

# *Scientific Discovery through Advanced Computing*

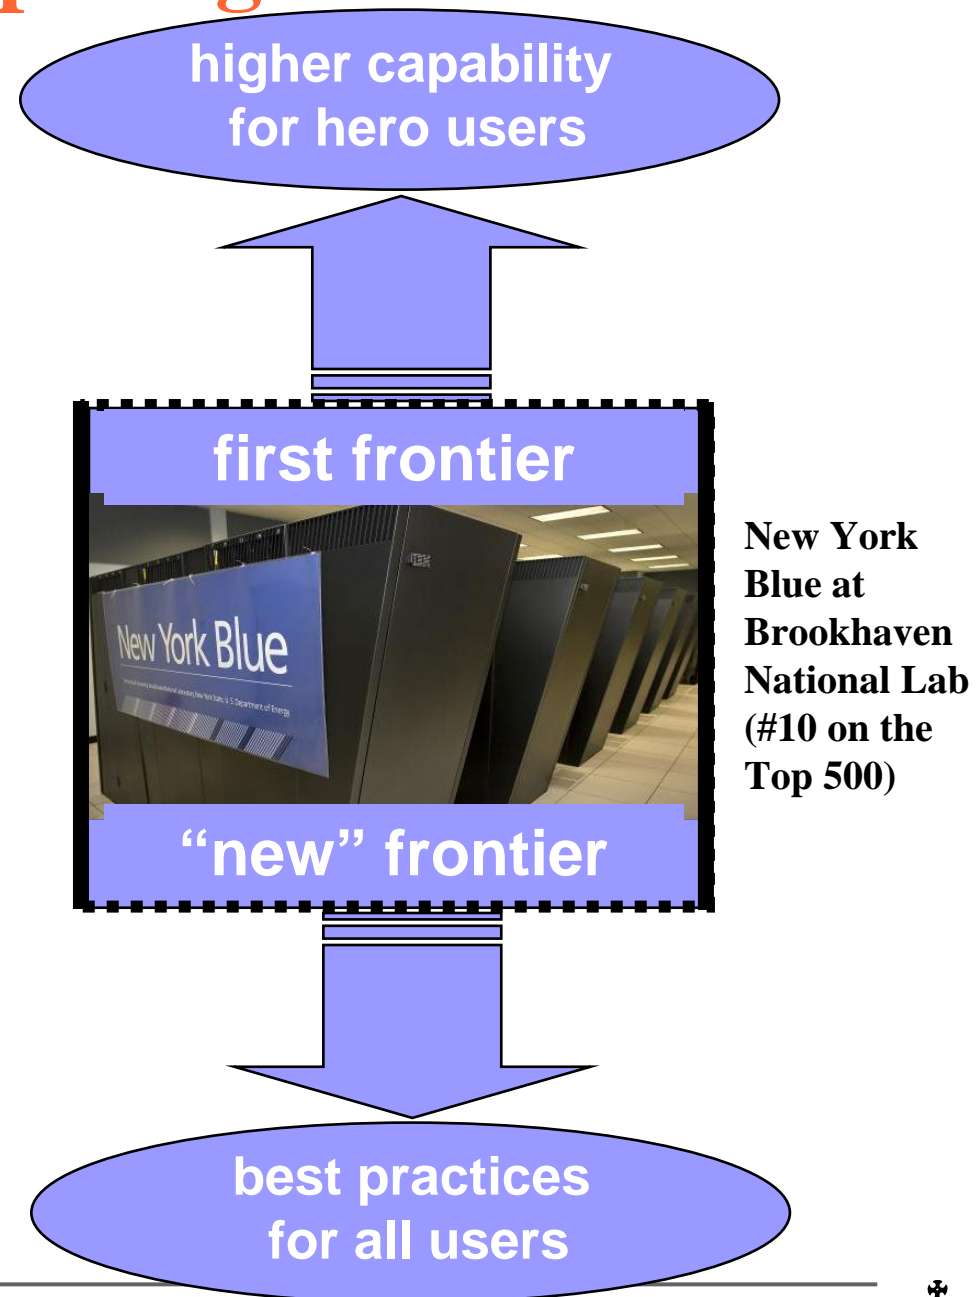
**David Keyes**

**Department of Applied Physics & Applied Mathematics  
Columbia University**



# High-performance computing: two frontiers

- **Two frontiers**
  - raise the peak capability for simulation experts
  - lower the HPC simulation entry threshold for people who are expert in something else
- **Historically, rewards and attention go to the former**
- **We describe a cross-cutting effort, DOE's Scientific Discover through Advanced Computing (SciDAC) program that attempts the latter**



# Presentation plan

- **Are we ready to call simulation “science”?**
- **Motivation in favor**
  - **see also second talk “Petaflop/s, seriously” for supporting trends**
- **Hurdles to science by simulation**
- **Anatomy of a simulation program (U.S. DOE’s SciDAC initiative)**
  - ***caveat*: speaker does not officially represent the U.S. DOE**
- **Example of SciDAC synergy with the international fusion energy program**




# Three pillars of scientific understanding

- **Theory**
- **Experiment**
- **Simulation**  
“theoretical experiments”

Computational simulation :

“a means of scientific discovery that employs a computer system to simulate a physical system according to laws derived from theory and experiment”



# Can simulation produce more than “insight”?

**“The purpose of computing is *insight*, not numbers.”**

— R. W. Hamming (1961)

**“The computer literally is providing a new window through which we can observe the natural world in exquisite detail.”**


— J. S. Langer (1998)

**“What changed were simulations that showed that the new ITER design will, in fact, be capable of achieving and sustaining burning plasma.”**

— R. L. Orbach (2003, in Congressional testimony about why the U.S. should rejoin the International Thermonuclear Energy Reactor (ITER) consortium)



# Can simulation lead to scientific discovery?



Experimental PIV measurement

Instantaneous flame front imaged by density of inert marker



Simulation

Instantaneous flame front imaged by fuel concentration

Images c/o R. Cheng (left), J. Bell (right), LBNL, and NERSC  
**2003 SIAM/ACM Prize in CS&E (J. Bell & P. Colella)**



# Turbulent combustion example (PDE)

- **Simulation models and methods:**


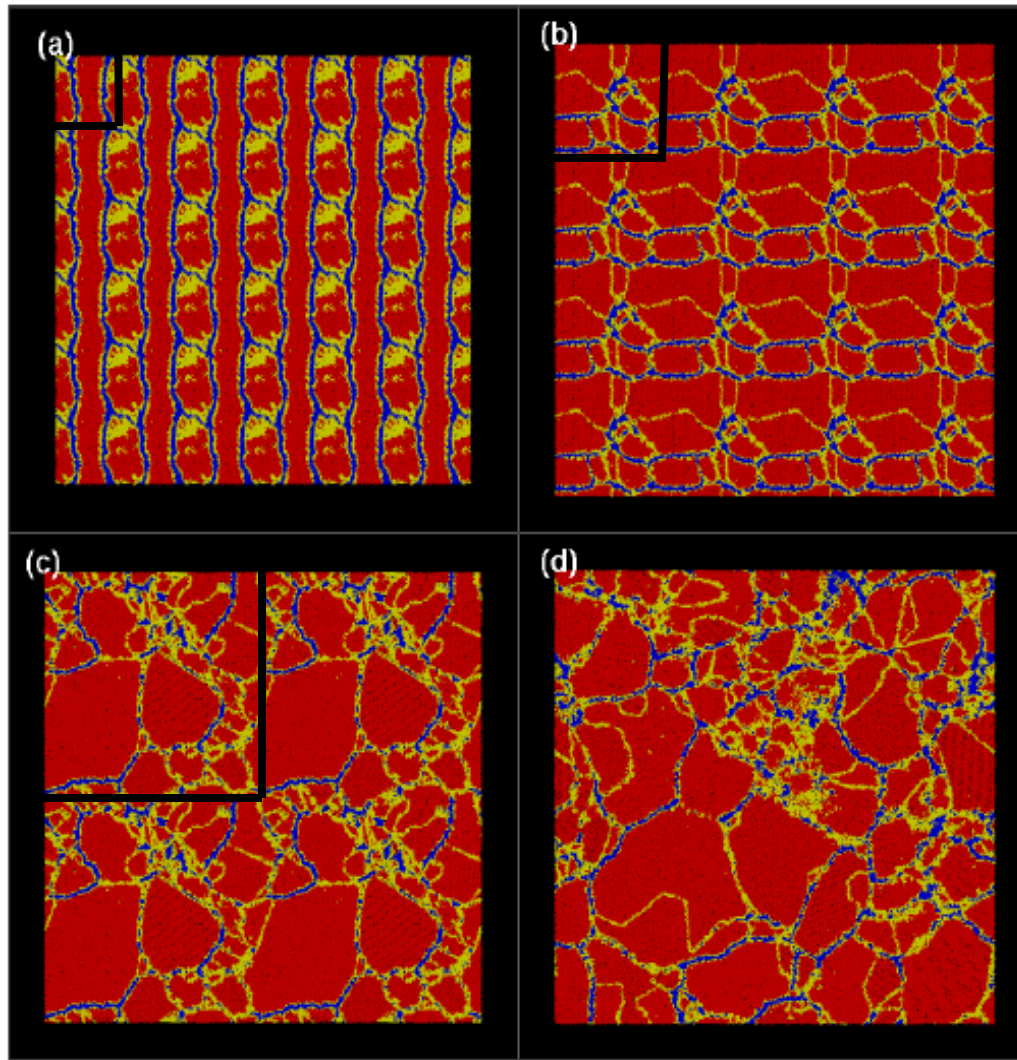
- *Detailed chemical kinetics w 84 reactions, 21 species*
- *Acoustically filtered compressible fluid model*
- *Adaptive mesh refinement,  $10^4 \times$  speedup*
- *Message-passing parallelism, 2048 procs*

This simulation sits at the pinnacle of numerous prior achievements in *experiment, theory, and computer science*

- **Reaction zone location a delicate balance of fluxes of: *species, momentum, internal energy***
- **Directly relevant to: *engines, turbines, furnaces, incinerators (energy efficiency, pollution mitigation)***
- **Component model of other computational apps: *firespread, stellar dynamics, chemical processing***
- **Theory, experiment, and simulation feed on and enrich each other**



# Phase change example (MD)



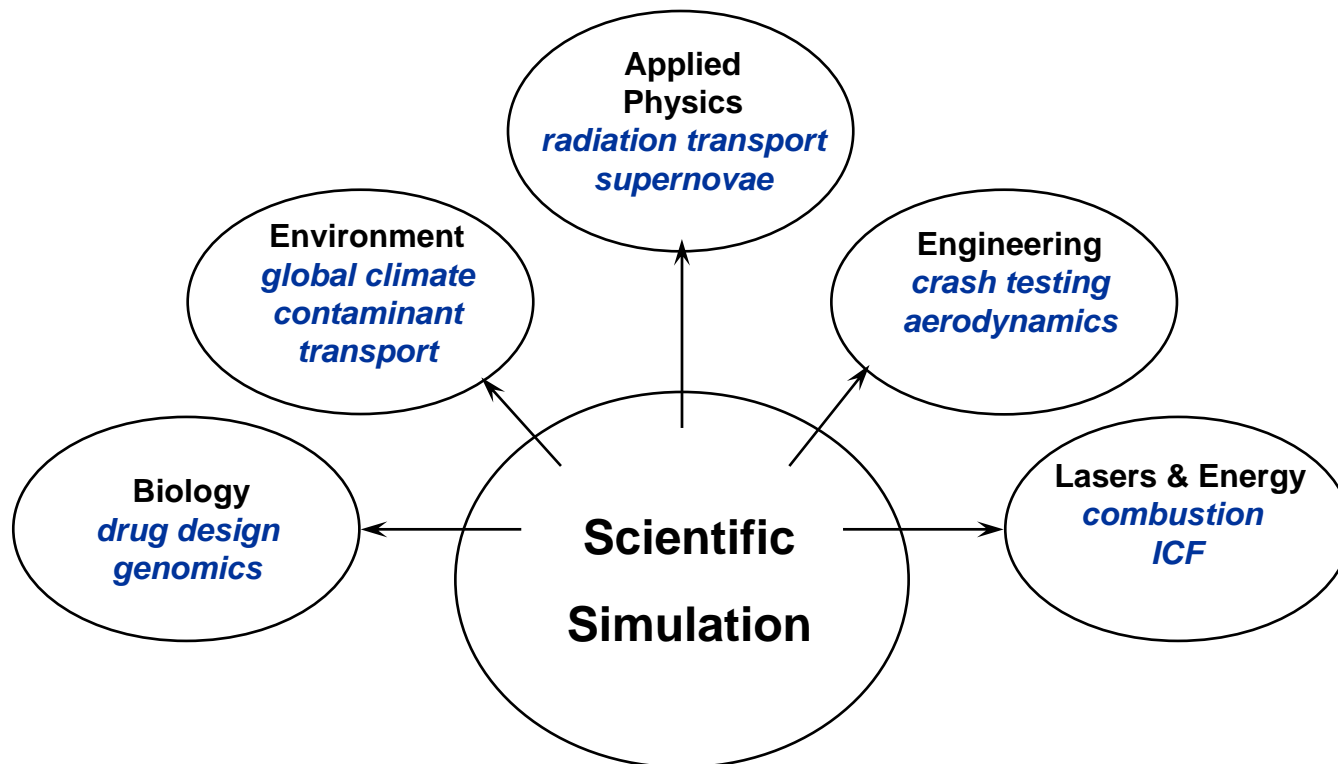
The size of the largest cluster in the system as a function of time, plotted for 64K (blue), 256K (pink), 2M (red), 8M (green), and 16M (black) atoms. The final doubling suggests that the grain size is no longer resolution-limited.

FIG. 10: Cross sectional images displaying the microstructure obtained in simulations containing (a) 64,000 atoms (b) 256,000 atoms, (c) 2,048,000 atoms and (d) 16,384,000 atoms after the start of the coarsening process. The three smaller sample images have been replicated using periodic boundary conditions to appear approximately the same size as the 16M atom simulation.





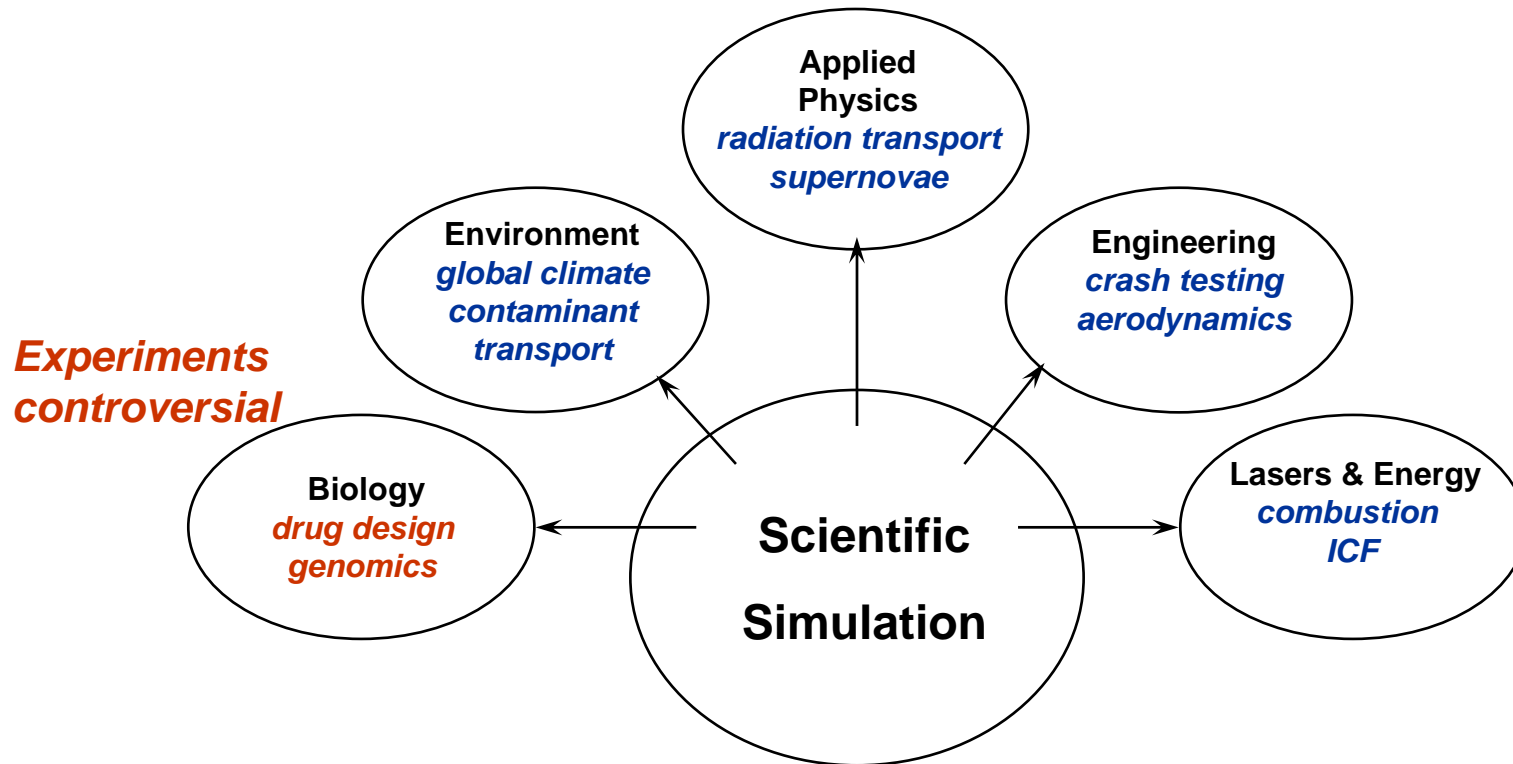
# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



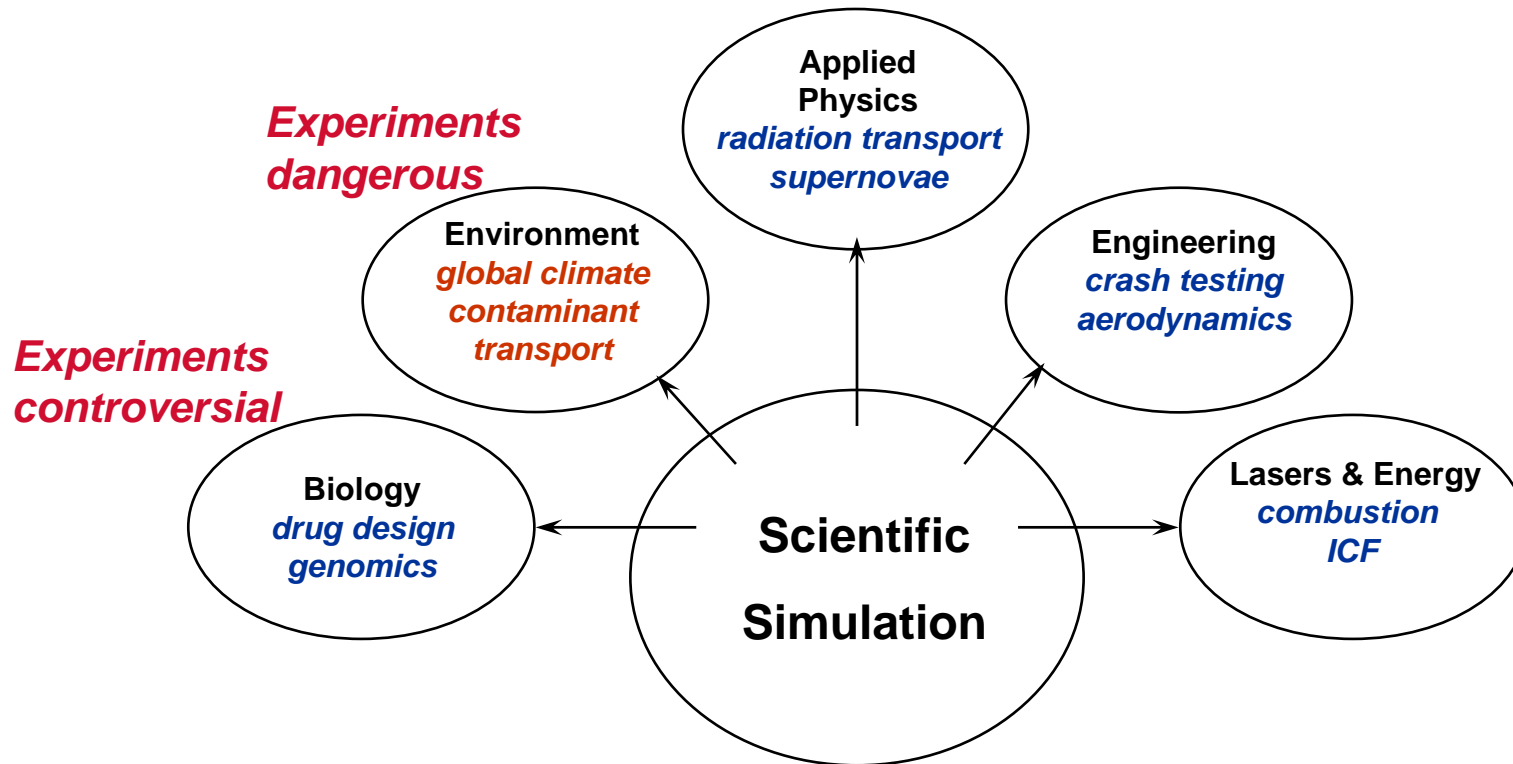
# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



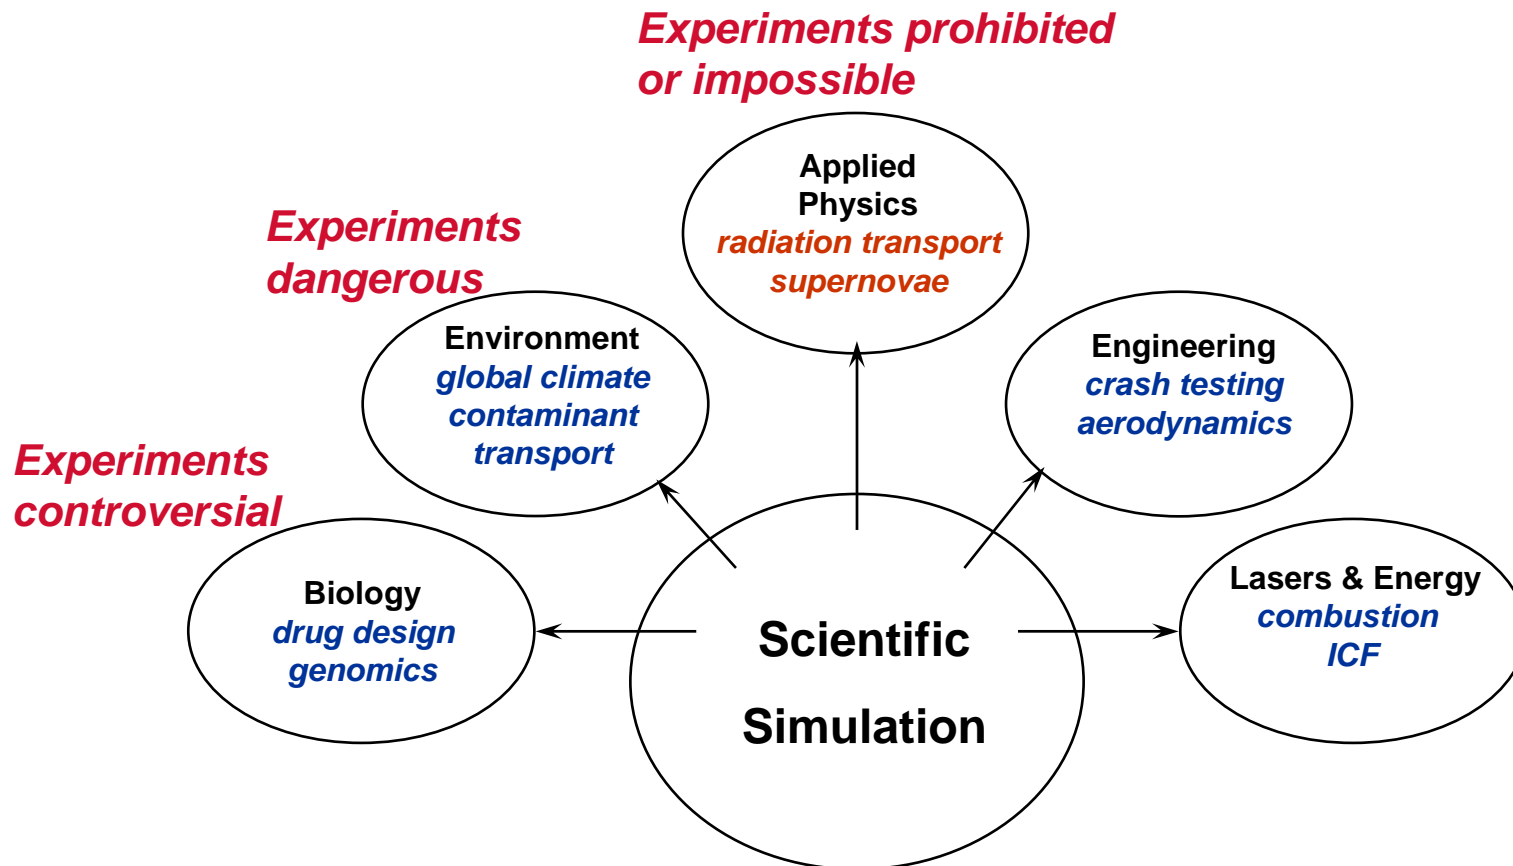
# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



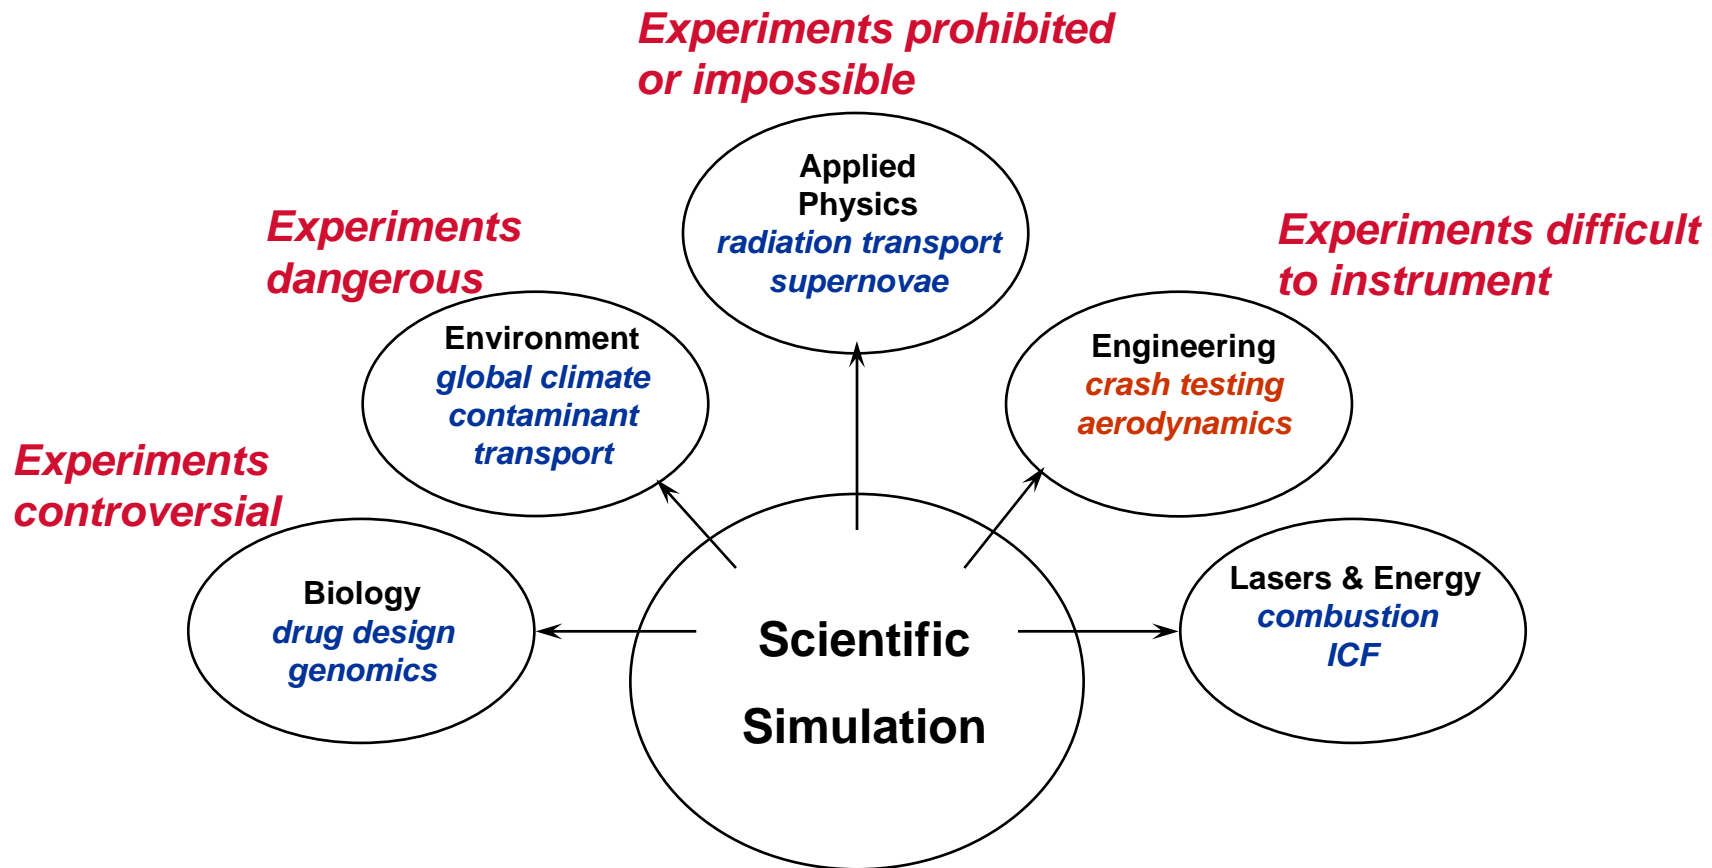
# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



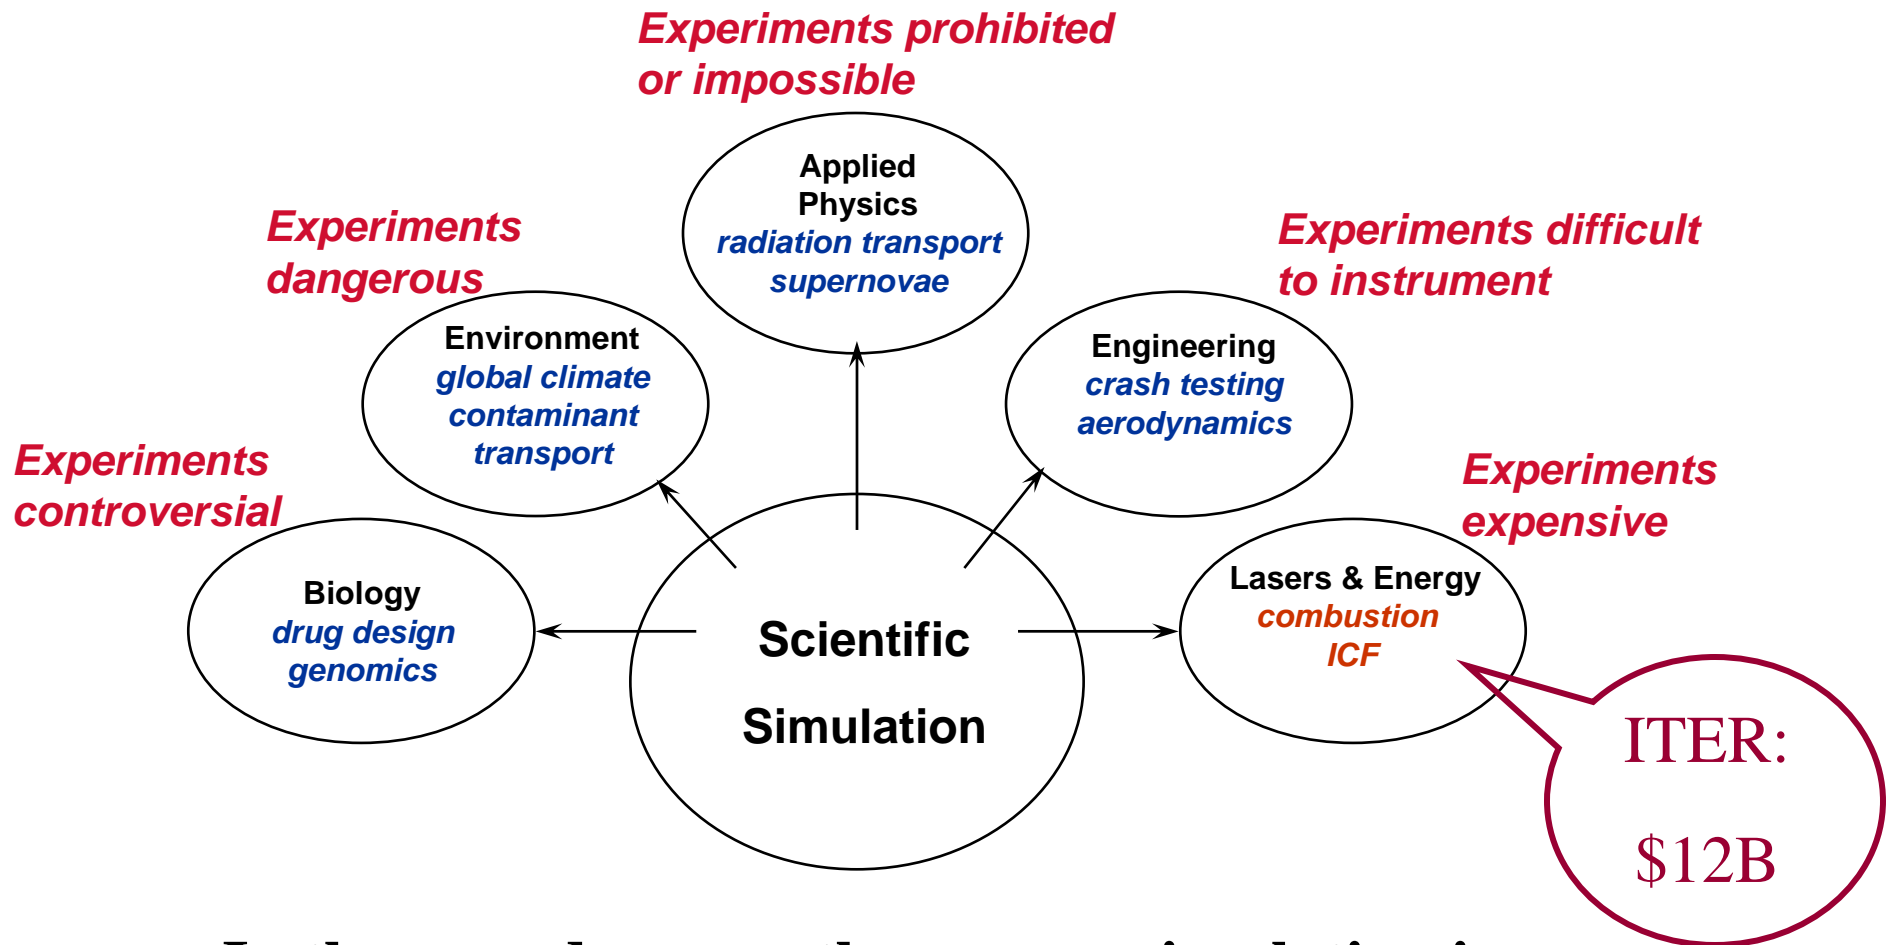
# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



# The imperative of simulation



**In these, and many other areas, simulation is an important complement to experiment.**



# Hurdles to simulation

- **“Triple finiteness” of computers**

- finite precision
- finite number of words
- finite processing rate

Need: stability,  
optimality of  
representation &  
optimality of work

- **Curse of dimensionality**

- Moore’s Law is quickly “eaten up” in 3 space dimensions plus time

Need adaptivity

- **Curse of uncertainty**

- models and inputs are often poorly known

Need UQ methods

- **Curse of knowledge explosion**

- no one scientist can track all necessary developments


Need good  
colleagues 😊



# The power of optimal algorithms

- Advances in algorithmic efficiency rival advances in hardware architecture
- Consider Poisson's equation on a cube of size  $N=n^3$

<i>Year</i>	<i>Method</i>	<i>Reference</i>	<i>Storage</i>	<i>Flops</i>
1947	GE (banded)	Von Neumann & Goldstine	$n^5$	$n^7$
1950	Optimal SOR	Young	$n^3$	$n^4 \log n$
1971	CG	Reid	$n^3$	$n^{3.5} \log n$
1984	Full MG	Brandt	$n^3$	$n^3$




- If  $n=64$ , this implies an overall reduction in flops of ~16 million\*

\*Six-months is reduced to 1 s





# Optimality from multilevel preconditioning




## A Multigrid V-cycle

### Restriction

transfer from  
fine to coarse  
grid

*coarser grid has fewer cells  
(less work & storage)*



*Recursively* apply this  
idea until we have an  
easy problem to solve




### Prolongation

transfer from coarse  
to fine grid

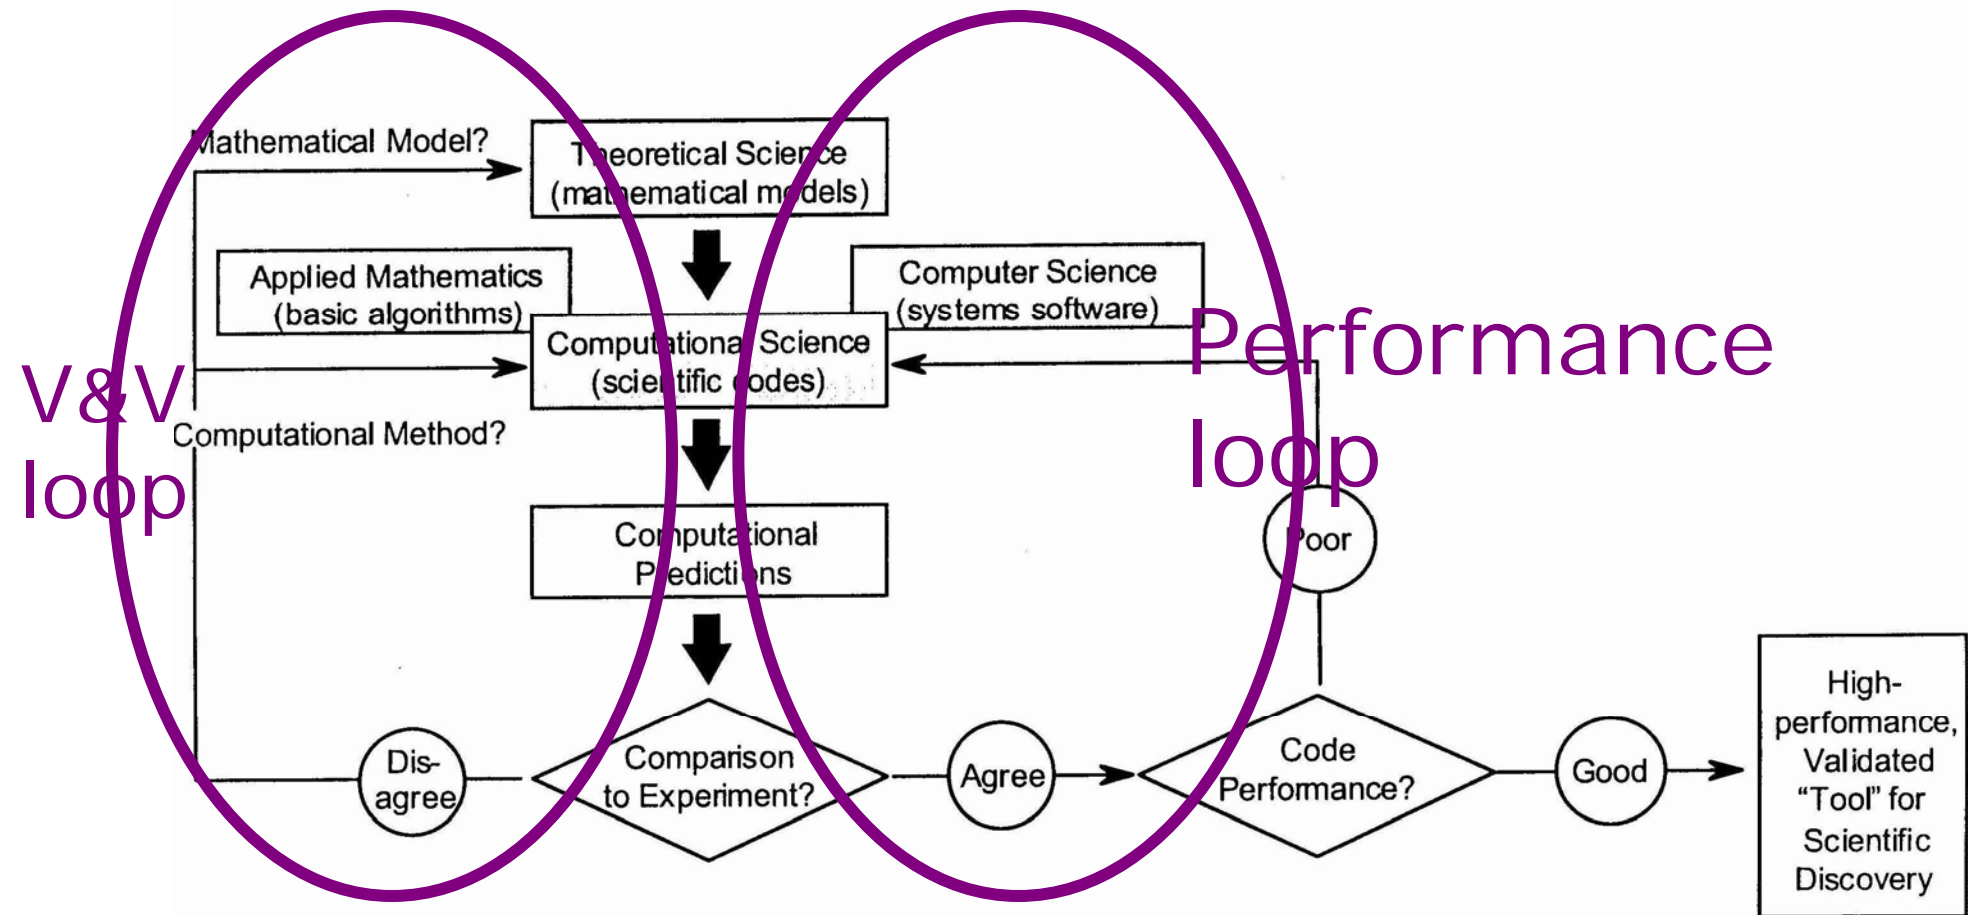


# Algorithms and Moore's Law

- This advance took place over a span of about 36 years, or 24 doubling times for Moore's Law
- $2^{24} \approx 16$  million  $\Rightarrow$  the same as the factor from algorithms alone!



# Designing a simulation code



# Important role of scientific software engineering defines our simulation era

(dates are symbolic)



1686

scientific models



1947

numerical algorithms








1976

computer architecture

1992






scientific software engineering




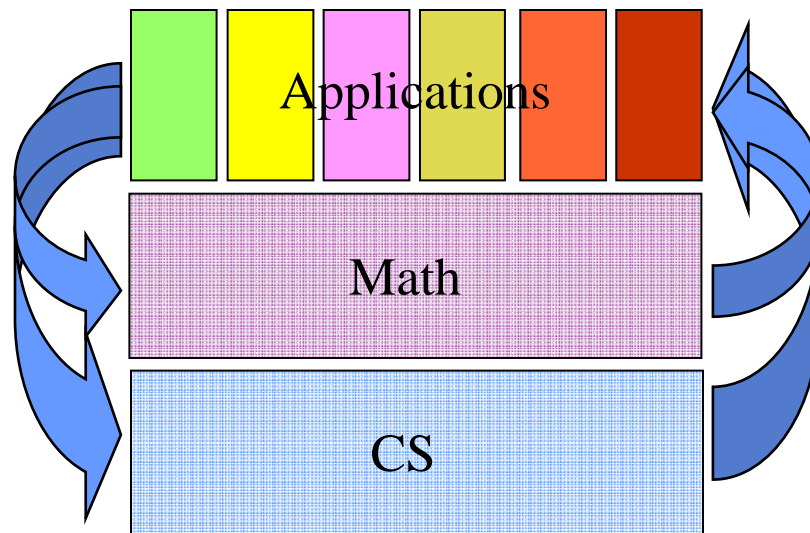
“Computational science is undergoing a phase transition.” – D. Hitchcock, DOE



# SciDAC: economy in general-purpose “ETs” for specialized “Apps”



Many  
applications  
drive



Enabling  
technologies  
respond to all



- **“Enabling technologies” groups to develop reusable software and partner with application groups**
- **In 2006 renewal, 49 projects share \$60M/year, divided between**
  - **applications projects**
  - **lab-based Centers for Enabling Technology (CETs)**
  - **academic-hosted “institutes”**
- **Plus, petaflop/s-scale IBM BlueGene machines at Berkeley and Argonne, and Cray XT machines available at Oak Ridge for SciDAC researchers**



# SciDAC's applied math “centers”

- **Interoperable Tools for Advanced Petascale Simulations (ITAPS)**  
PI: *L. Freitag-Diachin, LLNL*  
For complex domain geometry
- **Algorithmic and Software Framework for Partial Differential Equations (APDEC)**  
PI: *P. Colella, LBNL*  
For solution adaptivity
- **Combinatorial Scientific Computing and Petascale Simulation (CSCAPES)**  
PI: *A. Pothen, Old Dominion U*  
For partitioning and ordering
- **Towards Optimal Petascale Simulations (TOPS)**  
PI: *D. Keyes, Columbia U*  
For scalable solution







# SciDAC

Scientific Discovery through Advanced Computing

applications  
in high  
energy and  
nuclear  
physics




applications in  
biological and  
environmental  
research




scientific  
software and  
network  
infrastructure

applications  
in fusion  
energy  
science



applications  
in basic  
energy  
sciences



Fuel Cell Stack  
Startup Model

- Goals:
- Heat stack rapidly using air
  - Minimize thermal gradients and subsequent stresses

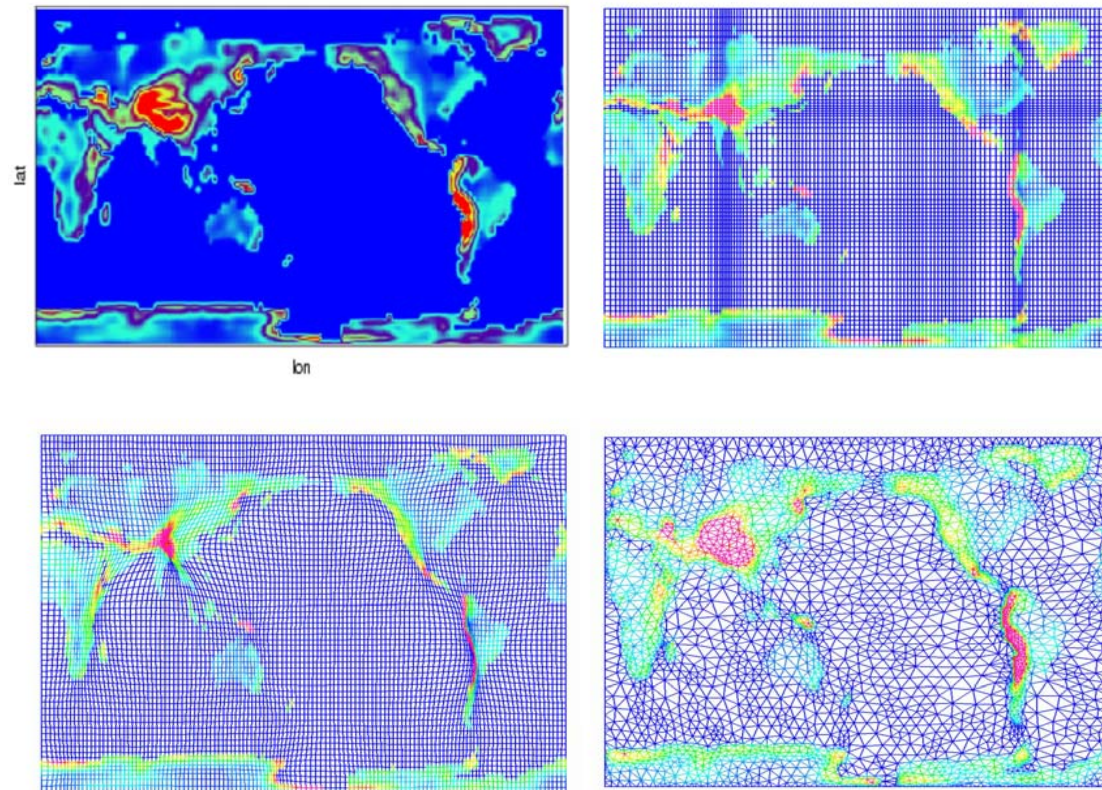




# ITAPS

## Interoperable Tools for Advanced Petascale Simulations



Develop framework for use of multiple mesh and discretization strategies within a single PDE simulation. Focus on high-quality hybrid mesh generation for representing complex and evolving domains, high-order discretization techniques, and adaptive strategies for automatically optimizing a mesh to follow moving fronts or to capture important solution features.



# APDEC

## Algorithmic and Software Framework for PDEs


Develop framework for PDE simulation based on locally structured grid methods, including adaptive meshes for problems with multiple length scales; embedded boundary and overset grid methods for complex geometries; efficient and accurate methods for particle and hybrid particle/mesh simulations.




# CSCAPES

## Combinatorial Scientific Computing and Petascale Simulation


Develop toolkit of partitioners, dynamic load balancers, advanced sparse matrix reordering routines, and automatic differentiation procedures, generalizing currently available graph-based algorithms to hypergraphs



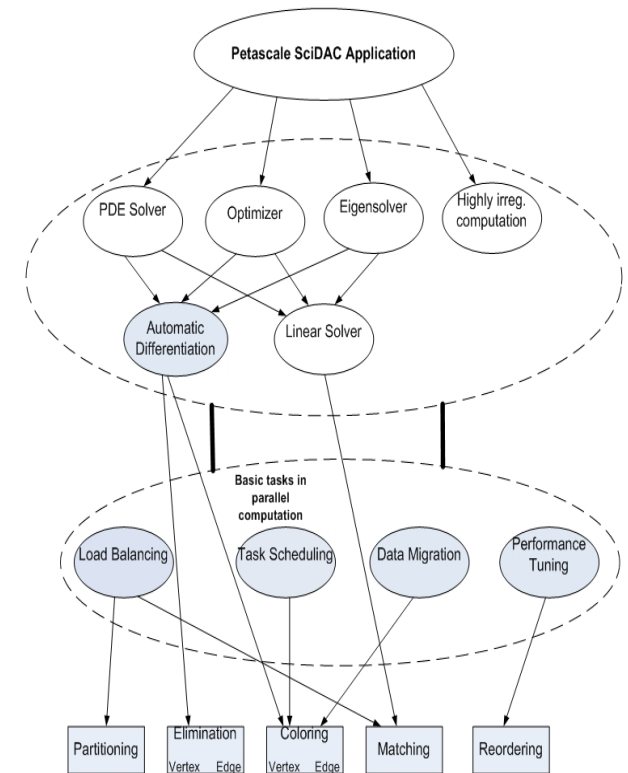
Contact detection



Particle Simulations




Linear solvers & preconditioners



# The TOPS Center for Enabling Technology spans 4 labs & 5 universities

**Our mission: Enable scientists and engineers to take full advantage of petascale hardware by overcoming the scalability bottlenecks traditional solvers impose, and assist them to move beyond “one-off” simulations to validation and optimization**




Columbia University



University of Colorado



University of Texas



Lawrence Livermore  
National Laboratory



Sandia National Laboratories




UCSD

University of California  
at San Diego




# TOPS software has taken a variety of applications to the architectural edge


- TOPS is at the heart of three Gordon Bell “Special” Prizes



1999  
fluids




2003  
seismic



2004  
mechanics


After new coarsening algorithm (red), nearly flat scaled speedup for Algebraic Multigrid




- Scales to the edge of BlueGene/L (131,072 processors, 2B unknowns)

- Powered numerous applications achievements in SciDAC-1

~5X speedup of plasma fusion code through linear solver replacement – like providing “next generation” computer




magneto-hydro-dynamics




accelerator design

Prototype shape optimization capability



Re part of “instanton”

QCD



Im part of “instanton”

Robust solution algorithm for zero quark mass, fine lattices



# Toolchain for PDE solvers in TOPS project

- Design and implementation of “solvers”

- Time integrators  
(w/ sens. anal.)

$$f(\dot{x}, x, t, p) = 0$$

- Nonlinear solvers  
(w/ sens. anal.)

$$F(x, p) = 0$$

- Constrained optimizers


$$\min_u \phi(x, u) \text{ s.t. } F(x, u) = 0, u \geq 0$$

- Linear solvers

$$Ax = b$$

- Eigensolvers

$$Ax = \lambda Bx$$



- Software integration

- Performance optimization

→ Indicates dependence




# Features of DOE's SciDAC initiative



- **Affirmation of importance of simulation**
  - for new scientific discovery, not just for “fitting” experiments
- **Recognition that leading-edge simulation is interdisciplinary**
  - physicists and chemists not supported to write their own software infrastructure; deliverables intertwined with those of math & CS experts
- **Commitment to distributed hierarchical memory computers**
  - new code must target this architecture type
- **Commitment to maintenance of software infrastructure (*rare* to find this 😊)**
- **Requirement of lab-university collaborations**
  - complementary strengths in simulation
  - 13 laboratories and about 50 universities involved



# SciDAC's Fusion Simulation Project: support of the international fusion program



J. Fusion Energy 20: 135-196 (2001)



Fusion by 2017; criticality by 2022

**“Big Iron” meets “Big Copper”**






# ITER: world's first magnetically confined burning plasma





ITER site in Cadaraches, France \*


## Magnetic Confinement



Magnetic Field Line



## The ITER Design: Poloidal Elevation



	ITER
Major radius	6.2 m
Minor radius	2.0 m
Plasma current	15 MA
Toroidal magnetic field	5.3T
Elongation / triangularity	1.85 / 0.49
Fusion power amplification	$\geq 10$
Fusion power	$\sim 400$ MW
Plasma burn duration	$\sim 400$ s

ITER parameters in Q = 10 reference inductive scenario

- China
  - Europe
  - India
  - Japan
  - Korea
  - Russia
  - USA
- See report:  
**“Simulation of Fusion Plasmas”**  
 (2007) *Plasma Science & Technology*,  
 29 authors,  
 Beijing 2006



# ITER challenges

- **Performance limited by plasma instabilities**
  - highest power production performance is near stability limits
  - can degrade magnetic containment
  - potentially damaging to the device
- **Important instabilities can be modeled (physicists believe) with magnetohydrodynamics and/or particle methods**
  - neoclassical tearing modes (NTMs)
  - edge-localized modes (ELMs)
- **High power radio frequency electromagnetic waves can influence stability**
  - triggering or suppressing
  - wave-plasma interactions are multiscale



# Taking on the ITER Challenge, Scientists Look to Innovative Algorithms, Petascale Computers

By Michelle Sipics

The promise of fusion as a clean, self-sustaining and essentially limitless energy source has become a mantra for the age, held out by many scientists as a possible solution to the world's energy crisis and a way to reduce the amounts of greenhouse gases released into the atmosphere by more conventional sources of energy. If self-sustaining fusion reactions can be realized and maintained long enough to produce electricity, the technology could potentially revolutionize energy generation and use.

ITER, initially short for International Thermonuclear Experimental Reactor, is now the official, non-acronymic name (meaning "the way" in Latin) of what is undoubtedly the largest undertaking of its kind. Started as a collaboration between four major parties in 1985, ITER has evolved into a seven-party project that finally found a physical home last year, when it was announced that the ITER fusion reactor would be built in Cadarache, in southern France. (The participants are the European Union, Russia, Japan, China, India, South Korea, and the United States.) In May, the seven initialed an agreement documenting the negotiated terms for the construction, operation, and decommissioning of the ITER tokamak, signifying another milestone for both the project itself and its eventual goal of using fusion to facilitate large-scale energy generation for the world.

Problems remain, however—notably the years, and perhaps decades, of progress needed to attain such a goal. In fact, even *simulating* the proposed ITER tokamak is currently out of reach. But according to David Keyes, a computational mathematician at Columbia University and acting director of the Institute for Scientific Computing Research (ISCR) at Lawrence Livermore National Laboratory, the ability to perform such simulations may be drawing closer.

## Hardware 3, Software 9

"Fusion scientists have been making useful characterizations about plasma fusion devices, physics, operating regimes and the like for over 50 years," Keyes says. "However, to simulate the dynamics of ITER for a typical experimental 'shot' over scales of interest with today's most commonly used algorithmic technologies would require approximately  $10^{24}$  floating-point operations." That sounds bleak, given the 280.6 Tflop/s ( $10^{12}$  flops/s) benchmark performance of the IBM BlueGene/L at Lawrence Livermore National Laboratory—as of June the fastest supercomputer in the world. But Keyes is optimistic: "We expect that with proper algorithmic ingenuity, we can reduce this to  $10^{15}$  flops."

Optimizing the algorithms used, in other words, could lower the computing power required for some ITER simulations by an astounding nine orders of magnitude. Even more exciting, those newly feasible simulations would be at the petascale—ready to run on the petaflop/s supercomputers widely expected within a few years.


The ingenuity envisioned by Keyes even has a roadmap. Together with Stephen Jardin of the Princeton Plasma Physics Laboratory, Keyes developed a breakdown that explains where as many as 12 orders of magnitude of speedup will come from over the next decade: 1.5 from increased parallelism, 1.5 from greater processor speed and efficiency, four from adaptive gridding, one from higher-order elements, one from field-line following coordinates, and three from implicit algorithms.



# Scaling fusion simulations up to ITER

Small tokamak    Large tokamak    Huge tokamak

name	symbol	units	CDX-U	DIII-D	ITER
Field	$B_0$	Tesla	0.22	1	5.3
Minor radius	a	meters	.22	.67	2
Temp.	$T_e$	keV	0.1	2.0	8.
Lundquist no.	S		$1 \times 10^4$	$7 \times 10^6$	$5 \times 10^8$
Mode growth time	$\tau_A S^{1/2}$	s	$2 \times 10^{-4}$	$9 \times 10^{-3}$	$7 \times 10^{-2}$
Layer thickness	$a S^{-1/2}$	m	$2 \times 10^{-3}$	$2 \times 10^{-4}$	$8 \times 10^{-5}$
zones	$N_R \times N_\theta \times N_\phi$		$3 \times 10^6$	$5 \times 10^{10}$	$3 \times 10^{13}$
CFL timestep	$\Delta X / V_A$ (Explicit)	s	$2 \times 10^{-9}$	$8 \times 10^{-11}$	$7 \times 10^{-12}$
Space-time pts			$6 \times 10^{12}$	$1 \times 10^{20}$	$6 \times 10^{24}$



**International  
Thermonuclear  
Experimental  
Reactor**

**2017 – first  
experiments, in  
Cadaraches,  
France**

$10^{12}$  needed  
(explicit  
uniform  
baseline)



# Where to find 12 orders of magnitude in 10 years?

Hardware: 3

Software: 9

- 1.5 orders: increased processor speed and efficiency
- 1.5 orders: ...
- 1 order: ...
  - Same improvements bring many fewer elements
- 1 order: ...
  - ... calculation down to petascale ( $10^{15}$ )!
- 4 orders: ...
  - Zones r ... 1% of ITER volume and resolution req ... from them are  $\sim 10^2$  less severe
- 3 orders: implicit solvers
  - Mode growth time 9 orders longer than Alfvén-limited CFL



# Comments on JK roadmap

- **increased processor speed**
  - 10 years is 6.5 Moore doubling times
- **increased concurrency**
  - BG/L is already  $2^{17}$  procs, MHD now at ca.  $2^{12}$
- **higher-order discretizations**
  - low-order FE preconditioning of high-order discretizations (Orszag, Fischer, Manteuffel, etc.)
- **flux-surface following gridding**
  - evolve mesh to approximately follow flux surfaces
- **adaptive gridding**
  - adapt mesh to concentrate points in high-gradient regions
- **implicit solvers**
  - we propose Newton-like fully implicit, with Krylov/MG innards



# SciDAC solver collaboration examples

- **Meeting physicists *at* a well-defined traditional interface**
  - **Magnetic fusion energy – swapping in new linear solvers**
- **Collaborating with physicists *across* traditional interfaces**
  - **Accelerator design – multidisciplinary design optimization**
  - **Quantum chromodynamics – research prototyping of new algorithm**



# Illustrations from computational MHD

- **M3D code (Princeton)**
  - multigrid replaces block Jacobi/ASM preconditioner for optimality
  - new algorithm callable across  $Ax=b$  interface
- **NIMROD code (General Atomics)**
  - direct elimination replaces PCG solver for robustness
  - scalable implementation of old algorithm for  $Ax=b$


The fusion community may use more cycles on unclassified U.S. DOE computers than any other (e.g., 32% of all cycles at NERSC in 2003). Well over 90% of these cycles are spent solving *linear systems* in M3D and NIMROD, which are prime U.S. code contributions to the designing of ITER.





# NIMROD: direct elim. for robustness

- **NIMROD code**
  - high-order finite elements
  - complex, nonsymmetric linear systems with 10K-100K unknowns in 2D (**>90% exe. time**)
- **TOPS collaboration**
  - replacement of diagonally scaled Krylov with **SuperLU**, a supernodal parallel sparse direct solver
  - 2D tests run 100× faster; 3D production runs are ~5× faster




*c/o D. Schnack, et al.*



# M3D: multigrid for optimality


## ● M3D code

- unstructured mesh, hybrid FE/FD discretization with C0 elements
- Sequence of real scalar systems (**>90% exe. time**)






## ● TOPS collaboration

- replacement of additive Schwarz (ASM) preconditioner with algebraic multigrid (AMG) from Hypre
- achieved mesh-independent convergence rate
- $\sim 5\times$  improvement in execution time



# Resistive MHD prototype implicit solver

- *Magnetic reconnection*: the breaking and reconnecting of oppositely directed magnetic field lines in a plasma, replacing hot plasma core with cool plasma, halting the fusion process
- Replace explicit updates with implicit Newton-Krylov from **SUNDIALS** with factor of  $\sim 5\times$  in execution time



J. Brin et al., "Geospace Environmental Modeling (GEM) magnetic reconnection challenge," *J. Geophys. Res.* 106 (2001) 3715-3719.



# Some high-end simulation plans in SciDAC

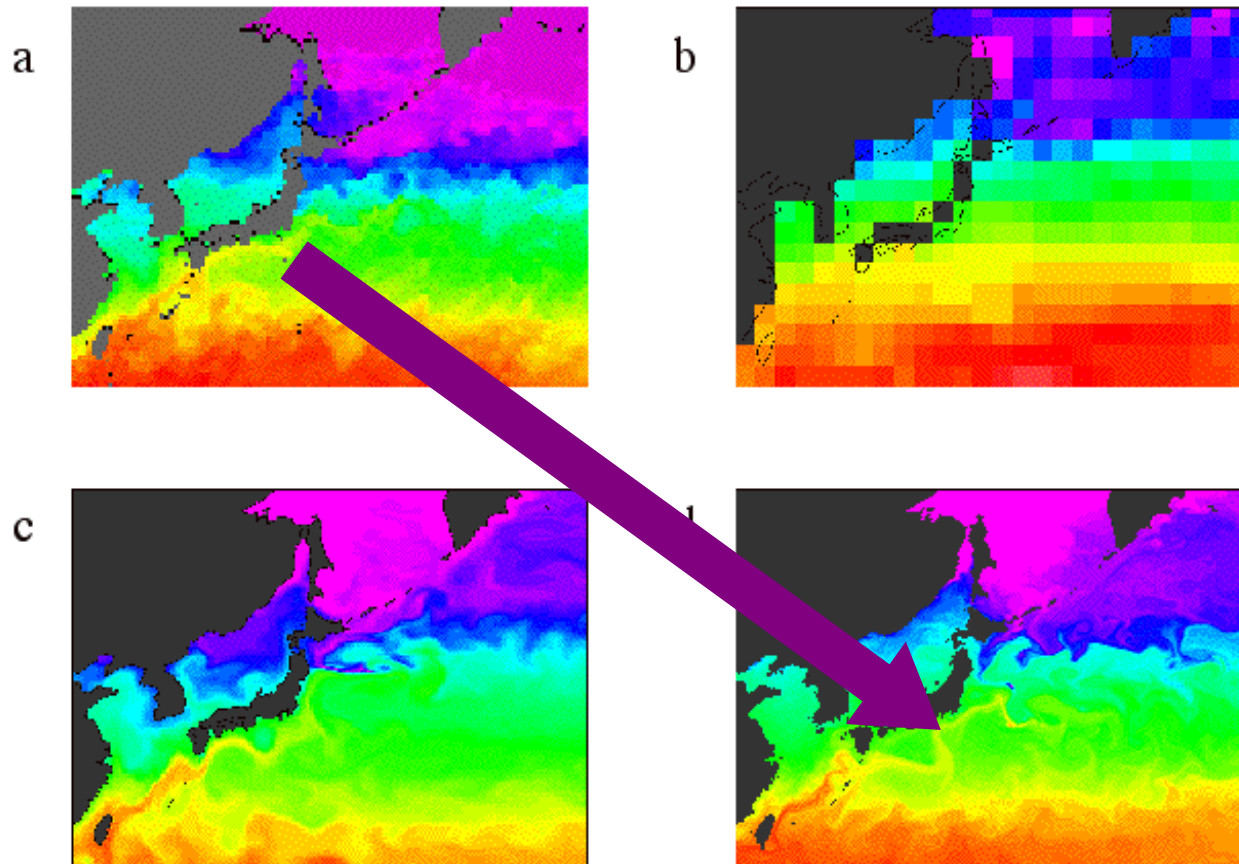
- **Understanding and predicting global climate change**
- **Exploring limits of the “Standard Model” of physics with quantum chromodynamics**
- **Designing billion-dollar accelerator facilities with mathematical optimization**
- **Probing the structure of supernovae for understanding of heavy element formation and standard candles**



# What would we do with 100-1000x more?

## *Example: predict future climates*

**Resolution of Kuroshio Current:** Simulations at various resolutions have demonstrated that, because equatorial meso-scale eddies have diameters  $\sim 10\text{-}200$  km, the grid spacing must be  $< 10$  km to adequately resolve the eddy spectrum. This is illustrated in four images of the sea-surface temperature. Figure (a) shows a snapshot from satellite observations, while the three other figures are snapshots from simulations at resolutions of (b)  $2^\circ$ , (c)  $0.28^\circ$ , and (d)  $0.1^\circ$ .



# What would we do with 100-1000x more?

## *Example: predict future climates*

- **Resolution**

- refine horizontal atmospheric scale from 160 to 40 km
- refine horizontal ocean scale from 105 to 15km

- **New physics**

- atmospheric chemistry
- carbon cycle (currently, carbon release is external driver)
- dynamic terrestrial vegetation (nitrogen and sulfur cycles and land-use and land-cover changes)

- **Improved representation of subgrid processes**

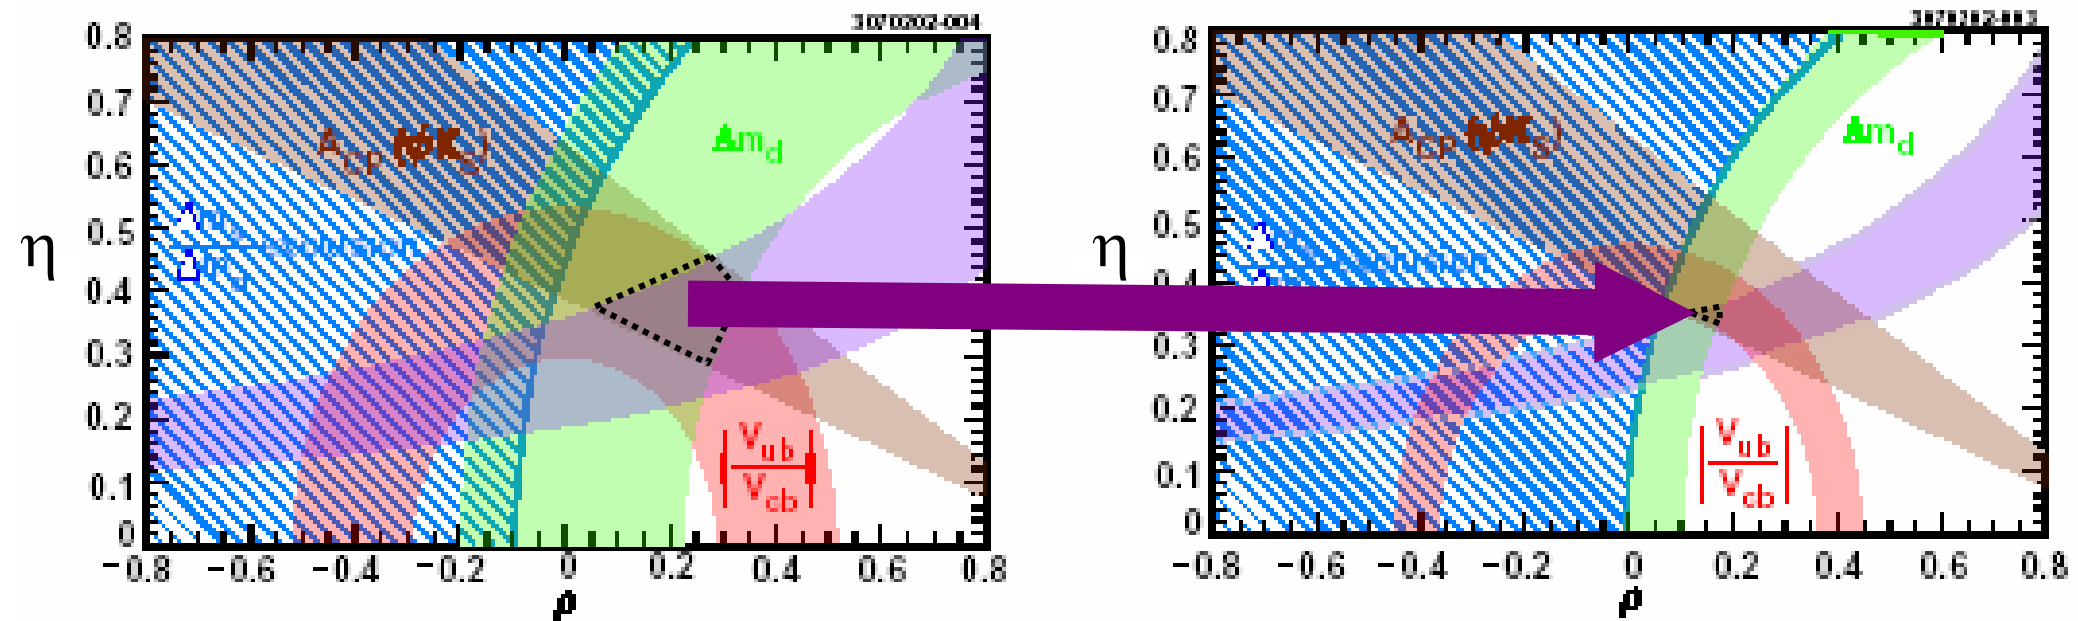
- clouds
- atmospheric radiative transfer



# What would we do with 100-1000x more?

## *Example: probe structure of particles*

*Constraints on the Standard Model parameters  $\rho$  and  $\eta$ .* For the Standard Model to be correct, these parameters from the Cabibbo-Kobayashi-Maskawa (CKM) matrix must be restricted to the region of overlap of the solidly colored bands. The figure on the left shows the constraints as they exist today. The figure on the right shows the constraints as they would exist with no improvement in the experimental errors, but with lattice gauge theory uncertainties reduced to 3%.



# What would we do with 100-1000x more?

## *Example: probe structure of particles*

- **Resolution**

- take current 4D SU(3) quantum chromodynamics models from  $32 \times 32 \times 32 \times 16$  to  $128 \times 128 \times 128 \times 64$
- explore new 5D “domain wall fermion”

- **New physics**

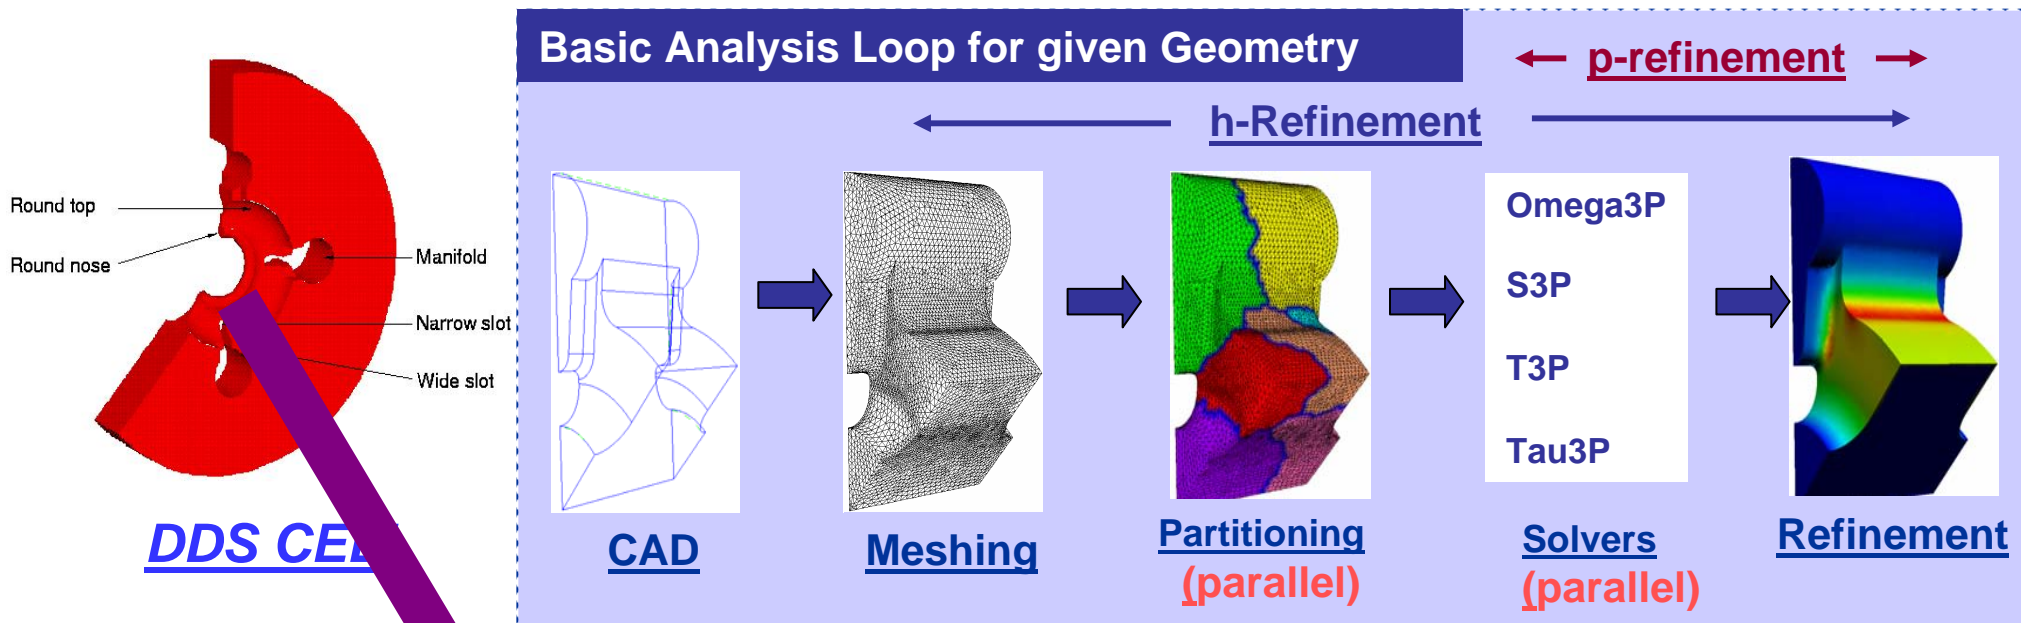
- “unquench” the lattice approximation: enable study of the gluon structure of the nucleon, in addition to its quark structure
- obtain chiral symmetry by solving on a 5D lattice in the domain wall Fermion formulation
- allow precision calculation of the spectroscopy of strongly interacting particles with unconventional quantum numbers, guiding experimental searches for states with novel quark and gluon structure






# What would we do with 100-1000x more?

## *Example: design accelerators*



*Next generation accelerators have complex cavities.* Shape optimization is required to improve performance and reduce operating cost.



# What would we do with 100-1000x more?

## *Example: design accelerators*

### ● Resolution

- complex geometry (long assemblies of damped detuned structure (DDS) cells, each one slightly different than its axial neighbor) requires unstructured meshes with hundreds of millions of degrees of freedom
- Maxwell eigensystems for interior elements of the spectrum must be solved in the complex cavity formed by the union of the DDS cells

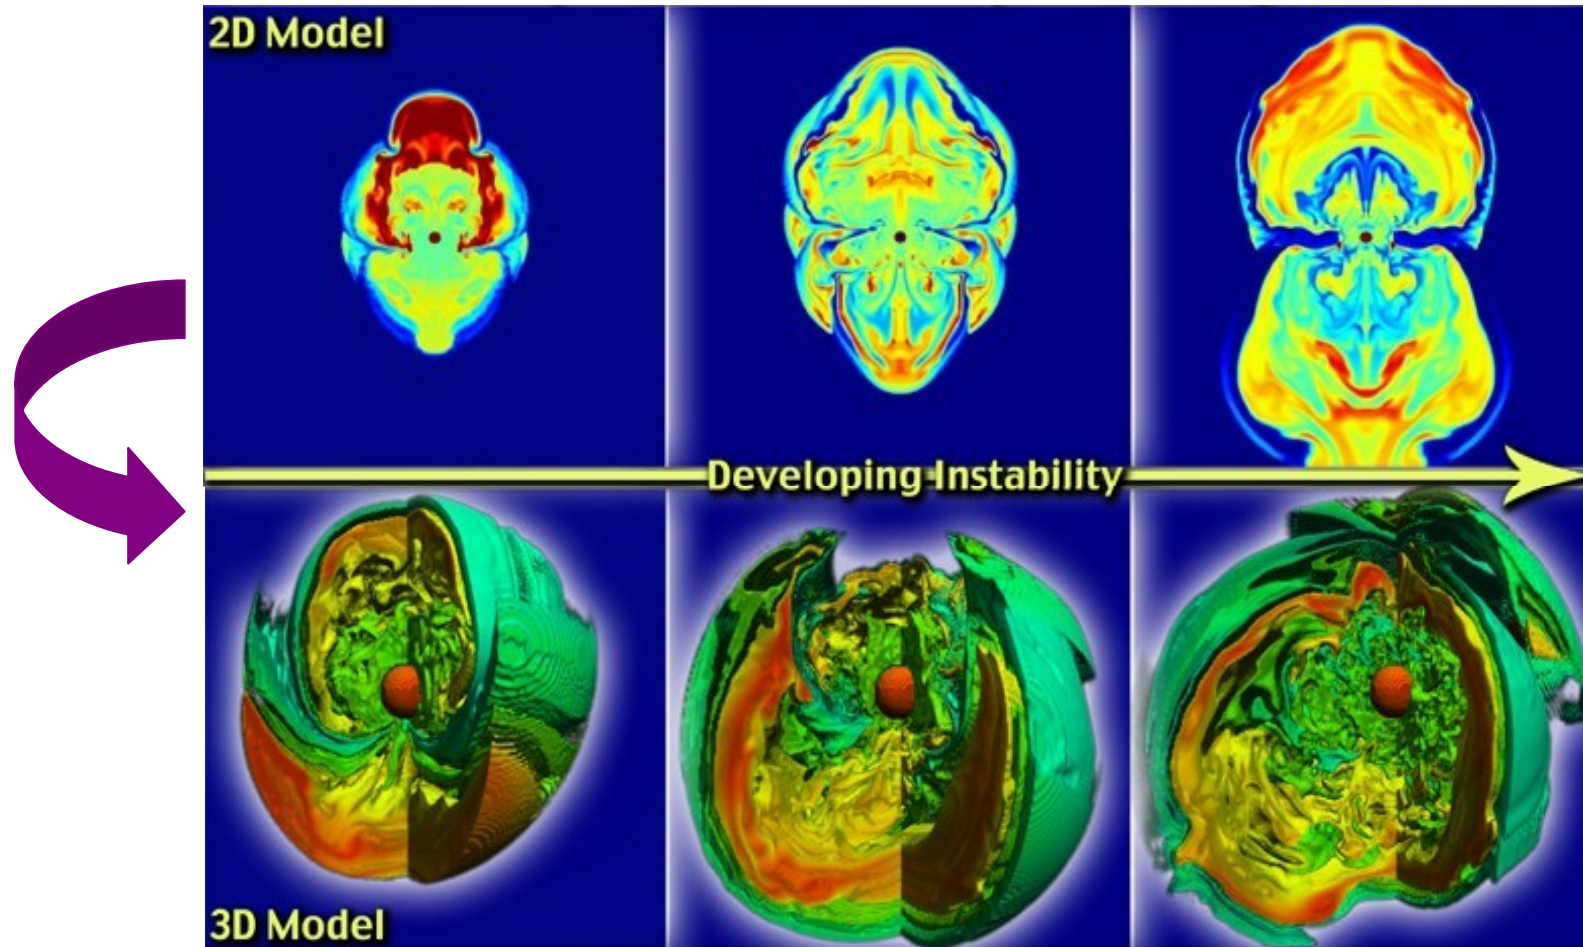
### ● Novel capability

- PDE-based mathematical optimization will replace expensive and slow trial and error prototyping approach
- each inner loop of optimization requires numerous eigensystem analyses



# What would we do with 100-1000x more?

## *Example: probe supernovae*



*Stationary accretion shock instability defines shape of supernovae and direction of emitted radiation.* Lower dimensional models produce insight; full dimensional models are ultimately capable of providing radiation signatures that can be compared with observations.



# What would we do with 100-1000x more?

## *Example: probe supernovae*

### ● Resolution

- current Boltzmann neutrino transport models are vastly under-resolved
- need at least  $512^3$  spatially, at least 8 polar and 8 azimuthal, and at least 24 energy groups energy groups per each of six neutrino types
- to discriminate between competing mechanisms, must conserve energy to within 0.1% over millions of time steps

### ● Full dimensionality

- current models capable of multigroup neutrino radiation are lower-dimensional; full 3D models are required



**A SCIENCE-BASED CASE FOR  
LARGE-SCALE SIMULATION**

VOLUME 1

OFFICE OF SCIENCE  
U.S. DEPARTMENT OF ENERGY

JULY 30, 2003



- **Chapter 1. Introduction**
- **Chapter 2. Scientific Discovery through Advanced Computing: a Successful Pilot Program**
- **Chapter 3. Anatomy of a Large-scale Simulation**
- **Chapter 4. Opportunities at the Scientific Horizon**
- **Chapter 5. Enabling Mathematics and Computer Science Tools**
- **Chapter 6. Recommendations and Discussion**

# Wrap up claims

- Simulation will become *increasingly cost-effective* relative to experiment, while never fully replacing experiment
- Simulation may define today's *limit to progress* in areas that are already theoretically well modeled
- Simulation *aids model refinement* in areas not already well modeled (via interplay with theory)
- Advanced simulation makes scientists and engineers *more productive* (can partially offset national disadvantage in workforce recruiting)




# Wrap up lessons from SciDAC

- Much high pay-off work to be done in large-scale simulation is *at the interface* between disciplines
- *Mission-oriented laboratories* and *idea-oriented universities* make good partners in developing the “science” of simulation



# On “Experimental Mathematics”



“There will be opened a gateway and a road to a large and excellent science into which minds more piercing than mine shall penetrate to recesses still deeper.”

*Galileo (1564-1642) on “experimental mathematics”*





# URLs

- **TOPS SciDAC project on solvers**

<http://www.scidac.gov/math/TOPS.html>

- **The SCaLeS report**

<http://www.pnl.gov/scales/>

