing

- Grid computing needs
  - To be transparently coupled into the scientific workflows
  - Needs to be reliable (every launched job should run)
  - Needs to improve performance (if resources are available, a launched job should start rapidly)
  - To inter-operate with the other levels at the data access level

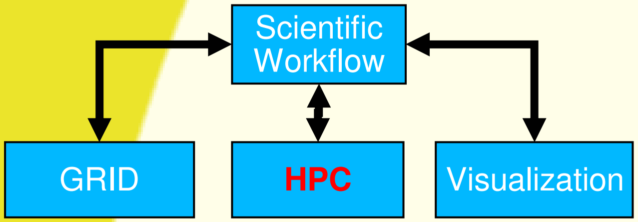

HPC

- Need to be transparently coupled into the workflows
- HPC facilities will need to deal with jobs coming from workflows
  - Negotiations about resource availability and expected response times
  - Deal with communications between the different parts of a workflow
  - Better integrate inter-operability of data between HPC and the other levels

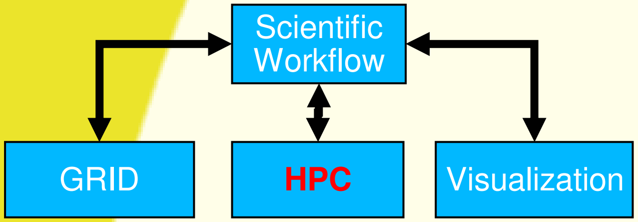

HPC, II

- Need to better prepare jobs to scale from the present 1k – 10k cpus to 100k – 1M cpus
  - This means at the development stage
    - Having machines exposing large numbers of cores to code developers interactively or almost interactively (< 5 minute turn around)
    - Helping the code developers with algorithms, …
    - Helping the code developers with better IO strategies
  - At the production stage
    - Having the resources available

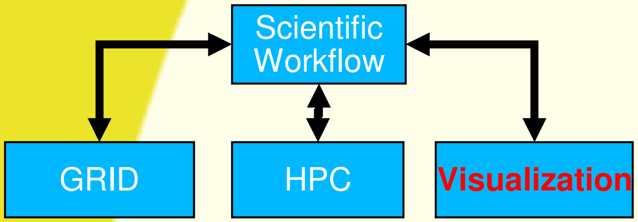

Visualization

- Challenges in visualization include
  - Visualizing the data flowing around a workflow so that the scientist can monitor/diagnose a running job
  - Deal with very large amounts of data produced in a distributed environment
  - Help provide the scientist with a better understanding of his/her results
  - Help the scientist by producing visualizations with that “Wow!” factor

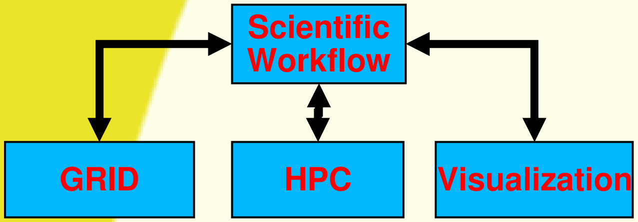

Over all

- Need to make it simpler to move data between the different parts of a workflow
  - Remote data access
  - High speed data transfers
  - Better integration
- Need to think now about very large data sets
  - What is the “best” relationship between petaflops and petabytes?

What has EUFORIA demonstrated?

- Ability to launch a GRID job from Kepler and then retrieve the results
- Ability to launch a HPC job from Kepler and then retrieve the results
- Ability of a GRID job to write results to the Gateway using the UAL
- Deployment of applications on the EUFORIA GRID
  - **EMC3-EIRENE partially shifting to Grid (TEXTOR DED)**
- Succesful support/optimization on HPC
  - BIT1 (sequential ➔ 512-1024 cores, DEISA time)
  - TYR (128 ➔ 4096 cores)
  - Building a framework for supporting users
- Coordinating and facilitating contacts with other activities
  - EFDA (ITM-TF, HPC-FF), DEISA, EGEE, PRACE,…
  - ITER (informal)