EFDA
EUROPEAN FUSION DEVELOPMENT AGREEMENT

Frank Jenko

Extreme Computing in support of ITER

Max Planck Institute
for Plasma Physics
Garching, Germany

Ulm University, Germany

DEISA PRACE Symposium 2009

Distributed
European
Infrastructure for
Supercomputing
Applications

PRACE

HPC Infrastructures for Petascale Applications
11-13 May, 2009, Amsterdam

The ITER project:
Some background

How the Sun shines: Fusion energy

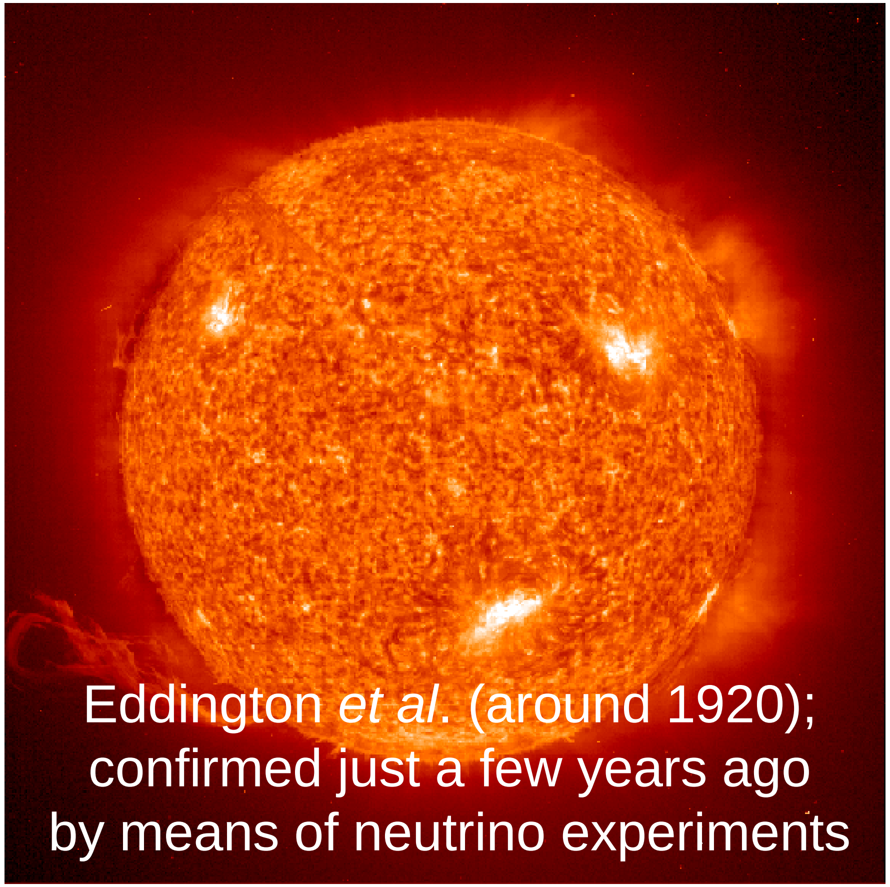

Eddington *et al*. (around 1920);
confirmed just a few years ago
by means of neutrino experiments

Fusion of 4 protons
in several stages

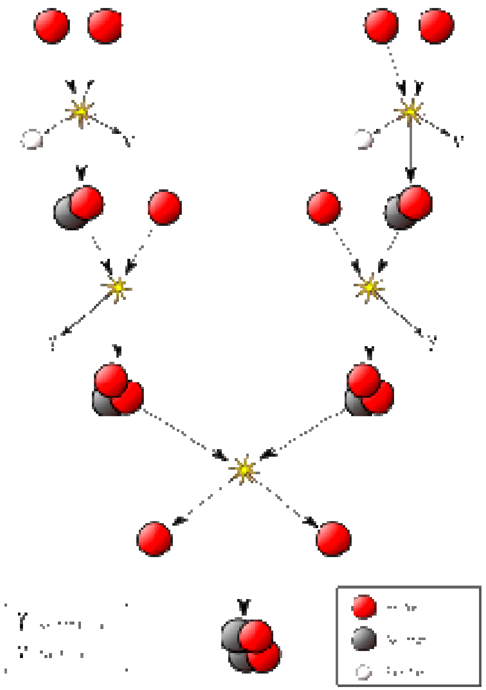

Idea: New source of CO_2 free energy on Earth for the 21st century and well beyond

Fusion energy in the laboratory

This process has by far the highest reaction rate under experimentally accessible conditions:

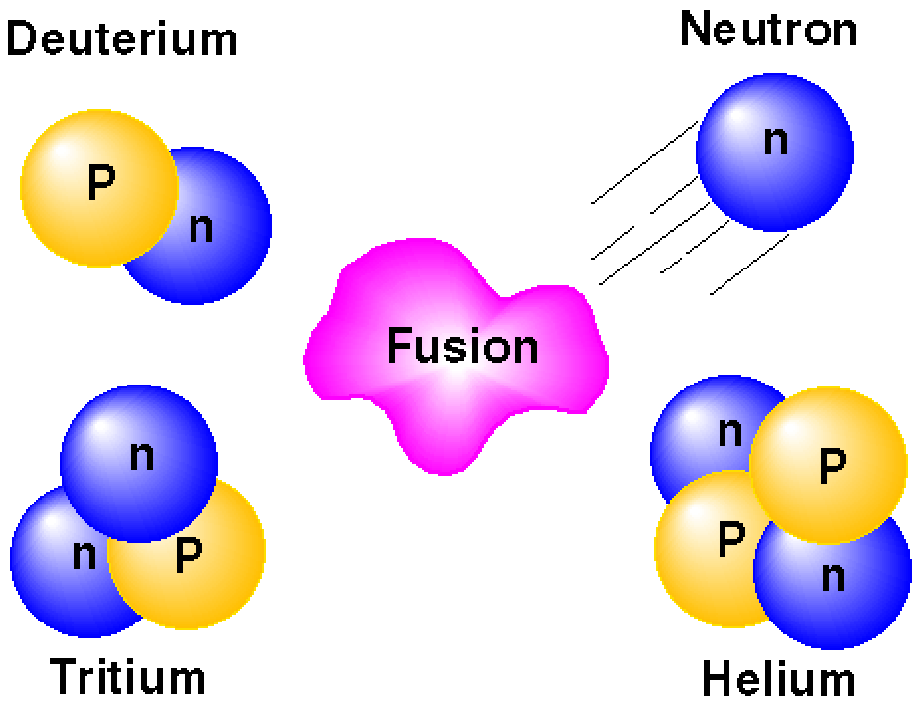

Still, temperatures of about 100 million degrees are required!
Thus, we are dealing with a fully ionized gas (plasma).

Magnetic confinement of plasmas

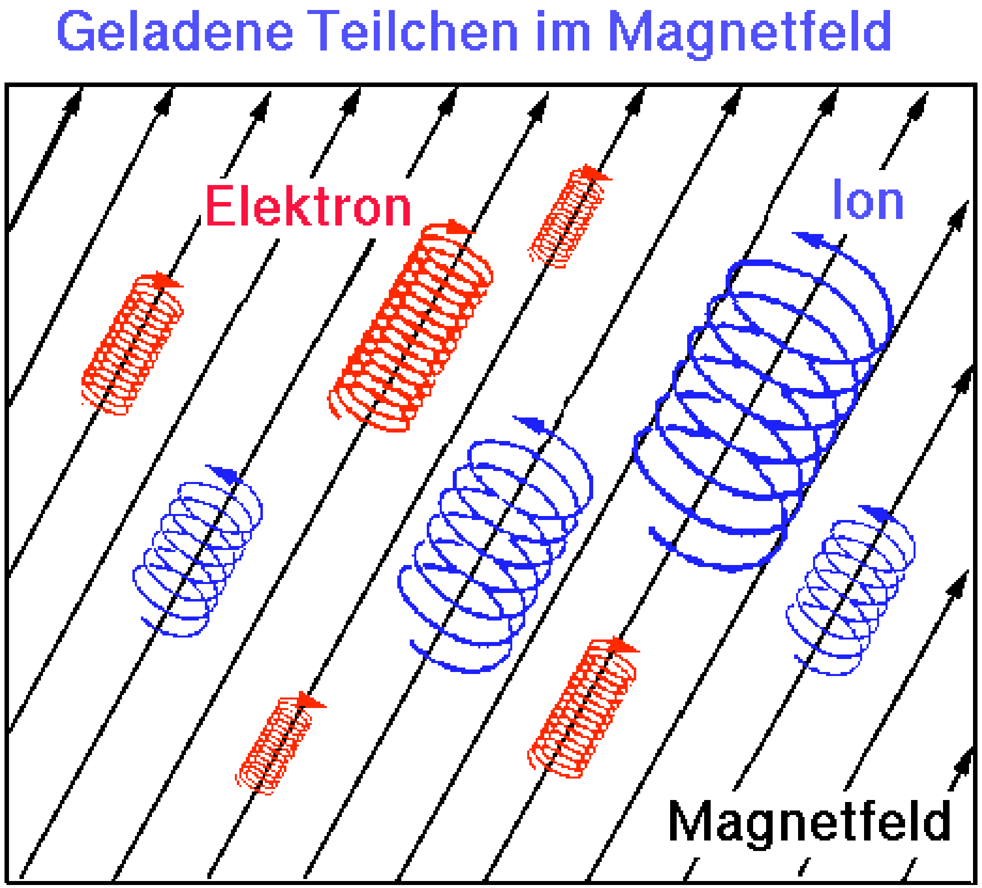

- ➢ electrically charged particles spiral around magnetic field lines
- ➢ perpendicular (not parallel) confinement
- ➢ bend magnetic field lines to a torus

Most advanced concept: Tokamak

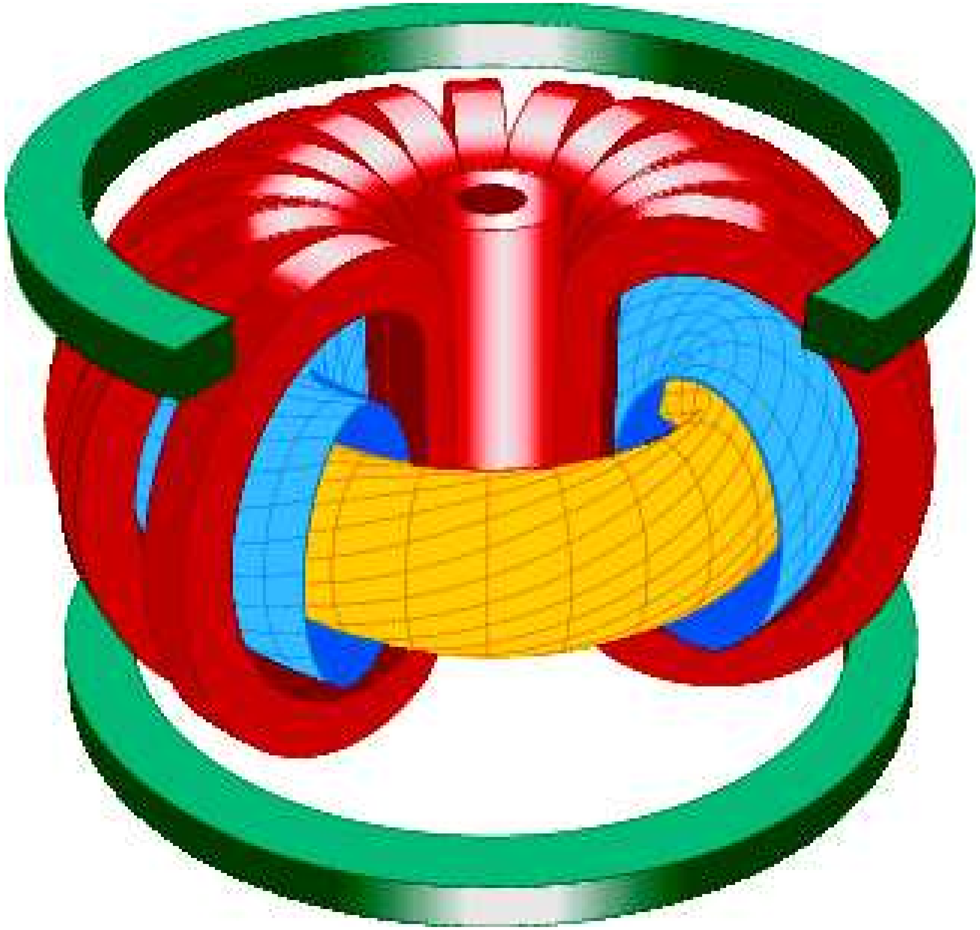

- Axisymmetric torus
- Magnetic field generation by a combination of external coils and internal currents
- Magnetic field lines span nested toroidal surfaces
- Plasma particles more or less follow the field lines

Fusion research: Towards ignition

α heating must compensate energy losses:

- **Electromagnetic radiation**
- **Turbulent transport**

Key requirements:

- **Large central pressure (limited by onset of large-scale instabilities)**

- **Large energy confinement time (limited by small-scale instabilities, i.e. turbulence)**

$$\tau_E = E_{plasma} / P_{loss}$$

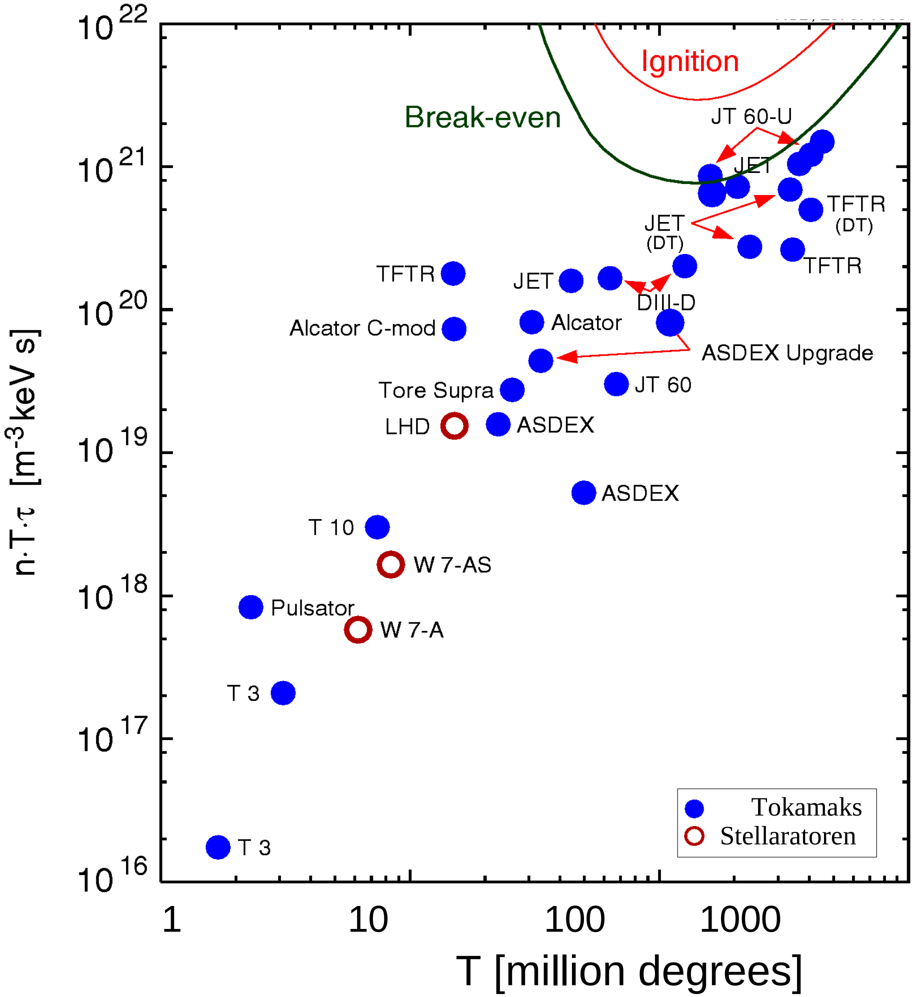

The JET device (operated by EFDA)

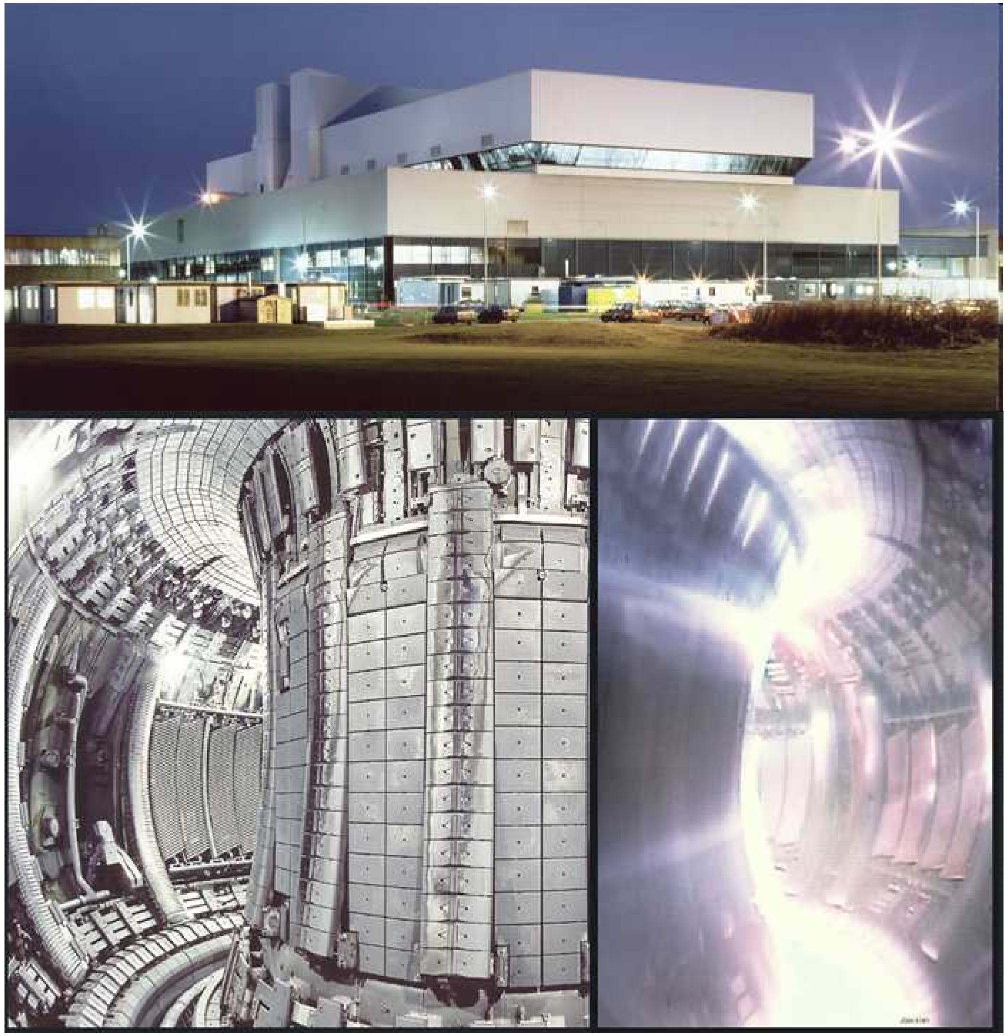

Joint **E**uropean **T**orus

Used by more than 20 European Countries

Currently the world’s largest fusion experiment

Located near Oxford, UK

World record: 16 MW of fusion power (1997)

ITER: The final step towards DEMO (a demonstration fusion power plant)

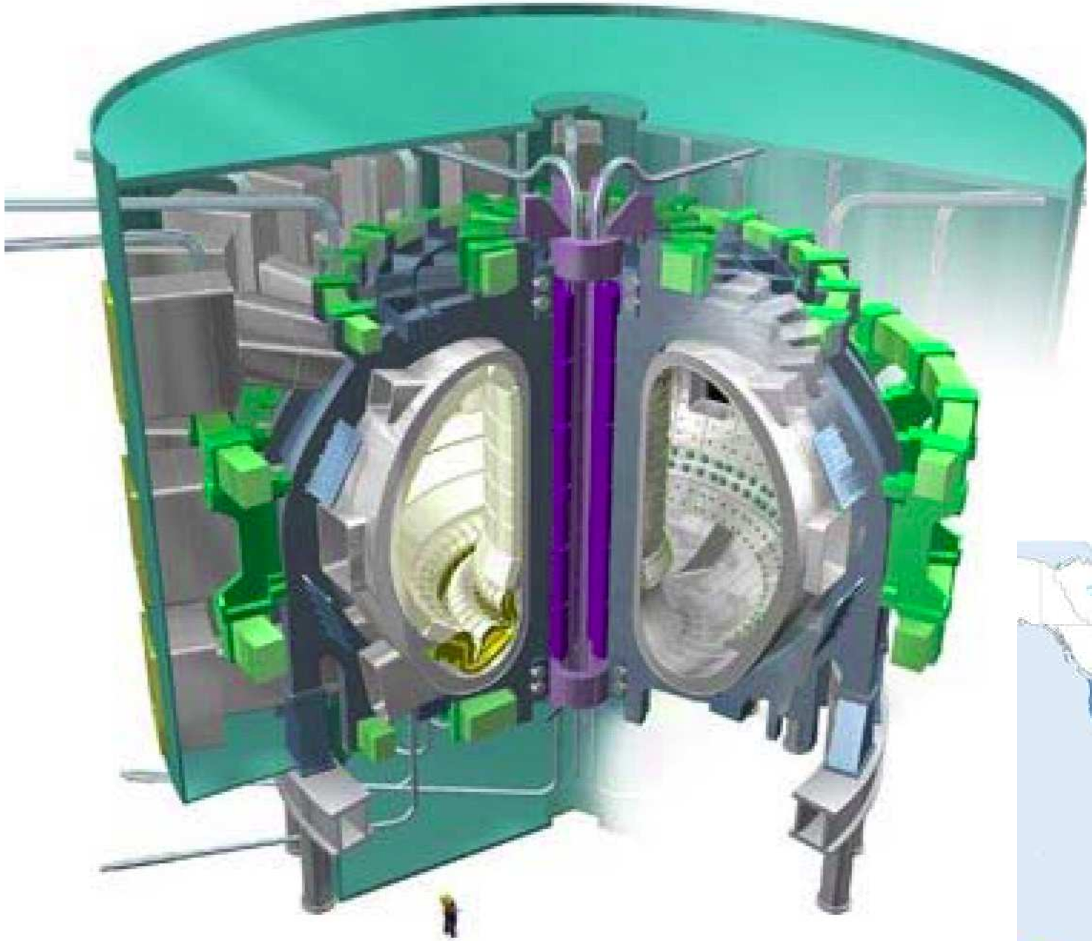

Goal: 500 MW
of fusion power

www.iter.org

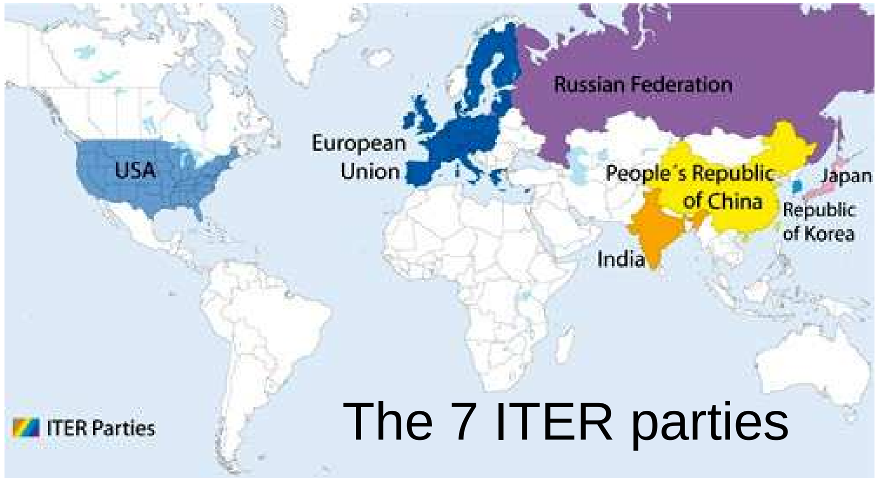

The 7 ITER parties

ITER is currently being built in Cadarache, Southern France

Operation is expected to start in 2018

Extreme Computing in support of ITER

3 key challenges for fusion physics

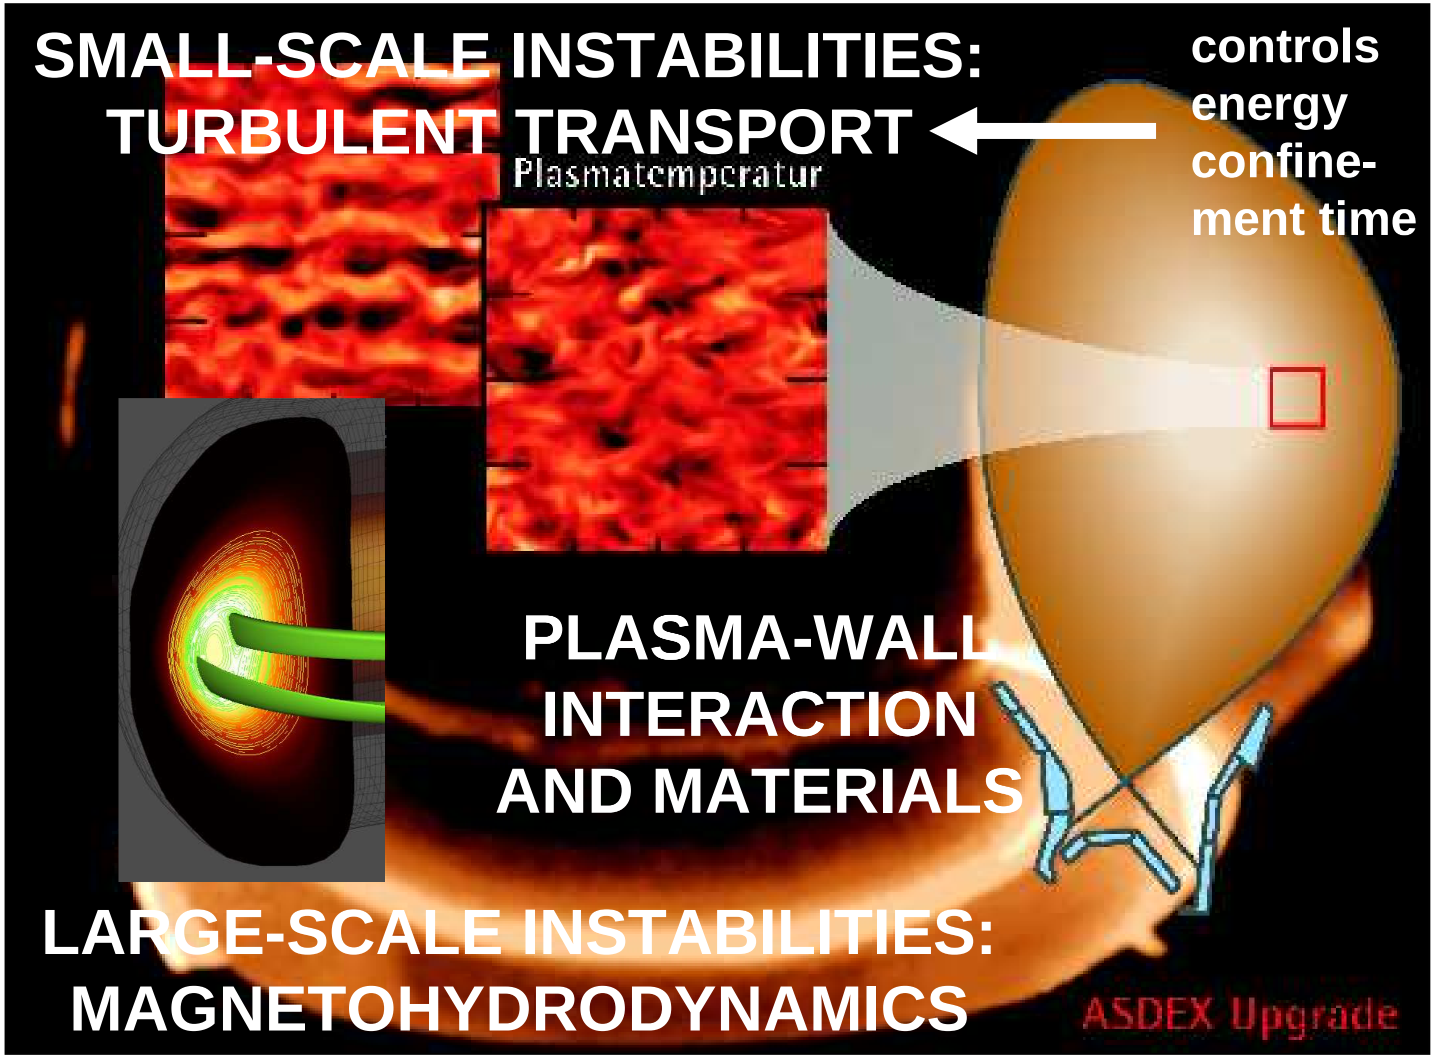

The huge range of spatial and temporal scales presents major challenges to theory and simulation

- Overlap in scales often means strong (simplified) ordering is not possible
- Effective simulations at the petascale and beyond are required to meet FES mission needs in dealing with burning plasma (BP)/ITER science challenges

J. Roberto

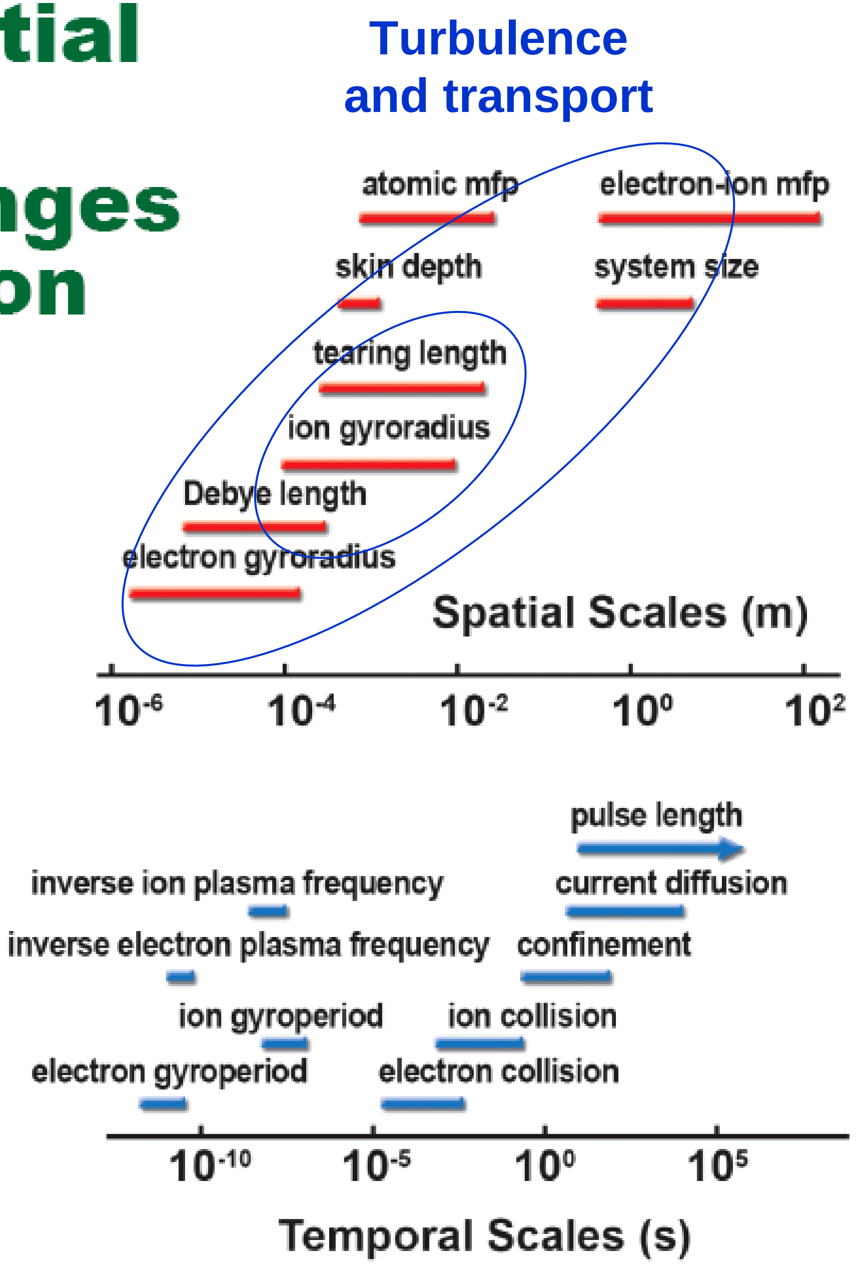

Turbulence – one of the most important unsolved problems in physics

According to a famous statement by Richard Feynman…

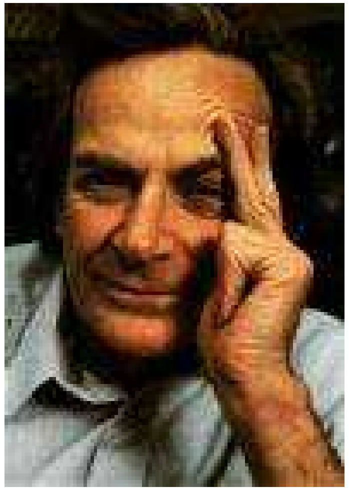

…and a survey by the British “Institute of Physics” among many of the leading physicists world-wide…

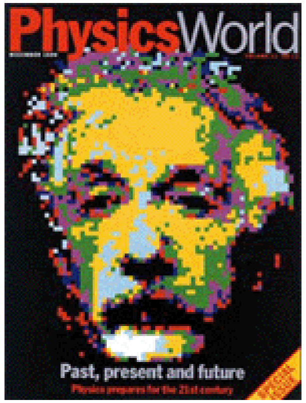

“Millennium Issue”
(December 1999)

TURBULENCE:
A challenging topic for both basic and applied research

What is turbulence?

Turbulence…

Leonardo
da Vinci
(1529)

- is a nonlinear phenomenon
- occurs (only) in open systems
- involves many degrees of freedom
- is highly irregular (chaotic) in space and time
- often leads to a (statistically) quasi-stationary state far from thermodynamic equilibrium

These properties make it a very complicated problem – neither Dynamical Systems Theory nor Statistics applies!

How to approach turbulence?

Many physicists – including Heisenberg, von Weizsäcker, Onsager, Feynman, and many others – have attempted to tackle turbulence **purely analytically** but with only **very limited success**.

Today, **supercomputers** help to unravel the “mysteries” of turbulence in the spirit of **John von Neumann**:

„There might be some hope to 'break the deadlock' by extensive, but well-planned, computational efforts...“

Reduced kinetic description

Dilute and/or hot plasmas are almost collisionless.

Therefore, (3D) fluid theory is not applicable,
and one has to use a (reduced) kinetic description!

Vlasov-Maxwell equations

$$\left[\frac{\partial}{\partial t}+\mathbf{v}\cdot\frac{\partial}{\partial \mathbf{x}}+\frac{q}{m}\left(\mathbf{E}+\frac{\mathbf{v}}{c}\times\mathbf{B}\right)\cdot\frac{\partial}{\partial \mathbf{v}}\right]f(\mathbf{x},\mathbf{v},t)=0$$

Removing the fast gyromotion

[Frieman, Chen, Lee, Hahm, Brizard *et al.*, 1980s]

$\omega \ll \Omega$

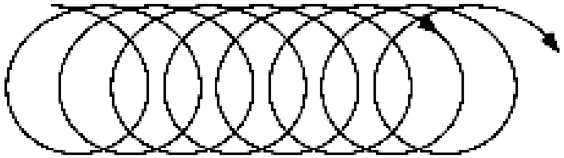

Charged rings as quasiparticles;
gyrocenter coordinates

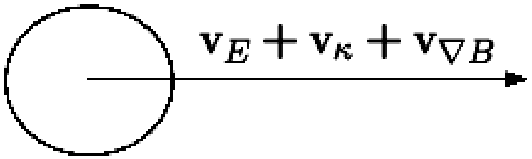

Nonlinear integro-differential equations in **5 dimensions**…

The nonlinear gyrokinetic equations

$$f = f(\mathbf{X}, v_\parallel, \mu; t)$$

Advection/Conservation equation

$$\boxed{\frac{\partial f}{\partial t} + \dot{\mathbf{X}} \cdot \frac{\partial f}{\partial \mathbf{X}} + \dot{v}_\parallel \frac{\partial f}{\partial v_\parallel} = 0}$$

$$\dot{\mathbf{X}} = v_\parallel \mathbf{b} + \frac{B}{B_\parallel^*} \left(\frac{v_\parallel}{B} \bar{\mathbf{B}}_{1\perp} + \mathbf{v}_\perp \right)$$

$$\mathbf{v}_\perp \equiv \frac{c}{B^2} \bar{\mathbf{E}}_1 \times \mathbf{B} + \frac{\mu}{m\Omega} \mathbf{b} \times \nabla (B + \bar{B}_{1\parallel}) + \frac{v_\parallel^2}{\Omega} (\nabla \times \mathbf{b})_\perp$$

$$\dot{v}_\parallel = \frac{\dot{\mathbf{X}}}{m v_\parallel} \cdot \left(e \bar{\mathbf{E}}_1 - \mu \nabla (B + \bar{B}_{1\parallel}) \right)$$

X = gyrocenter position
$v_\parallel$ = parallel velocity
μ = magnetic moment

Appropriate field equations

$$\frac{n_1}{n_0} = \enclose{circle}{\frac{\bar{n}_1}{n_0}} - \left(1 - \| I_0^2 \| \right) \frac{e\phi_1}{T} + \| x I_0 I_1 \| \frac{B_{1\parallel}}{B}$$

$$\nabla_\perp^2 A_{1\parallel} = -\frac{4\pi}{c} \sum \enclose{circle}{\bar{J}_{1\parallel}}$$

$$\frac{B_{1\parallel}}{B} = -\sum \epsilon_\beta \left(\frac{\enclose{circle}{\bar{p}_{1\perp}}}{n_0 T} + \| x I_1 I_0 \| \frac{e\phi_1}{T} + \| x^2 I_1^2 \| \frac{B_{1\parallel}}{B} \right)$$

Turbulent fluctuations are quasi-2D

Reason: Strong background magnetic field

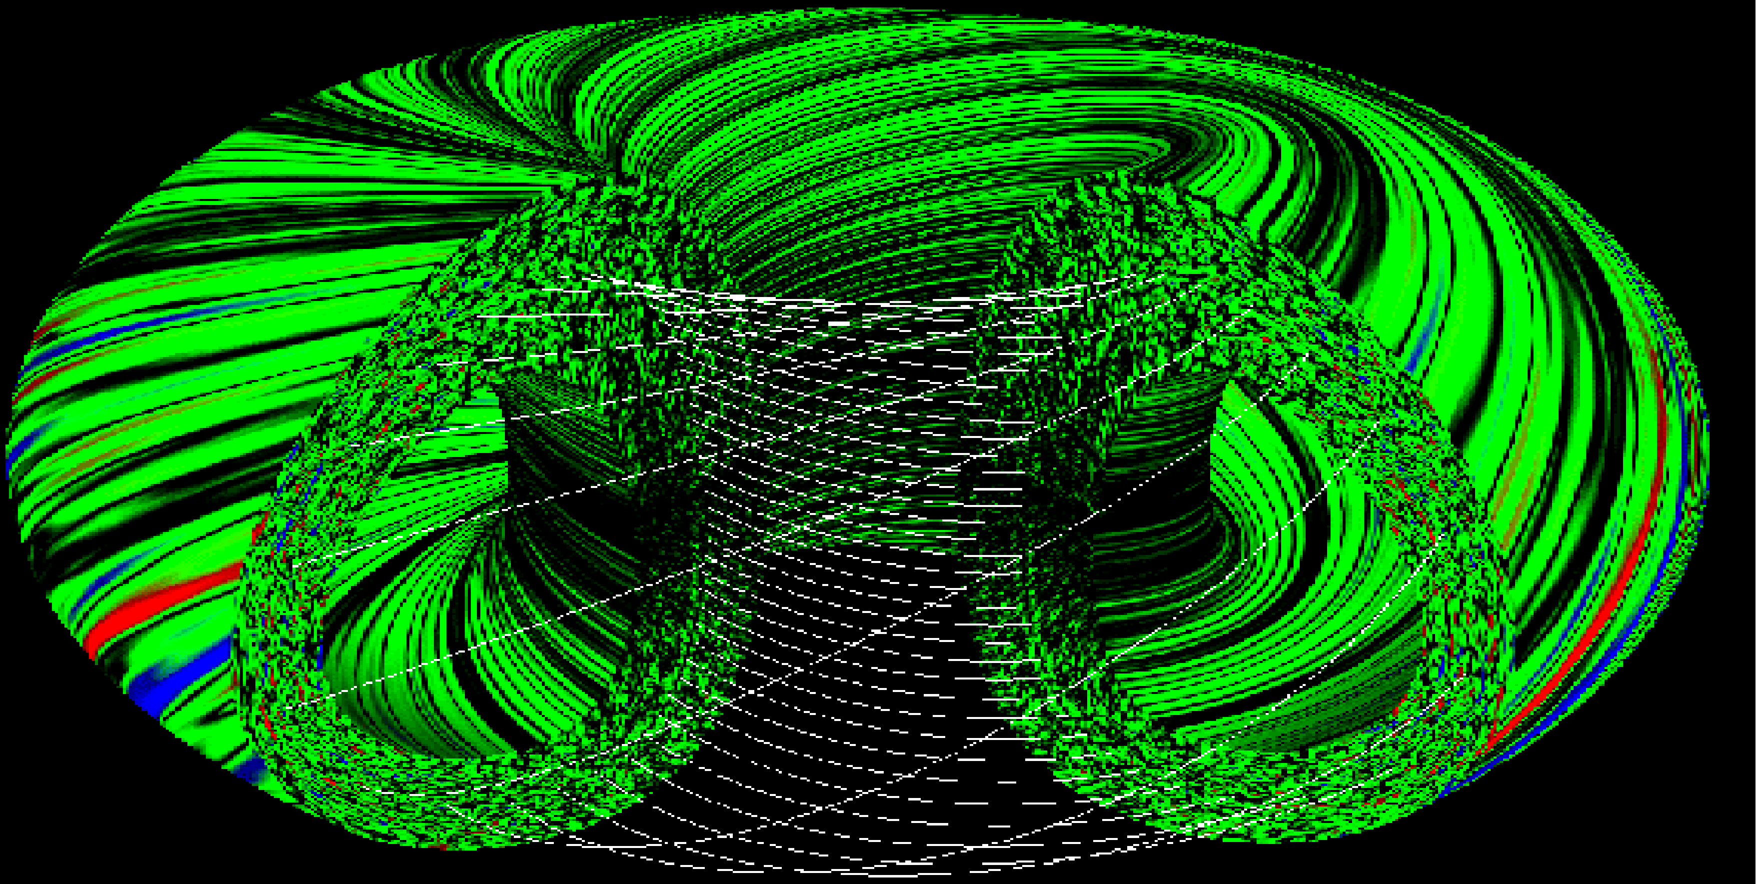

Possible simulation volume: flux tube, annulus, full (or fractional) torus

G.Hammett

Major theoretical speedups

relative to original Vlasov/pre-Maxwell system on a naïve grid, for ITER $1/\rho_* = a/\rho \sim 1000$

Nonlinear gyrokinetic equations	
eliminate plasma frequency: $\omega_{pe}/\Omega_i \sim m_i/m_e$	x10^3
eliminate Debye length scale: $(\rho_i/\lambda_{De})^3 \sim (m_i/m_e)^{3/2}$	x10^5
average over fast ion gyration: $\Omega_i/\omega \sim 1/\rho_*$	x10^3
Field-aligned coordinates	
adapt to elongated structure of turbulent eddies: $\Delta_{\|}/\Delta_{\perp} \sim 1/\rho_*$	x10^3
Reduced simulation volume	
reduce toroidal mode numbers (i.e., 1/15 of toroidal direction)	x15
$L_r \sim a/6 \sim 160\ \rho \sim 10$ correlation lengths	x6
Total speedup	**x10^{16}**
For comparison: Massively parallel computers (1984-2009)	x10^7

Status quo in gyrokinetic simulation

- over the last decade or so, GK has emerged as the standard approach to plasma turbulence
- a variety of nonlinear GK codes is being used and (further) developed
- these codes differ in their numerical schemes and physics contents

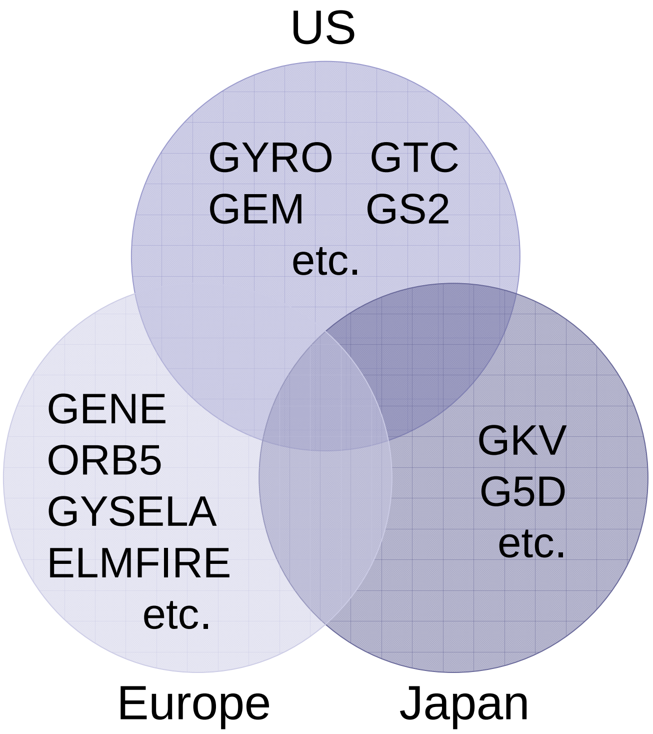

Code development is done in teams (involving computer experts)…

Code benchmarking

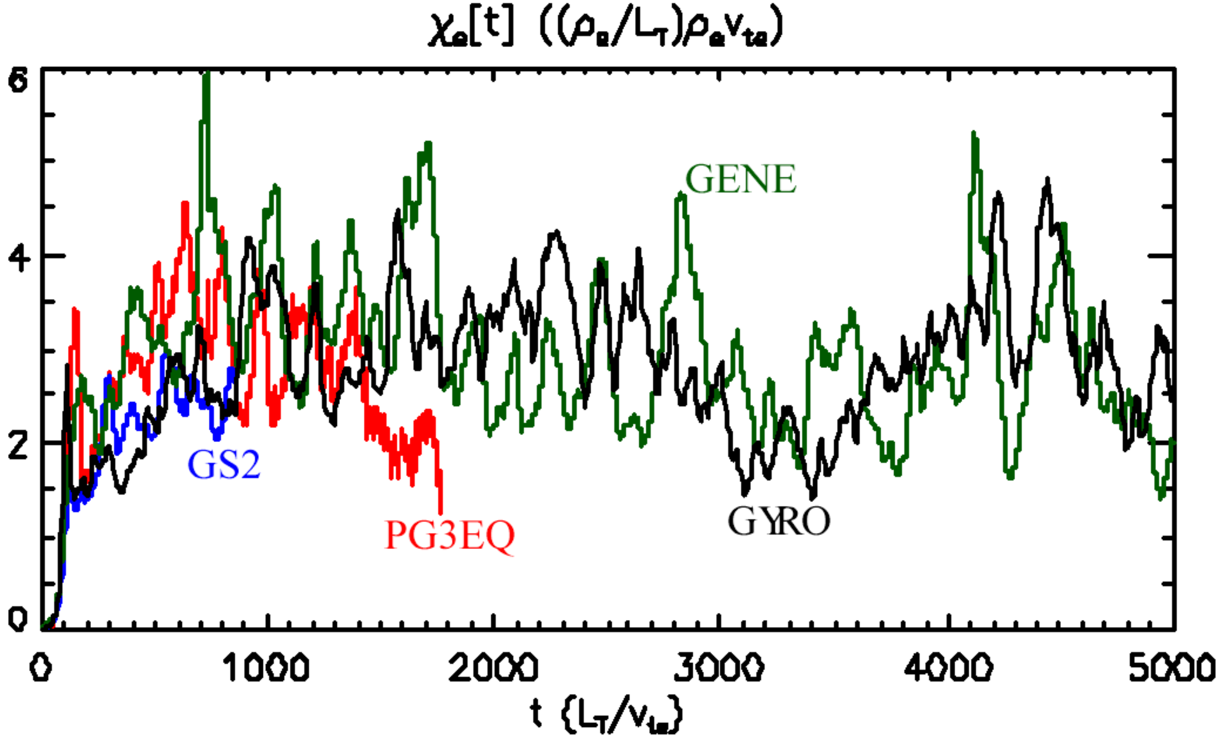

- Are we solving the equations right?
- Yes, according to a recent comparison between 4 codes (GENE, GYRO, GS2, PG3EQ)
- Such efforts are (sometimes) a bit painful but necessary

The GENE code:
A CFD-like approach

The simulation code GENE

- GENE is physically comprehensive and computationally efficient with applications in fusion research and astrophysics

- two main goals: deeper understanding of fundamental physics issues and direct comparisons with experiments/observations

- the differential operators are discretized via a combination of spectral, finite difference, finite element, and finite volume methods; the time stepping is done via a (non-standard) explicit Runge-Kutta method

- GENE is part of the European DEISA benchmark suite and the EU-Japanese IFERC benchmark suite

- GENE is developed cooperatively by an international team, and it is publicly available (www.ipp.mpg.de/~fsj/gene)

Gyrokinetic Electromagnetic Numerical Experiment

GENE parallelization

- parallelization/optimization strategy:
  - high-dimensional domain decomposition
  - optimal subroutines and processor layouts determined during initialization phase (à la FFTW)
  - either pure MPI or mixed MPI/OpenMP paradigm
  - GENE runs efficiently on a range of MPP architectures (e.g., IBM BlueGene, IBM p6, Cray XT4, SGI Altix)

- application support via DEISA:
  - parallelization and tuning
  - porting and scaling tests

Distributed
European
Infrastructure for
Supercomputing
Applications

Special thanks to R. Tisma, T. Dannert, and H. Lederer

GENE on BlueGene/L (strong scaling)

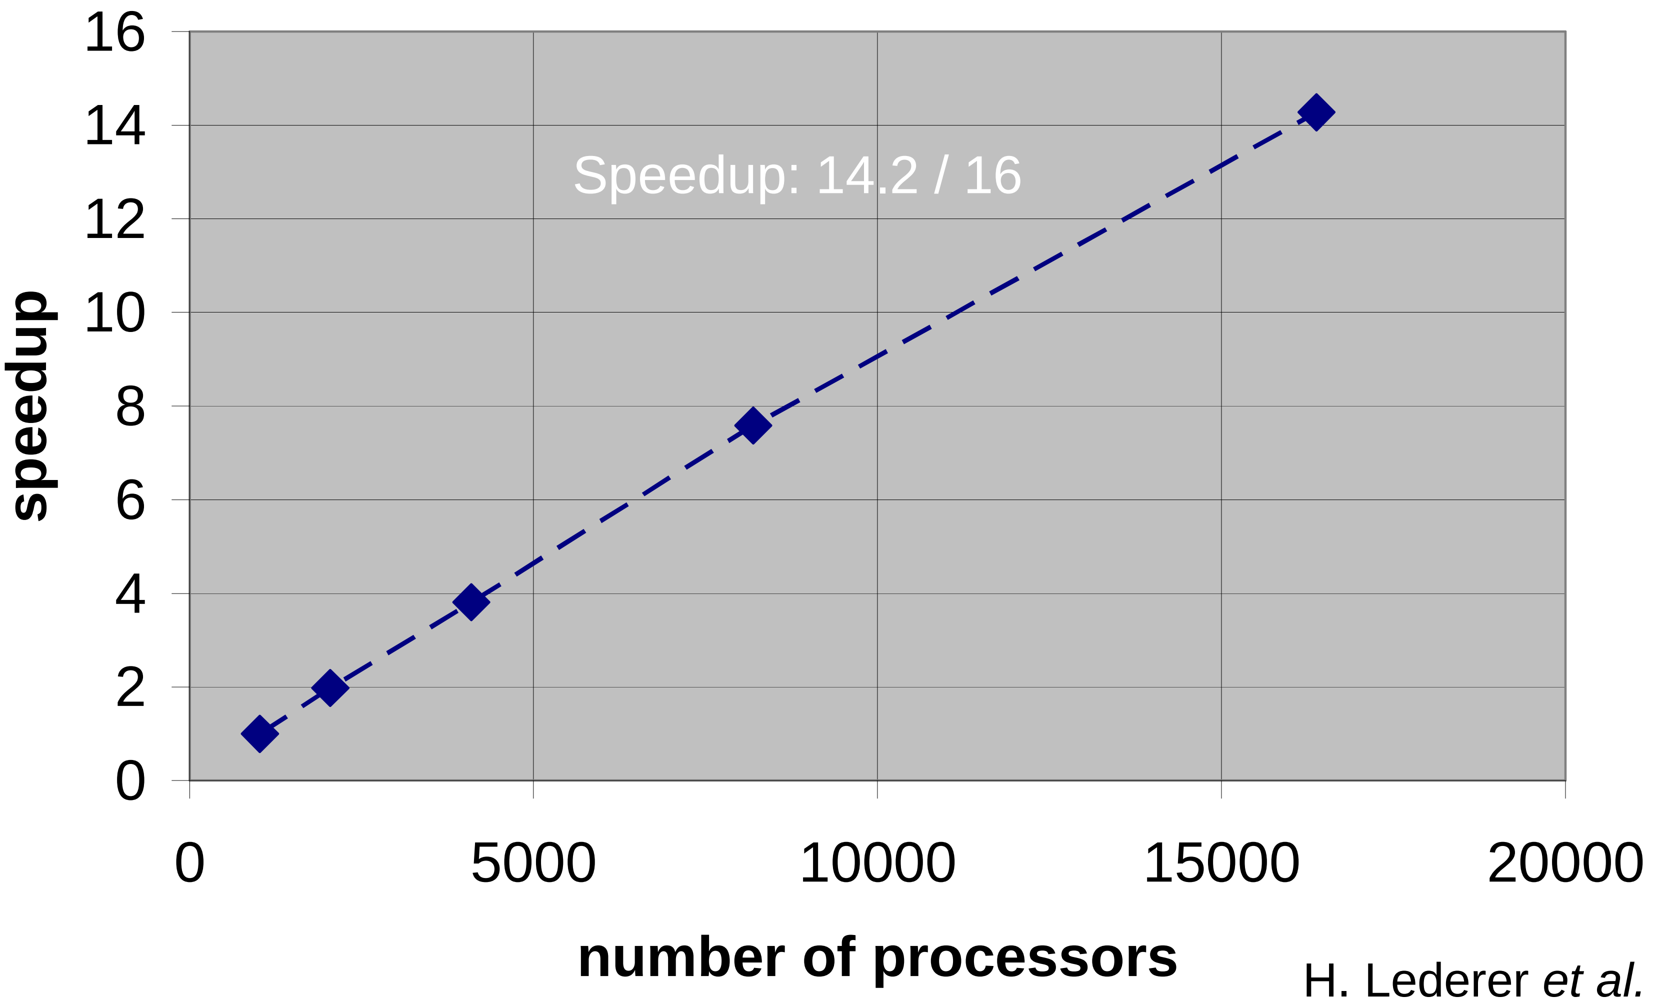

H. Lederer *et al.*

(problem size: ~300-500 GB;
measurements in co-processor mode at IBM Watson Research Center)

GENE on BlueGene/L (weak scaling)

BG/L at Rochester, Minnesota: 2 – 2k procs **BG/L at Watson Research C., NY: 2k – 32k procs**

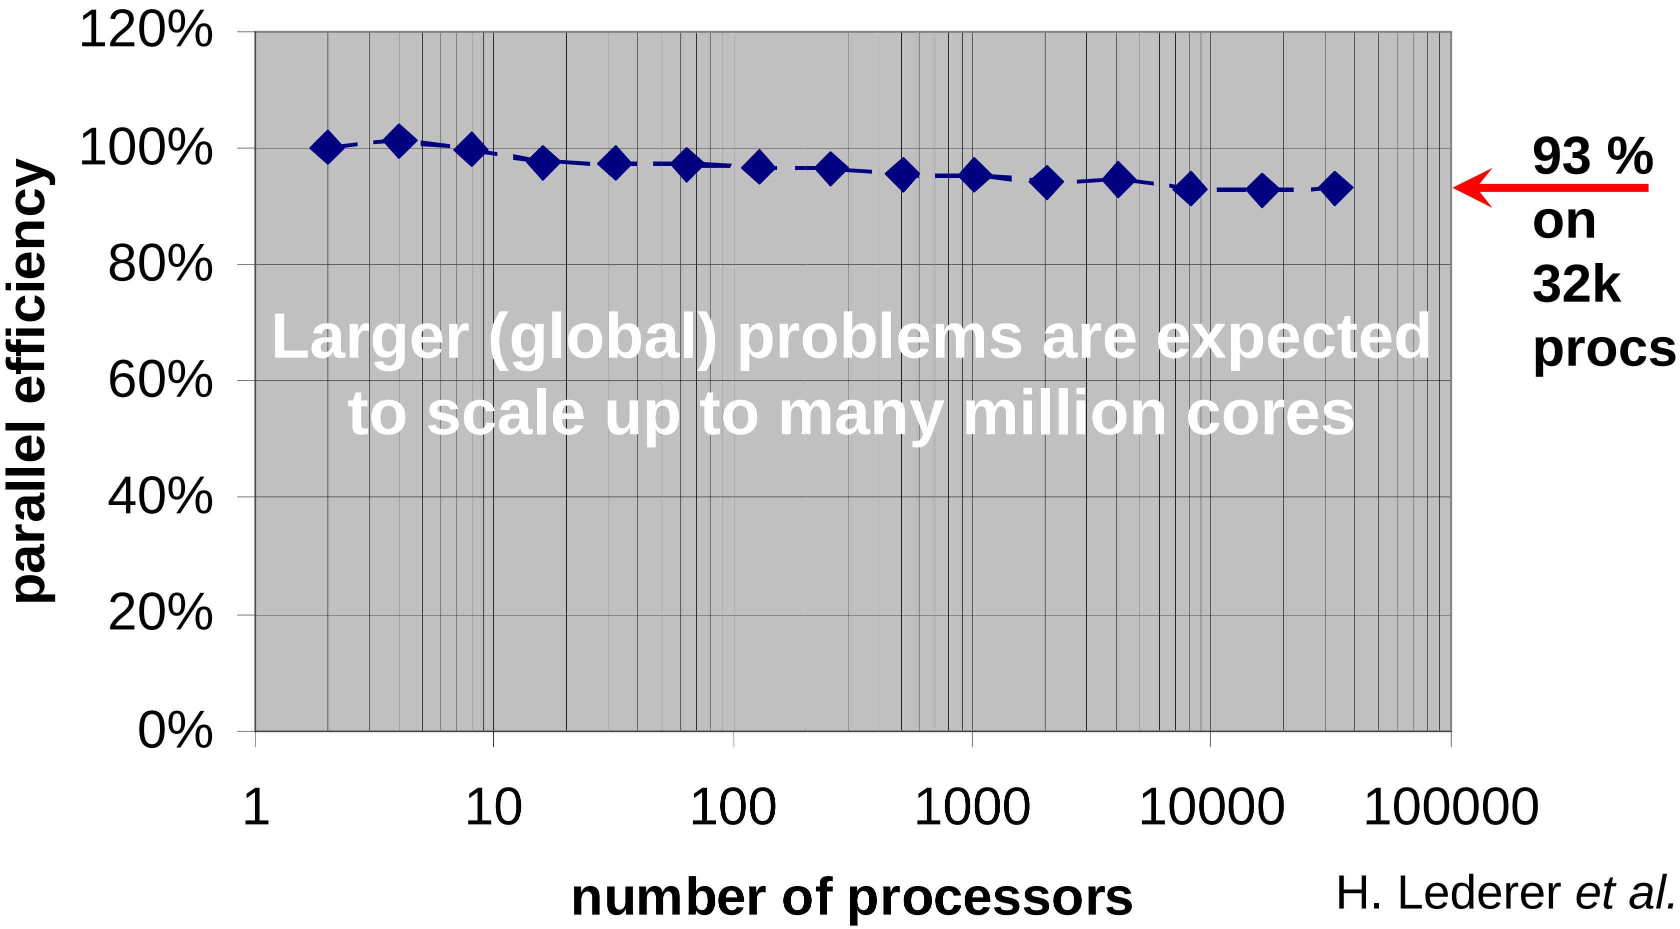

Larger (global) problems are expected to scale up to many million cores

93 % on 32k procs

H. Lederer *et al.*

Weak scaling of GENE (15 points covering 4 orders of magnitude)

(problem ~200 MB/proc; measurements in virtual node mode, **normalized to 2 processors**)

Plasma turbulence “in action”

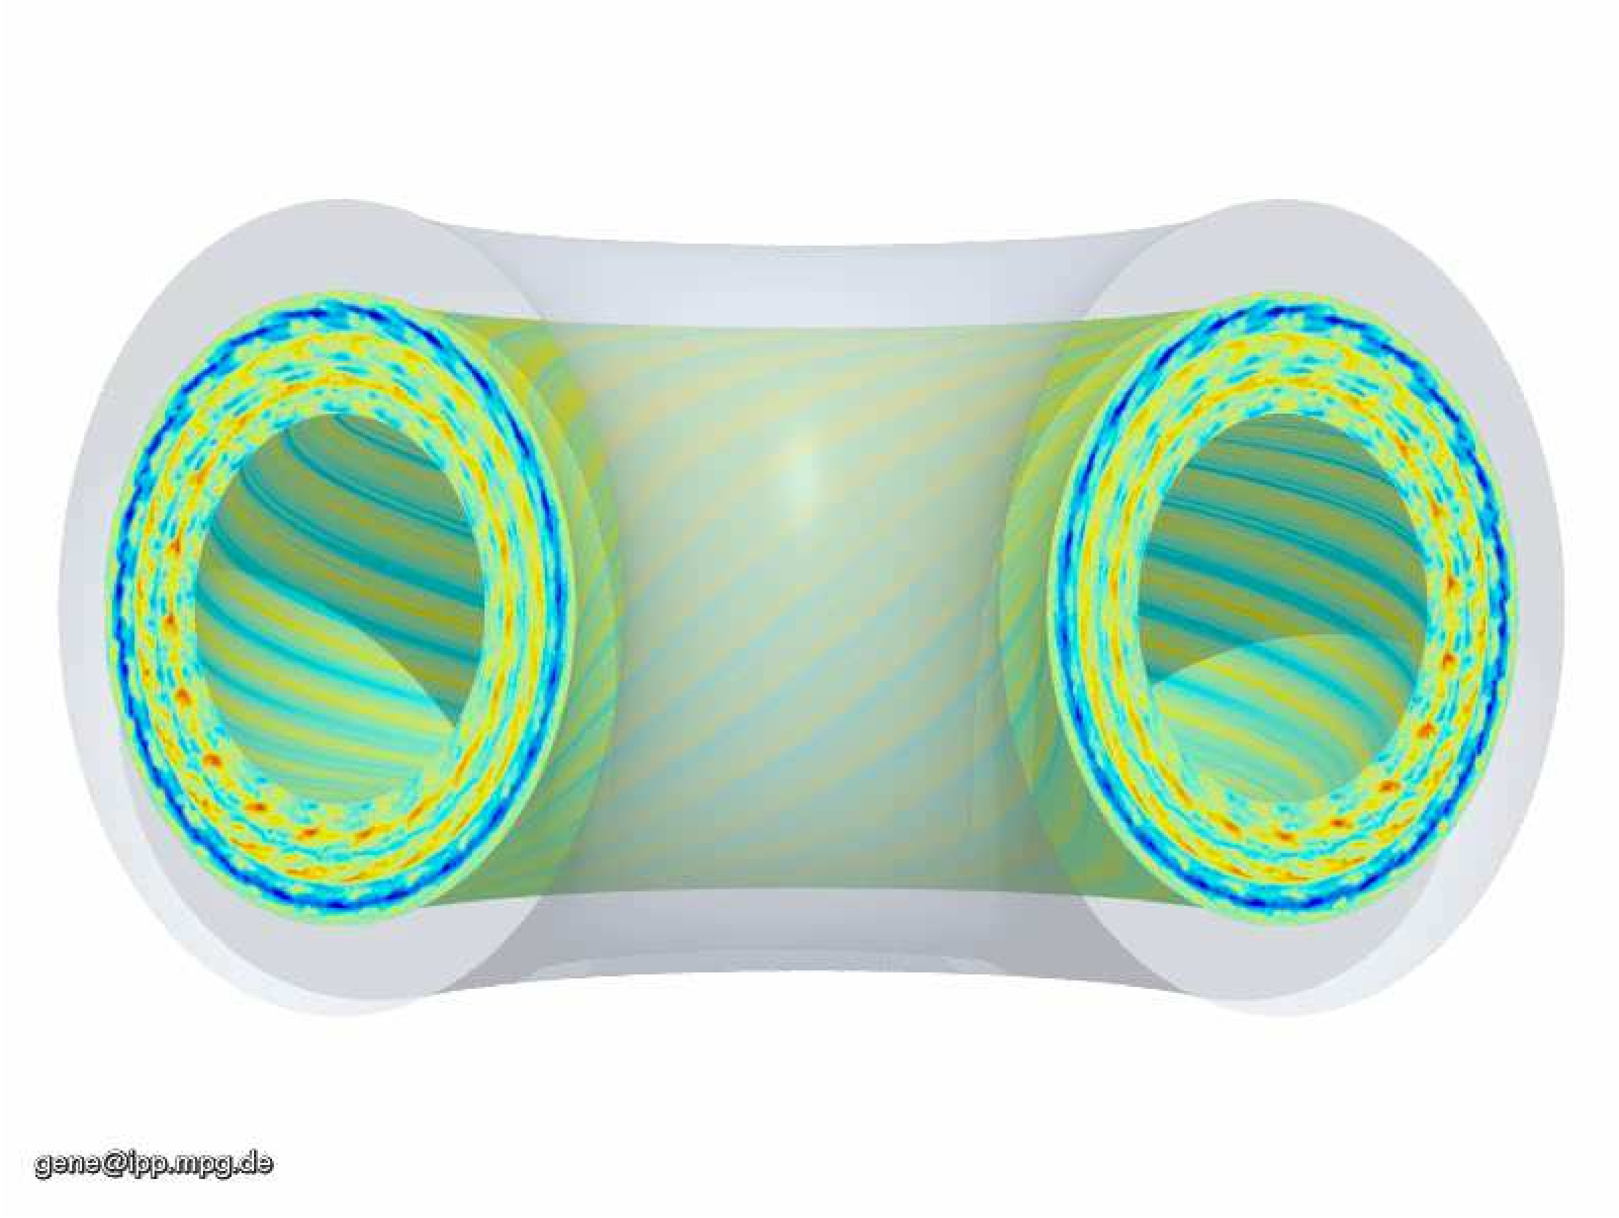

Some insights gained by means of GK simulation

Drive by various microinstabilities

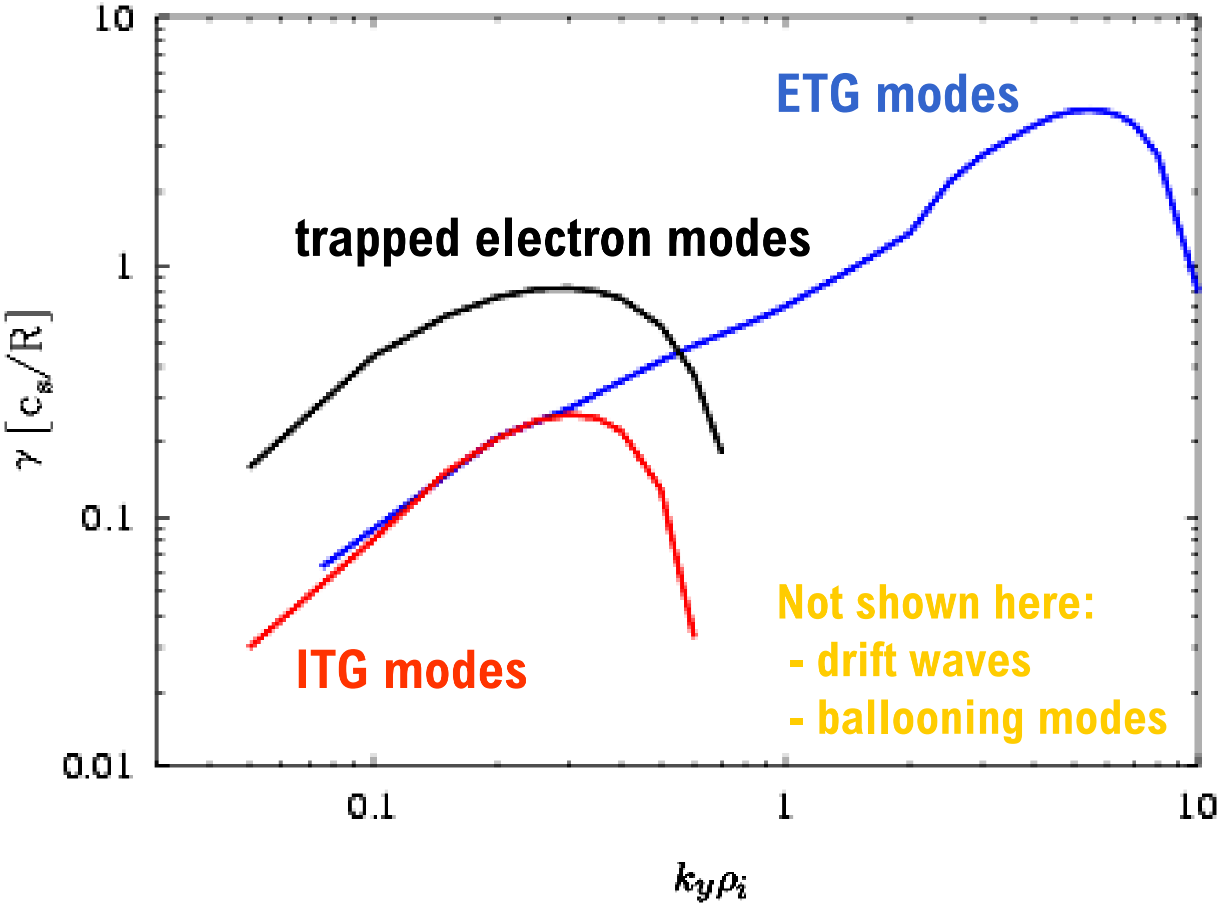

Insights: - Plasma turbulence is not universal

- Characteristic length scale is the gyroradius

Saturation of ITG modes: Zonal flows

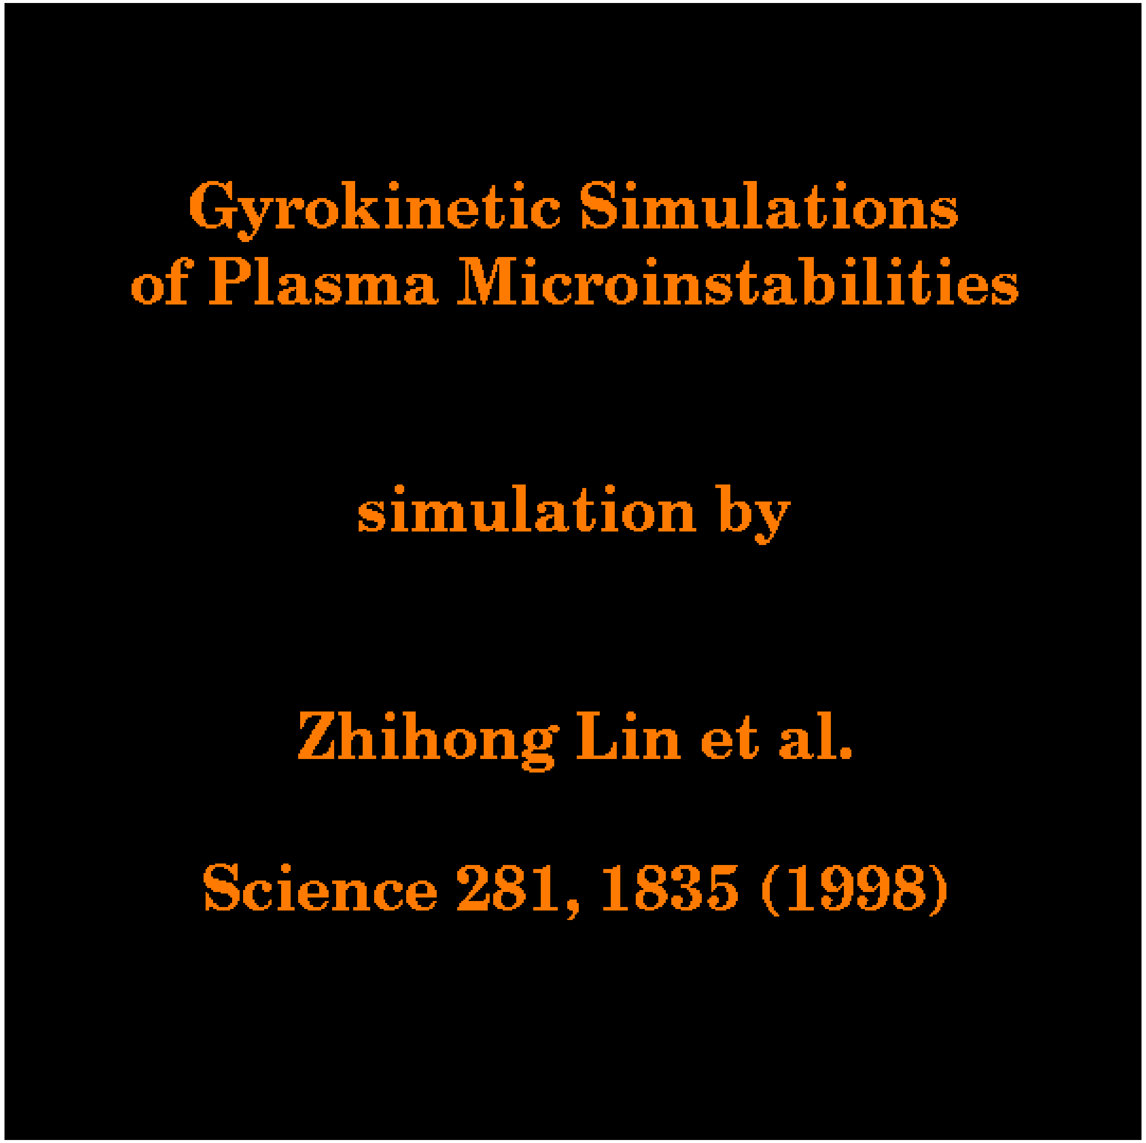

Gyrokinetic Simulations
of Plasma Microinstabilities

simulation by

Zhihong Lin et al.

Science 281, 1835 (1998)

Structure formation

Emergence of **zonal ExB flows** (due to symmetry breaking!)

They are linearly neutrally stable but excited nonlinearly

Zonal flows in geo-/astrophysics

Effect on turbulent transport

Zonal flows may reduce or even suppress the turbulent transport by means of vortex shearing

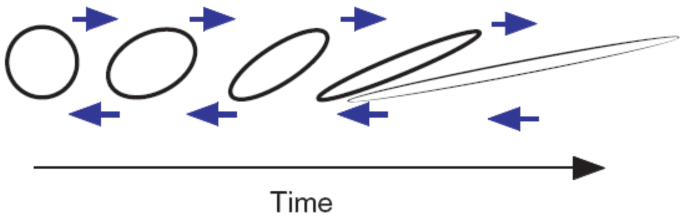

Trying to match experiments: The role of ITG turbulence

[Candy & Waltz, PRL 2003]

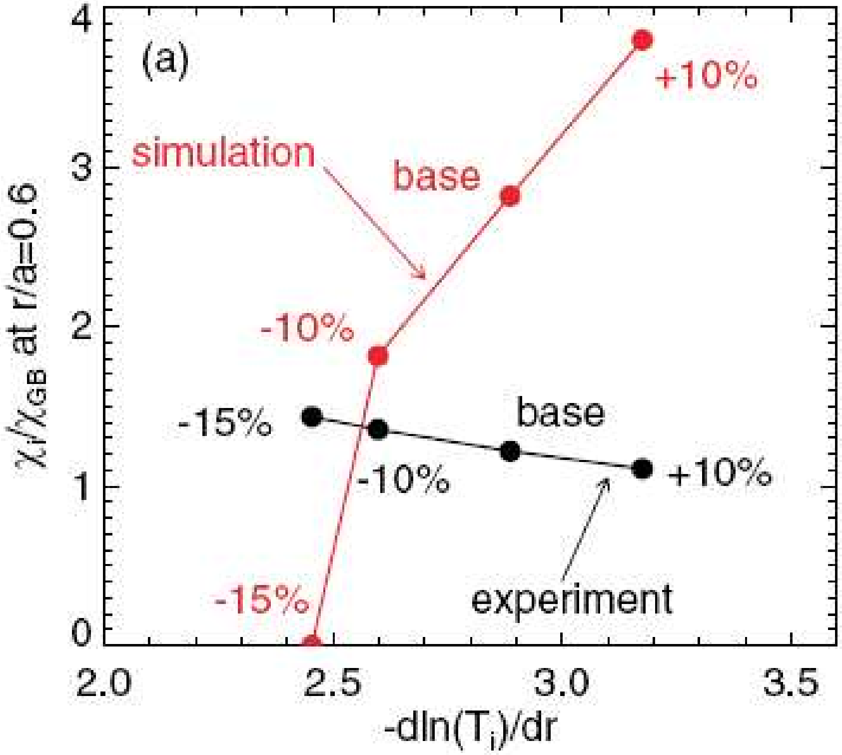

- Direct comparisons with experimental data are encouraging

- *Example:* ITG-induced ion heat fluxes of DIII-D L-mode discharge are recovered within experimental uncertainty

- Note: Ion temperature profiles tend to be stiff

Saturation of TEMs: "eddy damping"

Merz & Jenko, PRL 2008

Low-k_y drive range: large transport contributions, but small random noise; here, one finds:

$$\boxed{\mathcal{N}l[g] \simeq D(-k_\perp^2)g = D\nabla_\perp^2 g}$$

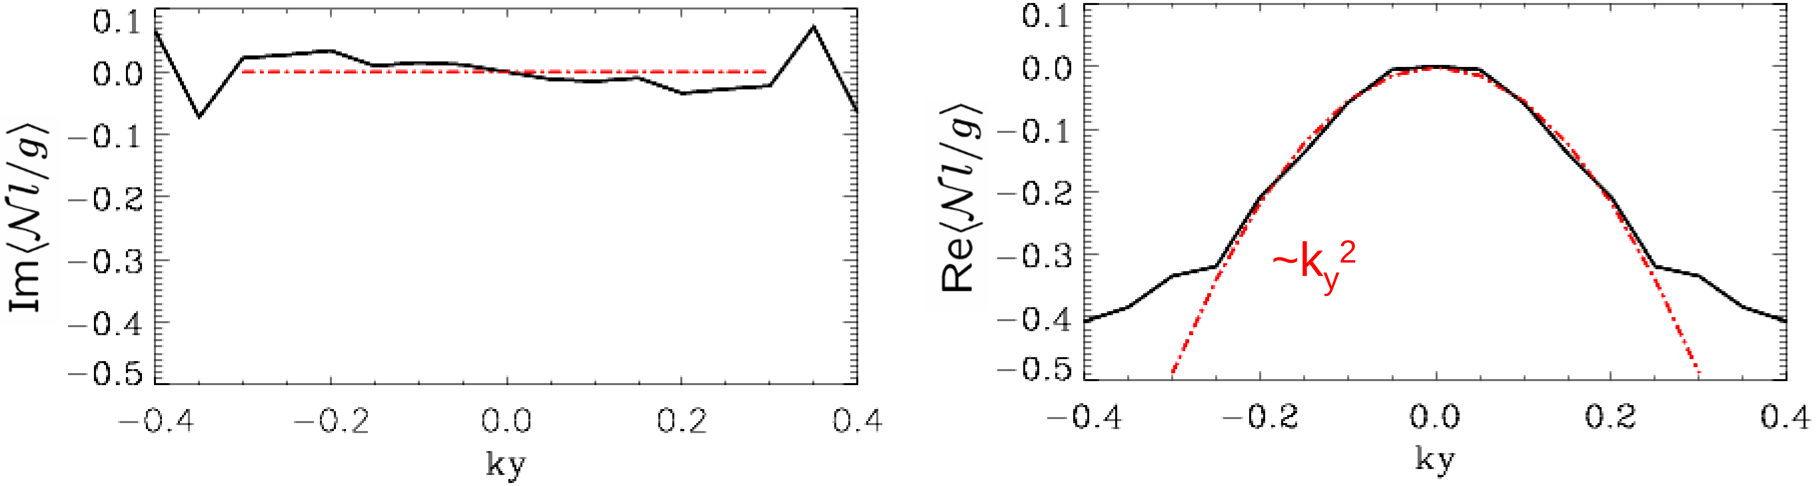

This result allows for the construction of reduced (quasilinear) models.

Coexistence of ITG and ETG modes

box size: ~64 ρ_i resolution: ~2ρ_e

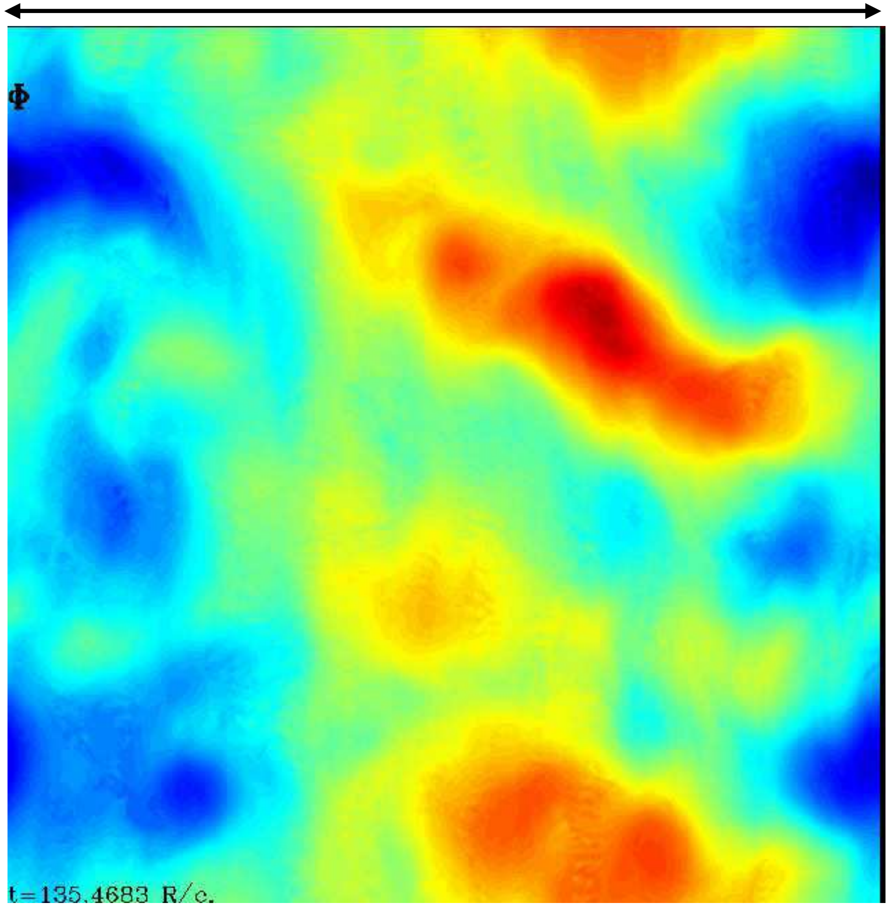

Reduced mass ratio (400), but still > 100,000 CPU-h.

[Görler & Jenko, PRL 2008]

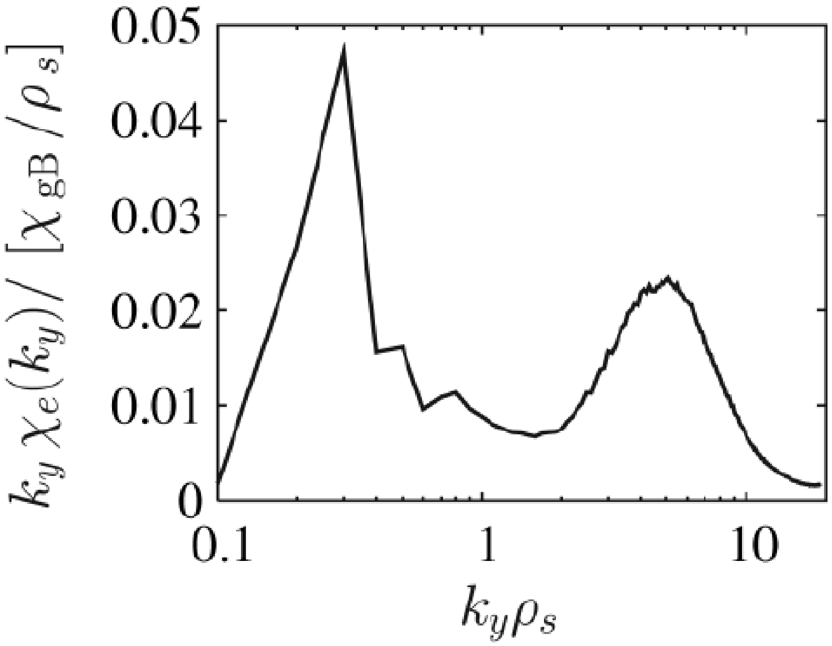

ITG/ETG turbulence: Large fraction of electron heat transport is carried by electron scales (cmp. recent experiments).

Turbulence near magnetic islands

Motivation: Heat transport in the island and island stability

This requires a coupling of MHD-type dynamics and gyrokinetics…

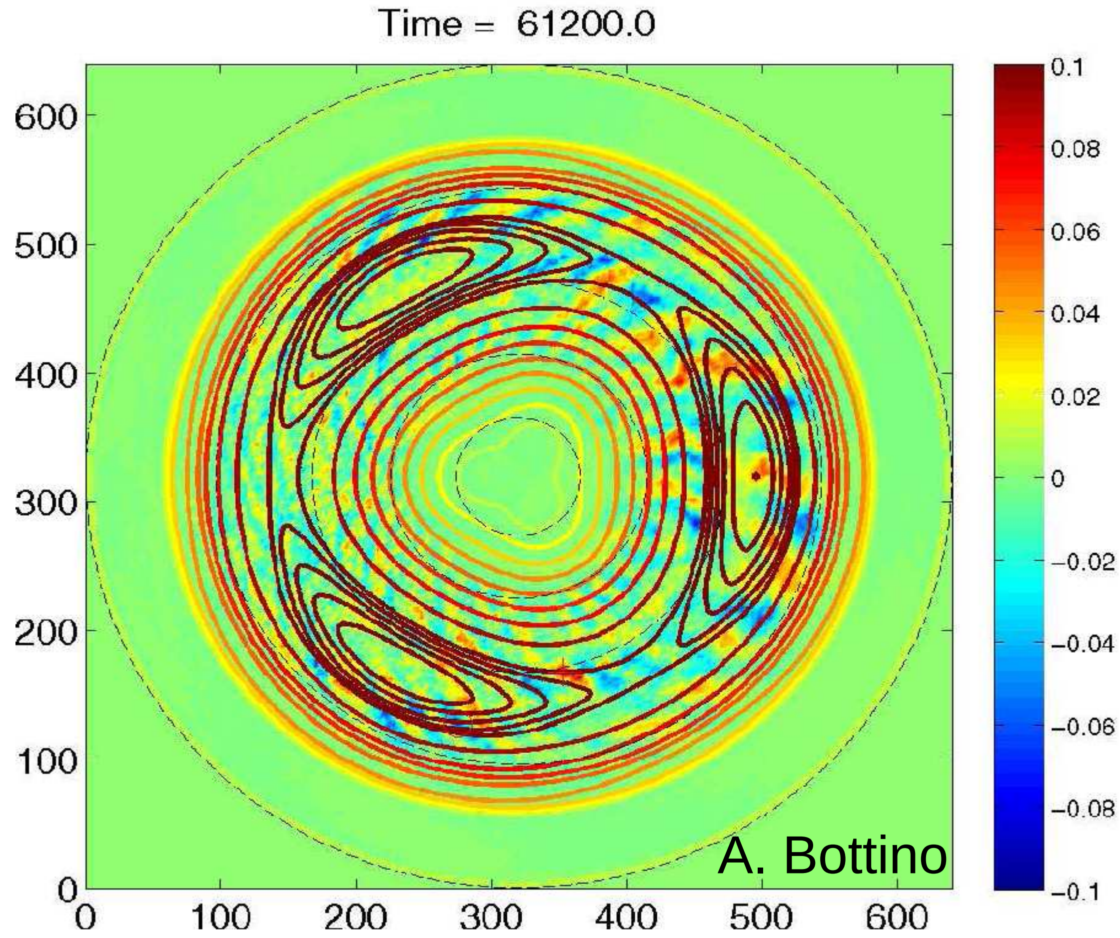

A. Bottino

Some final remarks: Future challenges

Some outstanding open issues

- Prediction of steady state density and temperature profiles
- Nonlocal and finite-size effects
- Core and edge transport barriers
- Reduced descriptions as well as multi-scale, multi-physics models (e.g., interaction of small/large-scale instabilities)

To tackle these issues, we are required and prepared to use the next generation(s) of HPC platforms (e.g., PRACE and DEISA supercomputers).

The ultimate goal? A virtual fusion plasma!

Comparisons of simulation and experiment based on model hierarchy

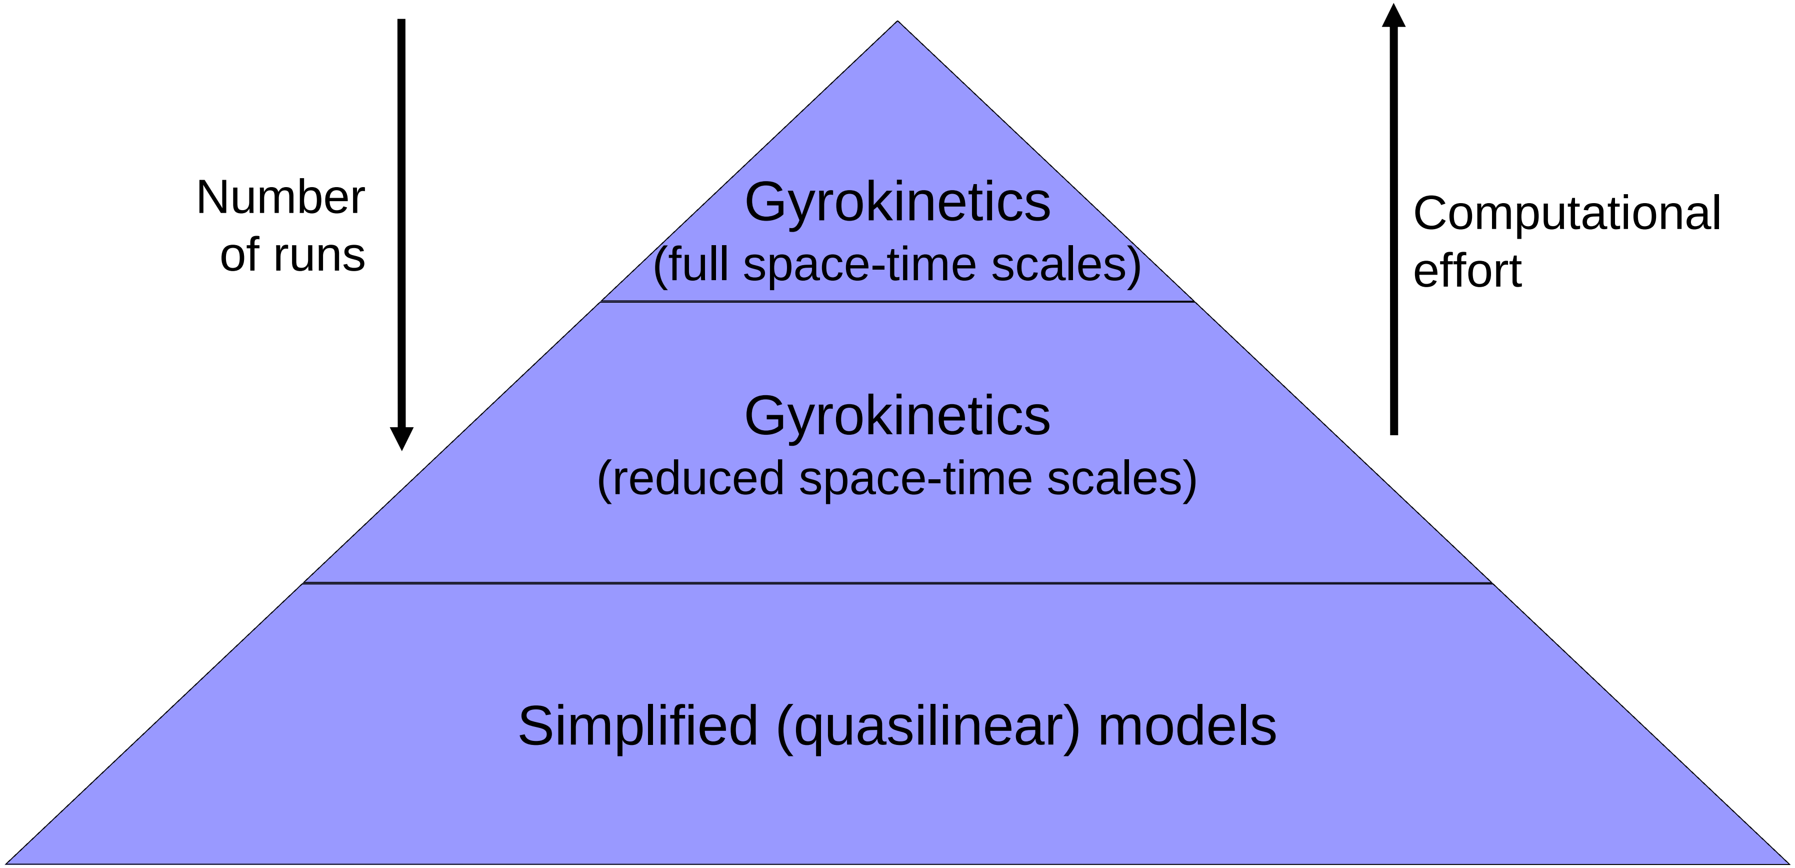

Number of runs

Gyrokinetics
(full space-time scales)

Computational effort

Gyrokinetics
(reduced space-time scales)

Simplified (quasilinear) models

Space-time scale reduction and/or quasilinear modelling is necessary!

A concerted European effort

- These computational challenges are – to some degree – addressed cooperatively, guided by EFDA

EFDA
EUROPEAN FUSION DEVELOPMENT AGREEMENT

- New in 2009: A common 100 TF platform dedicated to fusion research (HPC-FF); the idea is to stimulate HPC activities in smaller European Associations and to foster international collaborations (e.g., with JET)

Conclusions

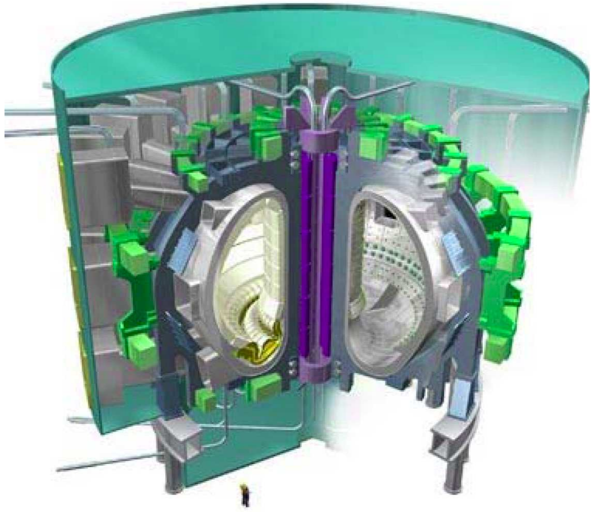

- The successful operation of ITER will depend on our capability to reliably predict, interpret, and optimize ITER discharges by means of Petascale (and Exascale) simulations

- Global ITER simulations (resolving only the ion gyroradius) will require >10 million CPU-h; including also electron gyro-radius scales or transport time scales will add a factor of 10^3

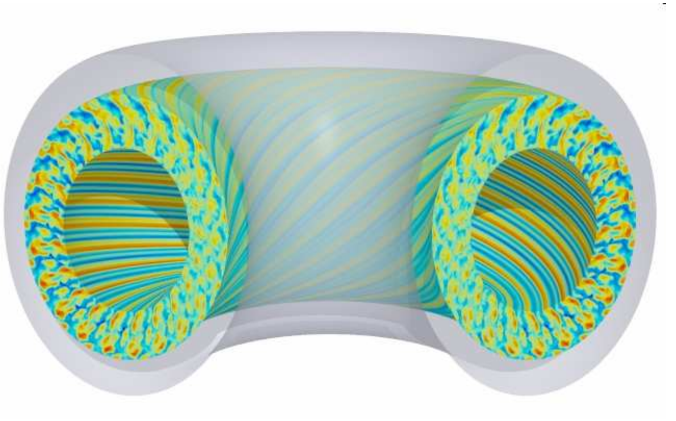

- The development of clean and unlimited fusion energy requires advanced computing capabilities like provided by DEISA & PRACE