



Scientific Discovery through Advanced Computing

David Keyes

Department of Applied Physics & Applied Mathematics

Columbia University

Presentation plan

- **Are we ready to call simulation “science”?**
 - **Supporting trends**
 - **Hurdles to science by simulation**
- **DOE’s SciDAC initiative**
 - **Anatomy of a scientific simulation program**
 - **A few “stories from the trenches”**
- **Outlook**
 - **Climate for simulation**
 - **Illustrative scientific opportunities**



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. intends to 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...

- **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**



Gedanken experiment:

How to use a jar of peanut butter as its price slides downward?

- In 2005, at \$3.20: make sandwiches
- By 2008, at \$0.80: make recipe substitutions for other oils
- By 2011, at \$0.20: use as feedstock for biopolymers, plastics, etc.
- By 2014, at \$0.05: heat homes
- By 2017, at \$0.0125: pave roads 😊



The cost of computing has been on a curve *much better than this* for two decades and promises to continue for at least one more. Like everyone else, scientists should plan increasing uses for it...



Gordon Bell Prize: “price performance”

<i>Year</i>	<i>Application</i>	<i>System</i>	<i>\$ per Mflops</i>
1989	Reservoir modeling	CM-2	2,500
1990	Electronic structure	IPSC	1,250
1992	Polymer dynamics	cluster	1,000
1993	Image analysis	custom	154
1994	Quant molecular dyn	cluster	333
1995	Comp fluid dynamics	cluster	278
1996	Electronic structure	SGI	159
1997	Gravitation	cluster	56
1998	Quant chromodyn	custom	12.5
1999	Gravitation	custom	6.9
2000	Comp fluid dynamics	cluster	1.9
2001	Structural analysis	cluster	0.24

Four orders
of magnitude
in 12 years



2005 update: another order of magnitude for various graphical applications



Gordon Bell Prize: “peak performance”

<i>Year</i>	<i>Type</i>	<i>Application</i>	<i>No. Procs</i>	<i>System</i>	<i>Gflop/s</i>
1988	PDE	Structures	8	Cray Y-MP	1.0
1989	PDE	Seismic	2,048	CM-2	5.6
1990	PDE	Seismic	2,048	CM-2	14
1992	NB	Gravitation	512	Delta	5.4
1993	MC	Boltzmann	1,024	CM-5	60
1994	IE	Structures	1,904	Paragon	143
1995	MC	QCD	128	NWT	179
1996	PDE	CFD	160	NWT	111
1997	NB	Gravitation	4,096	ASCI Red	170
1998	MD	Magnetism	1,536	T3E-1200	1,020
1999	PDE	CFD	5,832	ASCI BluePac	627
2000	NB	Gravitation	96	GRAPE-6	1,349
2001	NB	Gravitation	1,024	GRAPE-6	11,550


Four orders
of magnitude
in 13 years

2005 update: > 50 Tflop/s expected by at least three DOE-led teams on BG/L




Some recent Bell “special” prizes

- **2002 Bell Prize in “special category” went to an implicit, unstructured grid structural mechanics problem (static and vibrational analysis)**
 - **60 million degrees of freedom**
 - **1 Tflop/s sustained on 3000+ processors of IBM’s “ASCI White”**
 - **in production in engineering design and certification at Sandia**




Some recent Bell “special” prizes

- **2003 Bell Prize in “special category” went to explicit, unstructured grid geological parameter estimation problem (seismic inversion)**
 - 17 million degrees of freedom in mesh; 70 billion overall
 - 1 Tflop/s sustained on 2000+ processors of HP’s “Lemieux”
 - in production use in NSF-sponsored seismology research

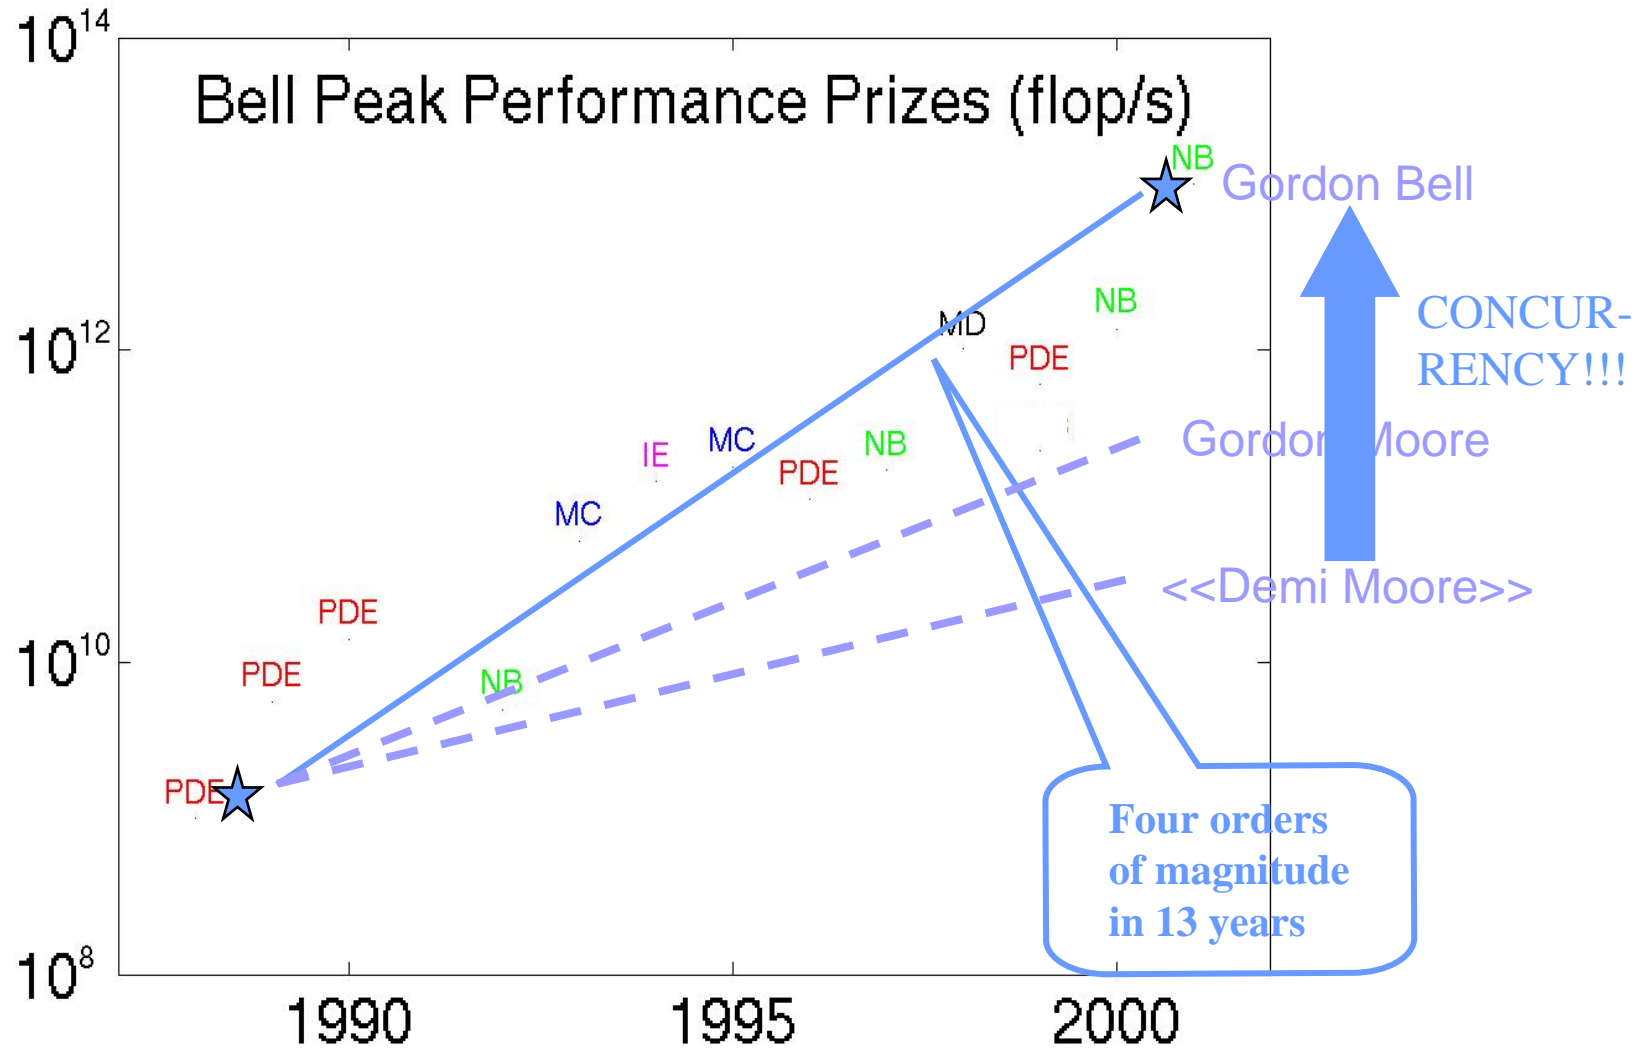


Some recent Bell “special” prizes

- 2004 Bell Prize in “special category” went to an implicit, unstructured grid large-deformation compression analysis for bone
 - half-billion degrees of freedom
 - 0.5 Tflop/s sustained on 4 thousand procs of ASCI White
 - in production at Berkeley bone mechanics lab



Gordon Bell Prize outpaces Moore's Law



IBM's BlueGene/L: 65536 dual procs, 180 Tflop/s

System
(64 cabinets, 64x32x32)

Cabinet
(32 Node boards, 8x8x16)

Node Board
(32 chips, 4x4x2)
16 Compute Cards

Compute Card
(2 chips, 2x1x1)

Chip
(2 processors)

2.8/5.6 GF/s
4 MB

5.6/11.2 GF/s
0.5 GB DDR

90/180 GF/s
8 GB DDR

2.9/5.7 TF/s
256 GB DDR


Present offer from IBM

Single cabinet
5.7 TFlop/s peak
\$2M in acad. consortium



Tflop/s-capable machines on “Top500”

<i>Year</i>	<i>#</i>
1997	1
1998	1
1999	3
2000	7
2001	17
2002	47
2003	131
2004	398



*Power
to the
people*

June 2005 list to be
published tomorrow at
ISC 2005 in Heidelberg!

ASCI Red

ASCI Blue

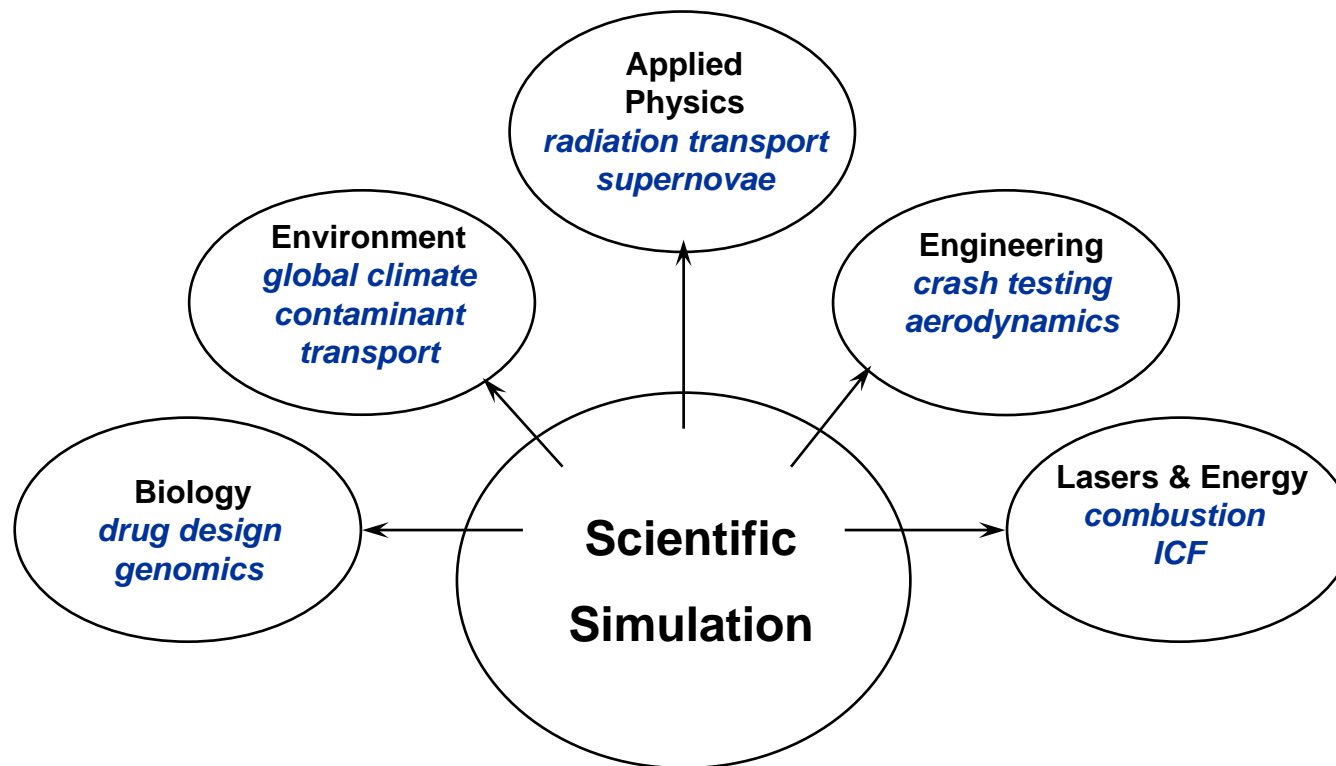
ASCI White

ESIM

BG/L



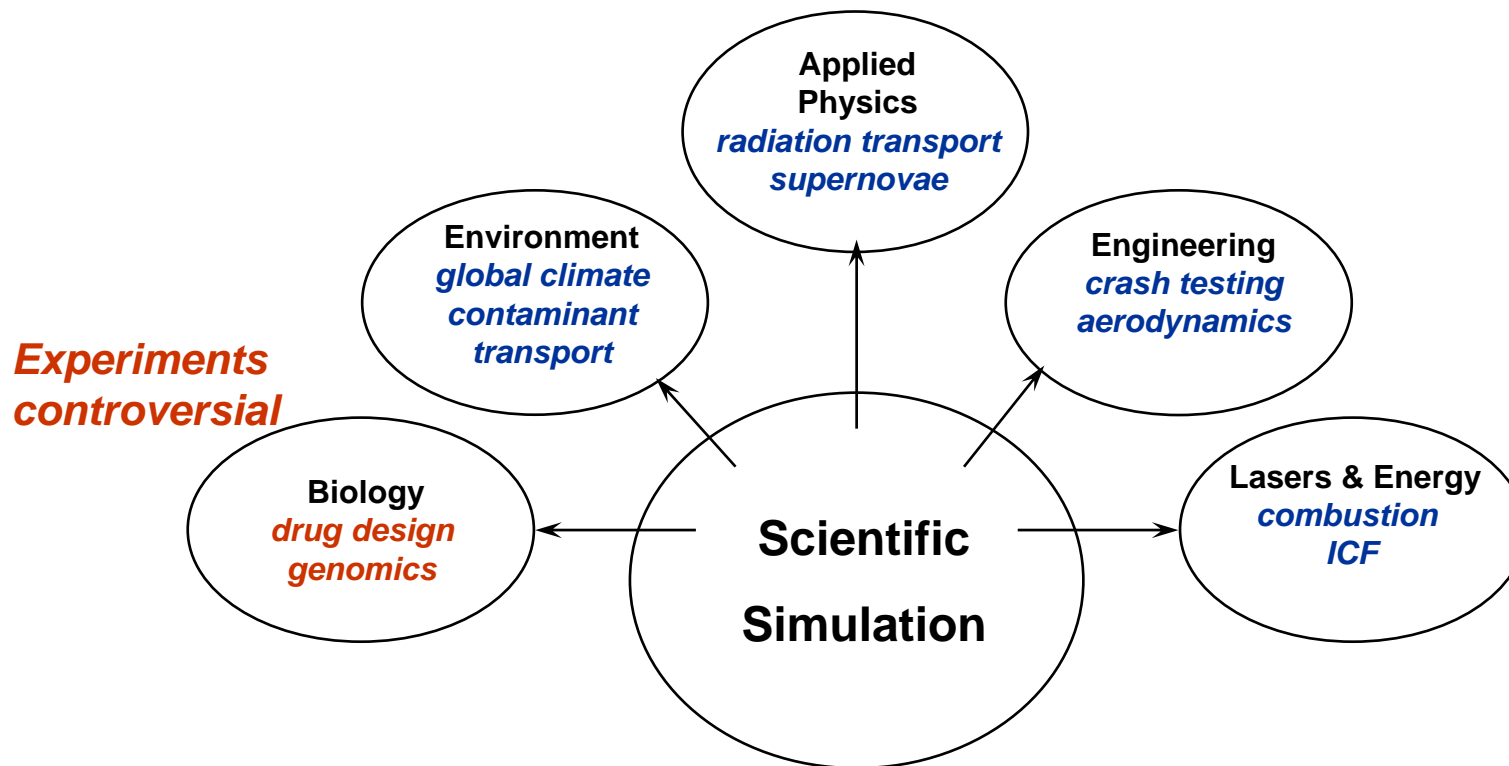
The imperative of simulation



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



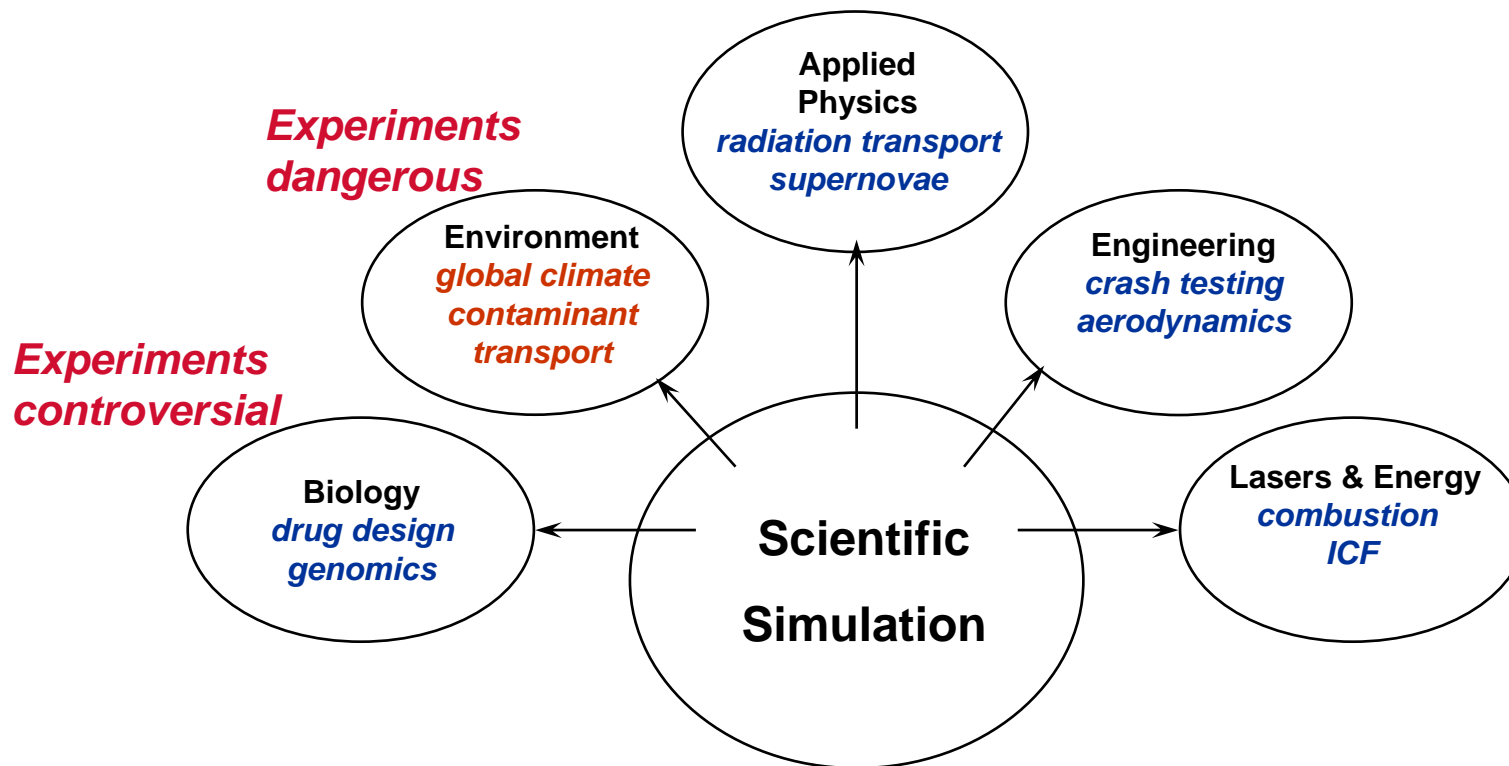
The imperative of simulation



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



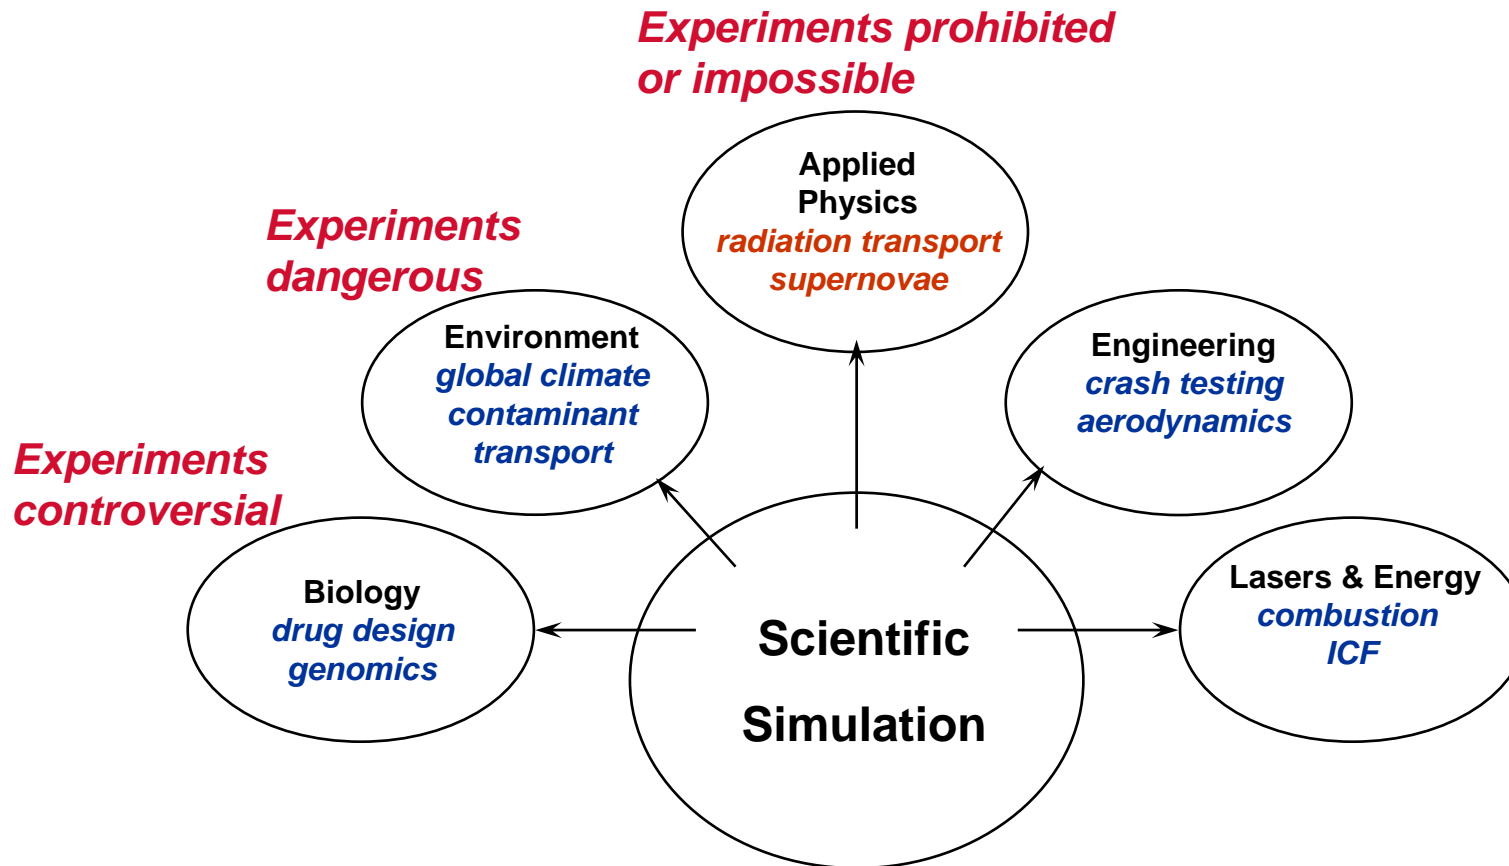
The imperative of simulation



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



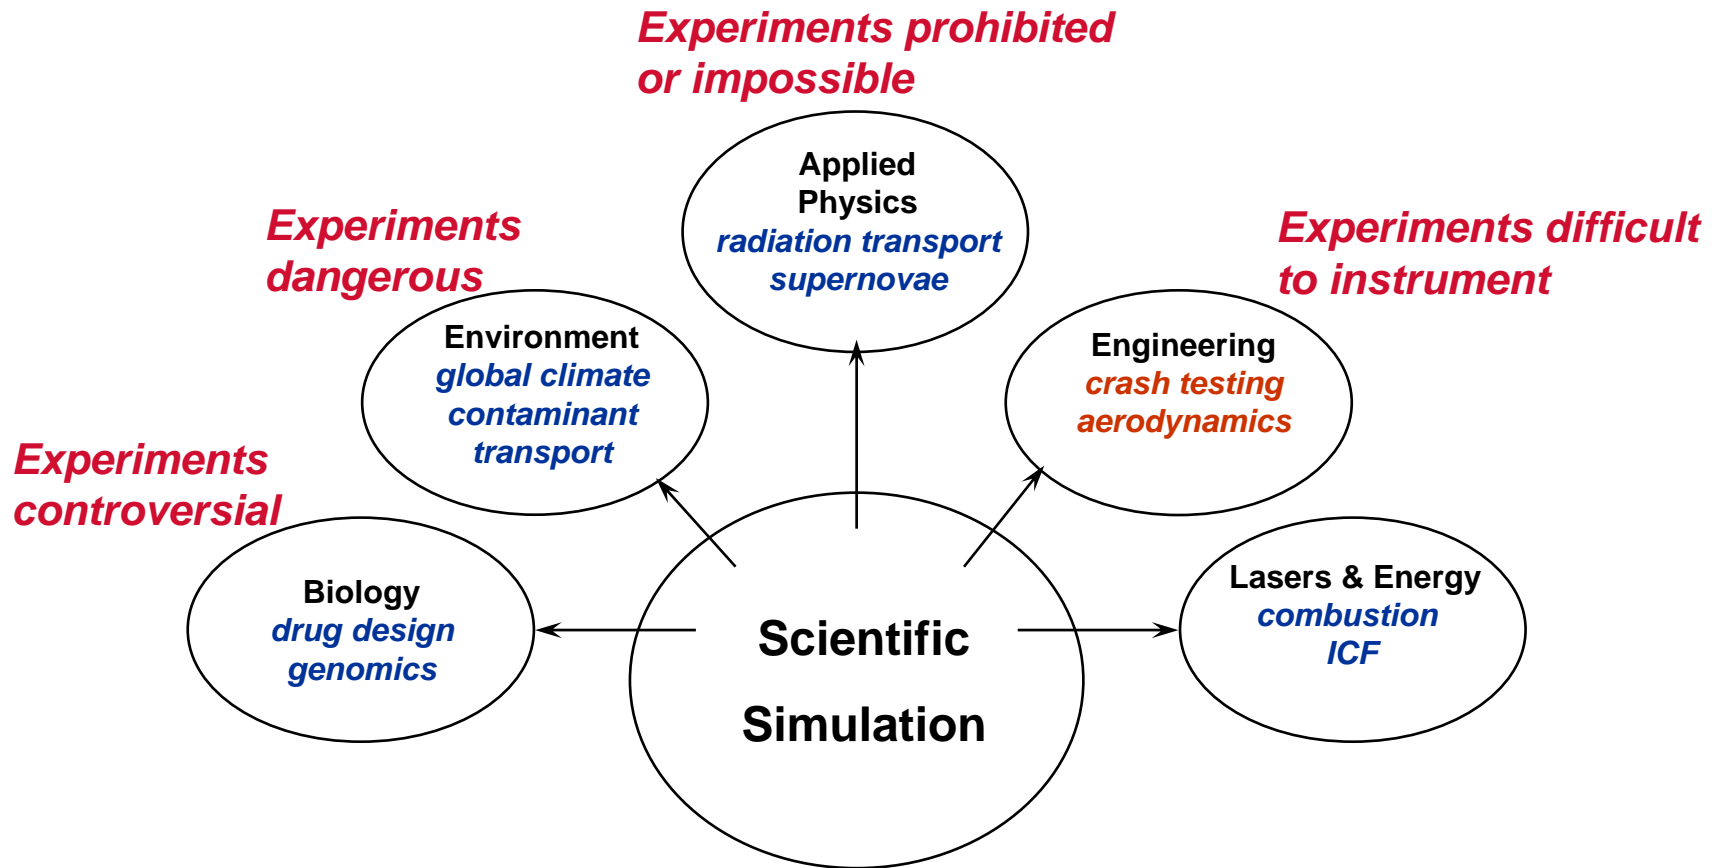
The imperative of simulation



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

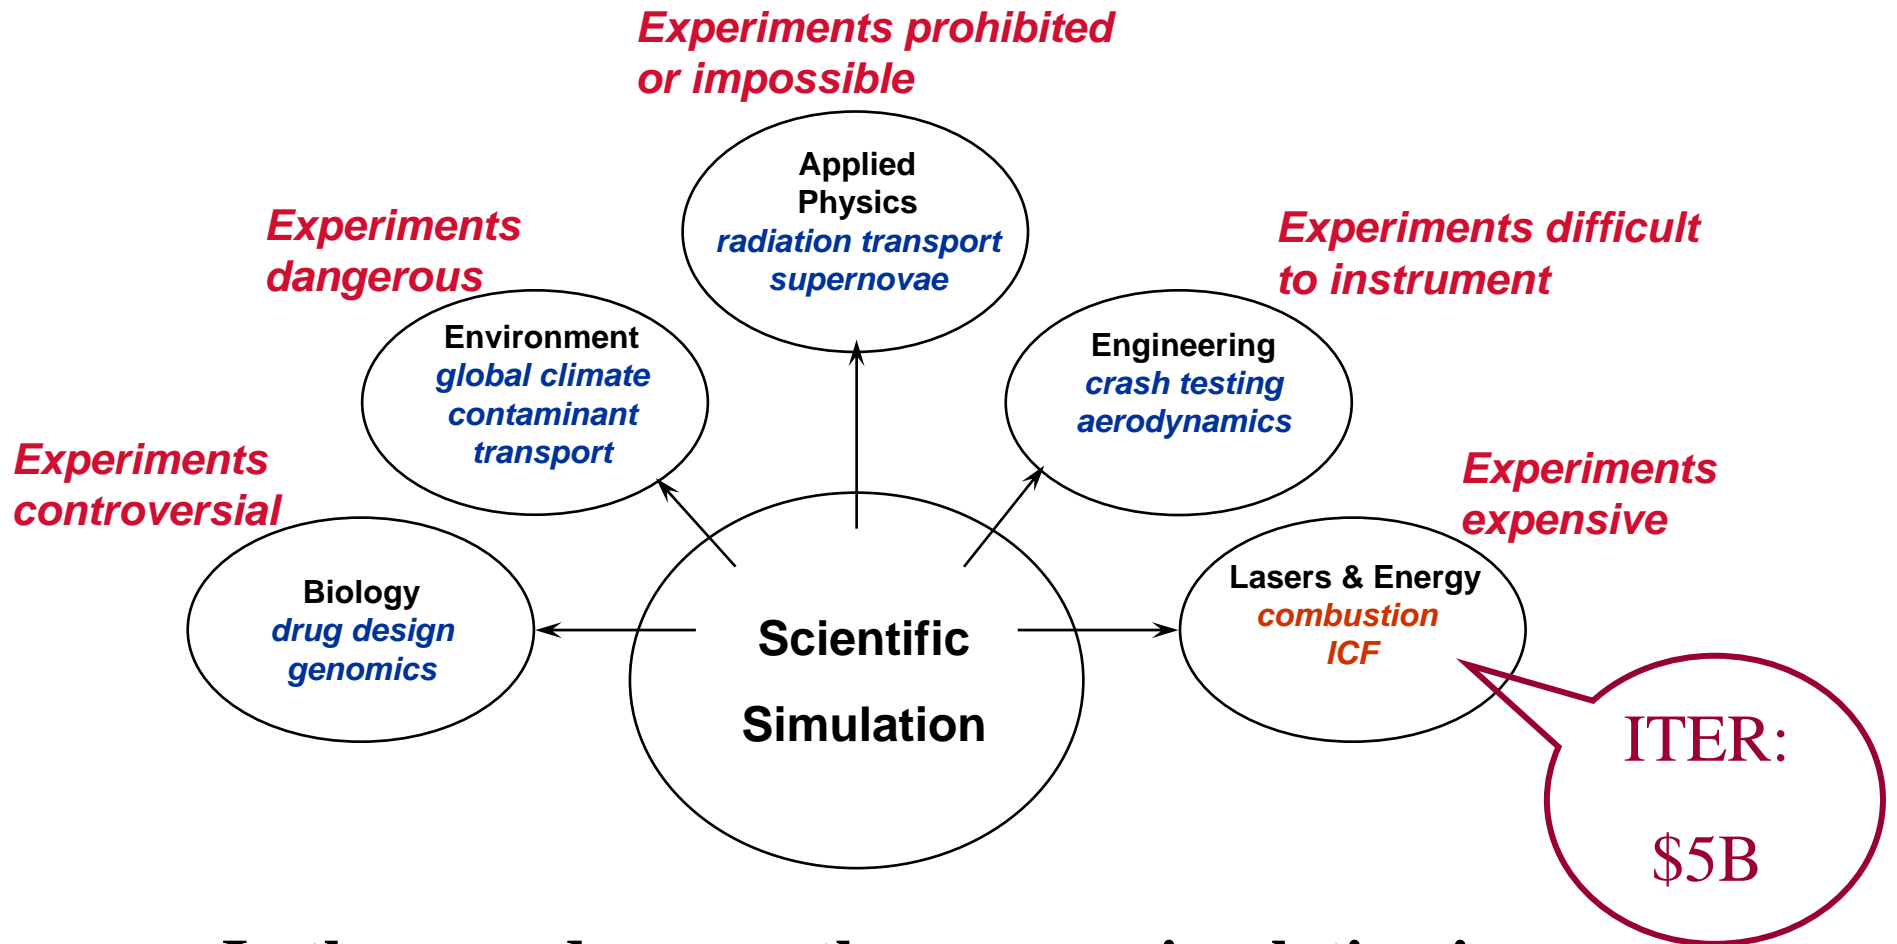


The imperative of simulation



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

- **Acceptance**

- models, inputs are often poorly known
- paltry standards for reproducibility

Need UQ methods
& fuller archiving

- **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

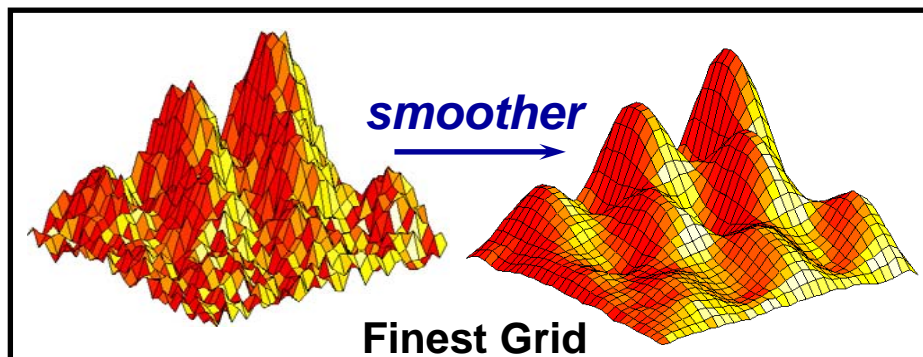


- If $n=64$, this implies an overall reduction in flops of ~16 million*

*Six-months is reduced to 1 s



Optimality from multilevel preconditioning

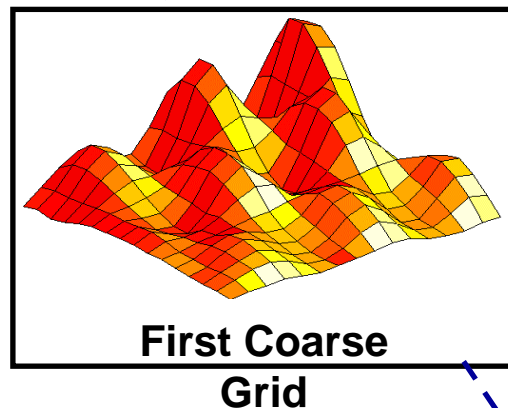


A Multigrid V-cycle

Restriction

transfer from fine to coarse grid

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



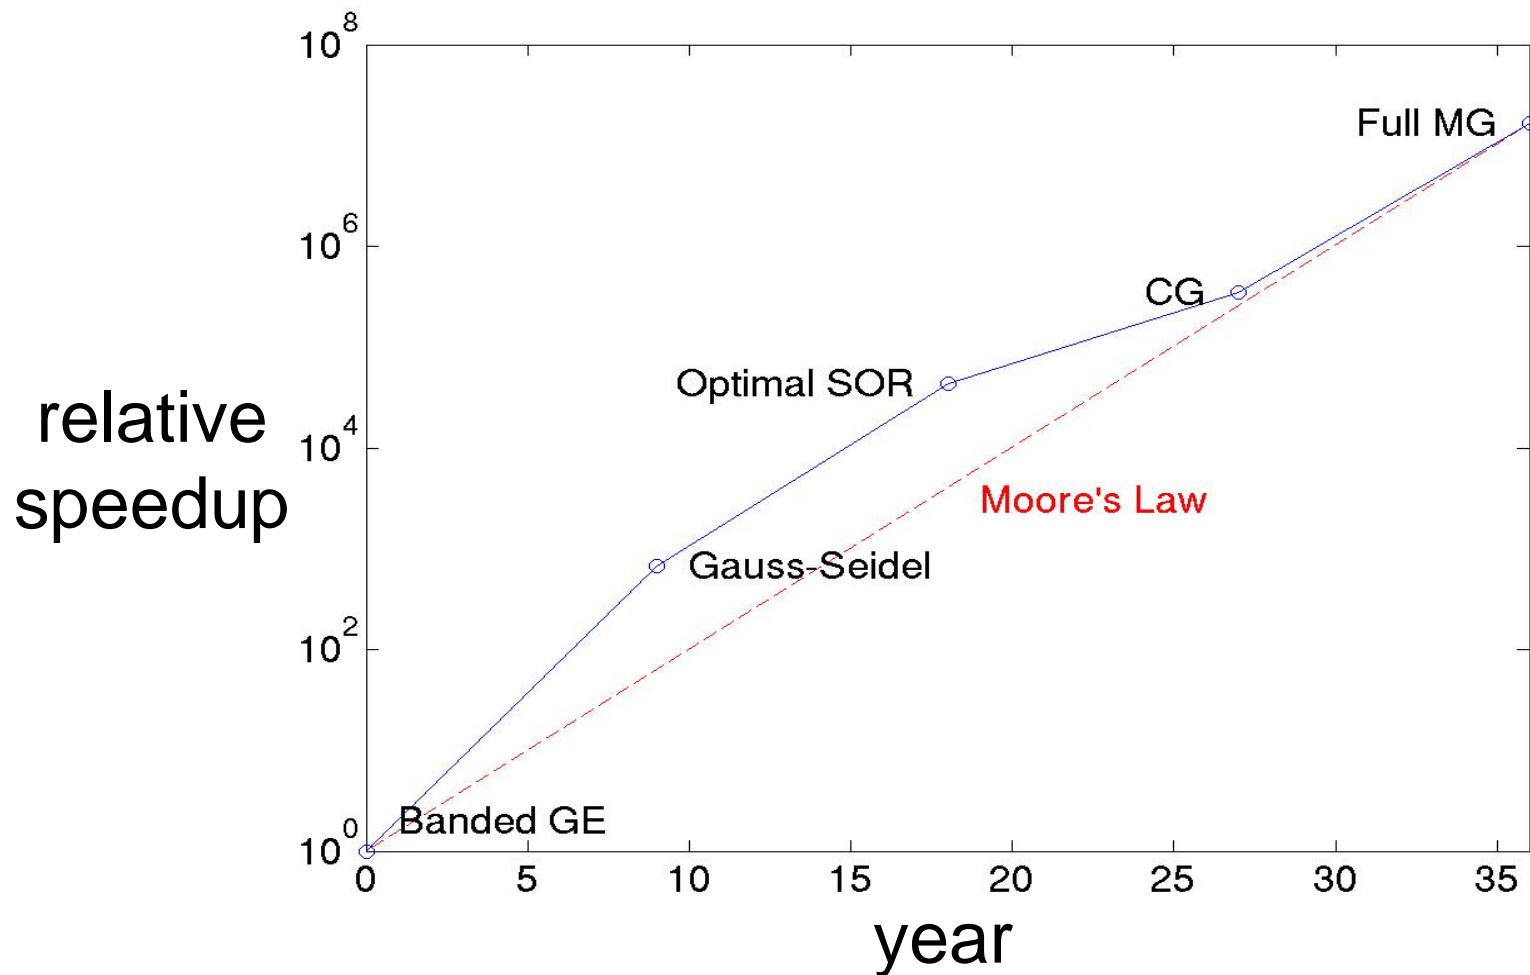
Recursively apply this idea until we have an easy problem to solve

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!



Today: a perfect season for simulation

(dates are symbolic)



1686

scientific models



1947

numerical algorithms







1976

computer architecture



1992

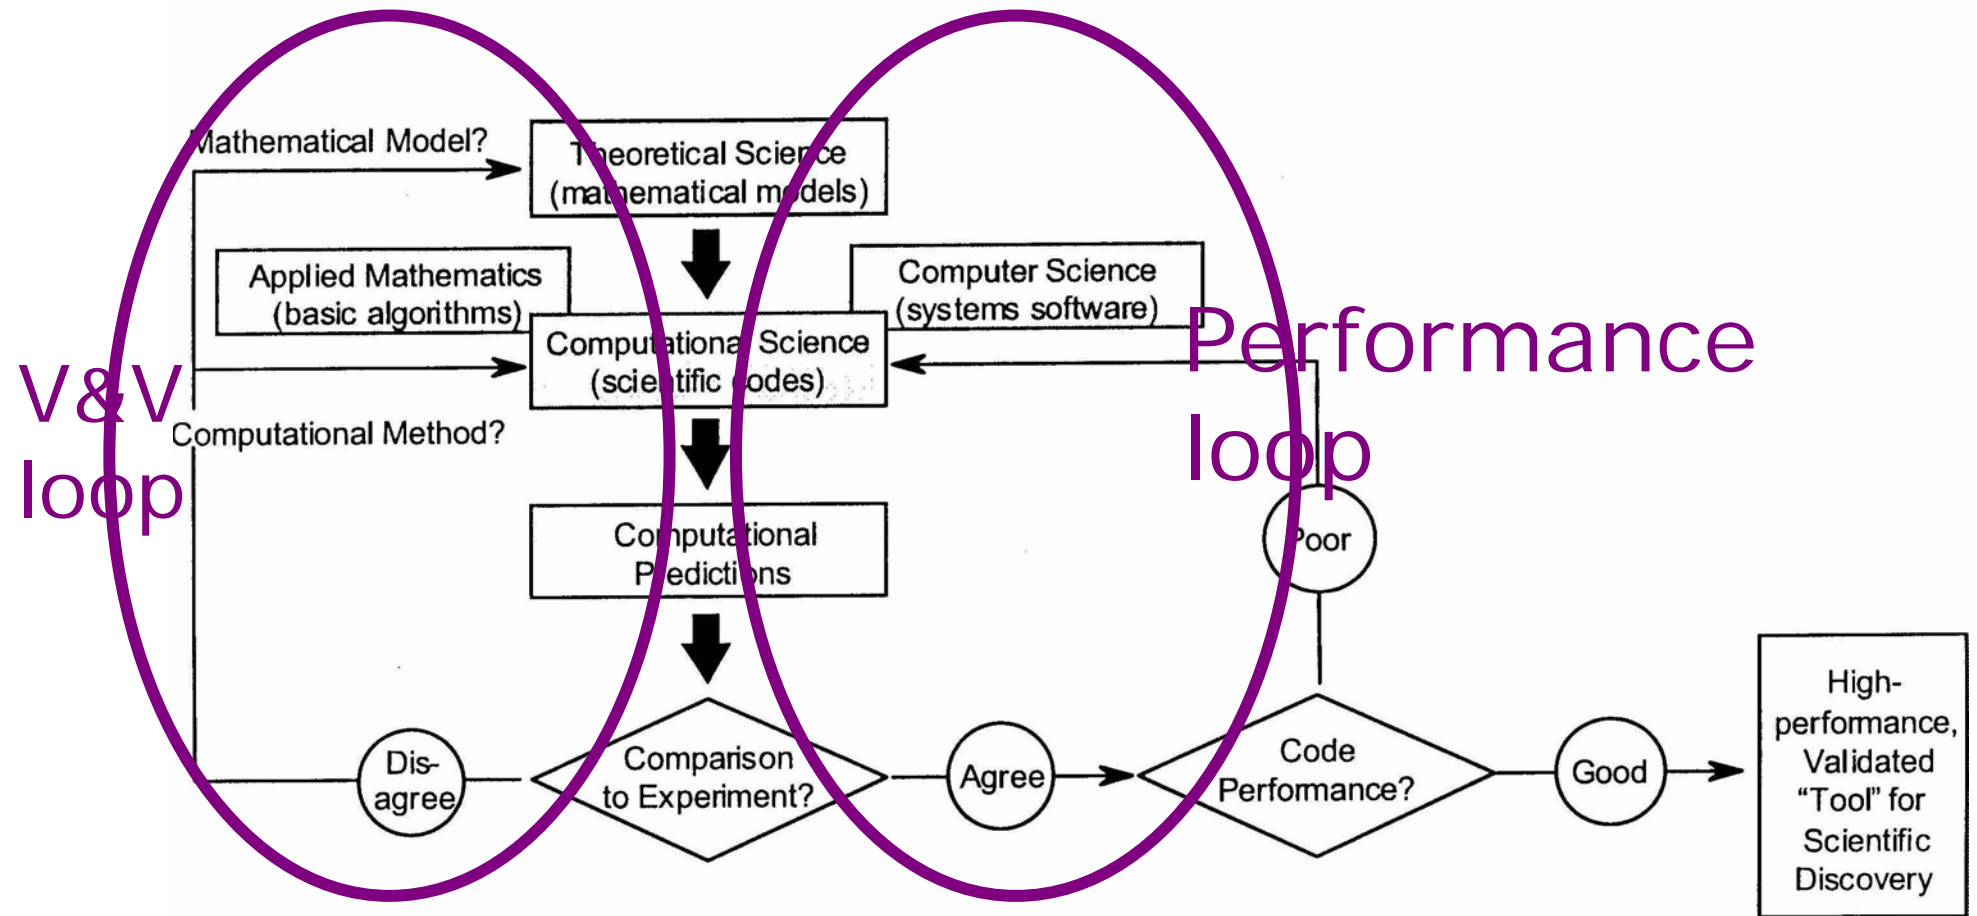
scientific software engineering



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



Designing a simulation code



How large-scale simulation is structured


- **Applications-driven**

- motivation is from applications to enabling technologies
- applications expose challenges, enabling technologies respond

- **Enabling technologies-intensive**

- in many cases, the application agenda is well-defined
- architecture, algorithms, and software represent bottlenecks

- **Most worthwhile development may be at the interface**




- **“Enabling technologies” groups to develop reusable software and partner with application groups**
- **From 2001 start-up, 51 projects share \$57M/year**
 - **Approximately one-third for applications**
 - **A third for “integrated software infrastructure centers”**
 - **A third for grid infrastructure and laboratories**
- **Plus, multi-Tflop/s IBM SP machines at NERSC and ORNL available for SciDAC researchers**




SciDAC

Scientific Discovery through Advanced Computing

4 projects in high energy and nuclear physics




14 projects in biological and environmental research




18 projects in scientific software and network infrastructure

5 projects in fusion energy science



10 projects will in basic energy sciences



Fuel Cell Stack Startup Model

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



SciDAC's Integrated Software Infrastructure Centers



Scientific Data Management



Common Component Architecture



Performance Engineering Research Center



Scalable Systems Software



Terascale Simulation Tools & Technologies

APDEC

Applied Partial Differential Equations Center



Terascale Optimal PDE Simulations




SciDAC's Grid Infrastructure Projects



DOE Science Grid



FusionGrid



Particle Physics Data Grid



GrPhyN



Earth System Grid

plus several others ...





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 50 universities in first round of projects



Collaborators in 5-year \$17M SciDAC scalable solvers project



Carnegie Mellon



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

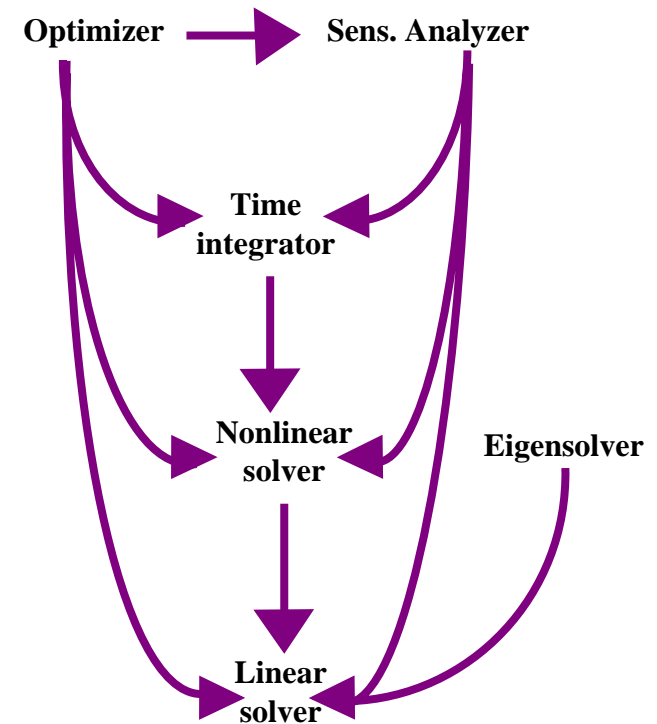
$$\min_u \phi(x, u) \text{ s.t. } F(x, u) = 0, u \geq 0$$

- Linear solvers

$$Ax = b$$

- Eigensolvers

$$Ax = \lambda Bx$$



- Software integration

- Performance optimization

→ Indicates dependence




Illustrations from computational MHD

- **M3D code (Princeton)**
 - multigrid for optimality
- **NIMROD code (General Atomics)**
 - direct elimination for robustness
- **AORSA code (Oak Ridge)**
 - new bases for storage economy

The fusion community uses more cycles on unclassified U.S. DOE computers than any other (32% of all cycles at NERSC in 2003). Well over 90% of the cycles are spent solving *linear systems* in each of these three codes, which are prime U.S. contributions to the designing of ITER.





Light a Star on Earth

The International Thermonuclear Experimental Reactor project (ITER)

The following slide is rated



for explicit equations.

No one admitted unless accompanied by a mathematician or physical scientist.

Magnetohydrodynamics:

Maxwell's equations of electromagnetics coupled to Navier-Stokes equations of fluid flow:

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}$$

$$\mathbf{E} = -\nabla \times \mathbf{A} + \eta \mathbf{J}$$

$$\mu_0 \mathbf{J} = \nabla \times \mathbf{B}$$

$$\frac{\partial n}{\partial t} + \nabla \cdot (n \mathbf{V}) = 0$$

$$\rho \left(\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V} \right) = \mathbf{J} \times \mathbf{B} - \nabla p + \nabla \cdot \nu \rho \nabla \mathbf{V}$$

$$\frac{n}{\gamma - 1} \left(\frac{\partial T}{\partial t} + \mathbf{V} \cdot \nabla T \right) = -p \nabla \cdot \mathbf{V} + \nabla \cdot n \left[(\chi_{\parallel} - \chi_{\perp}) \hat{\mathbf{b}} \hat{\mathbf{b}} + \chi_{\perp} \mathbf{I} \right] \cdot \nabla T + Q$$



Challenges in magnetic fusion


- **Conditions of interest possess two properties that pose great challenges to numerical approaches—anisotropy and stiffness**
 - **Anisotropy produces subtle balances of large forces, and vastly different parallel and perpendicular transport properties**
 - **Stiffness reflects the vast range of time-scales in the system: targeted physics is slow (~transport scale) compared to waves**
- **These have led to a family of codes specialized to numerous regimes (52 DOE codes inventoried in 2002)**



M3D: multigrid for optimality


● M3D code

- realistic toroidal geom., unstructured mesh, hybrid FE/FD discretization
- parallelized through domain decomposition w/PETSc (www.mcs.anl.gov/petsc)




● SciDAC accomplishment

- reformulation of Poisson solves to exploit symmetry and coefficient reuse
- replacement of additive Schwarz (ASM) preconditioner with algebraic multigrid (AMG) from Hypre (www.llnl.gov/CASC/linear_solvers)
- achieved mesh-independent convergence
- $4-5 \times$ improvement in solver execution time



NIMROD: direct elim. for robustness

- **NIMROD code**
 - high-order finite elements
 - complex, nonsymmetric linear systems with 10K-100K unknowns SciDAC accomplishment
- **SciDAC accomplishment**
 - replacement of diagonally scaled Krylov with SuperLU, a supernodal parallel sparse direct solver (crd.lbl.gov/~xiaoye/SuperLU)
 - 4-5 × improvement in solver execution time



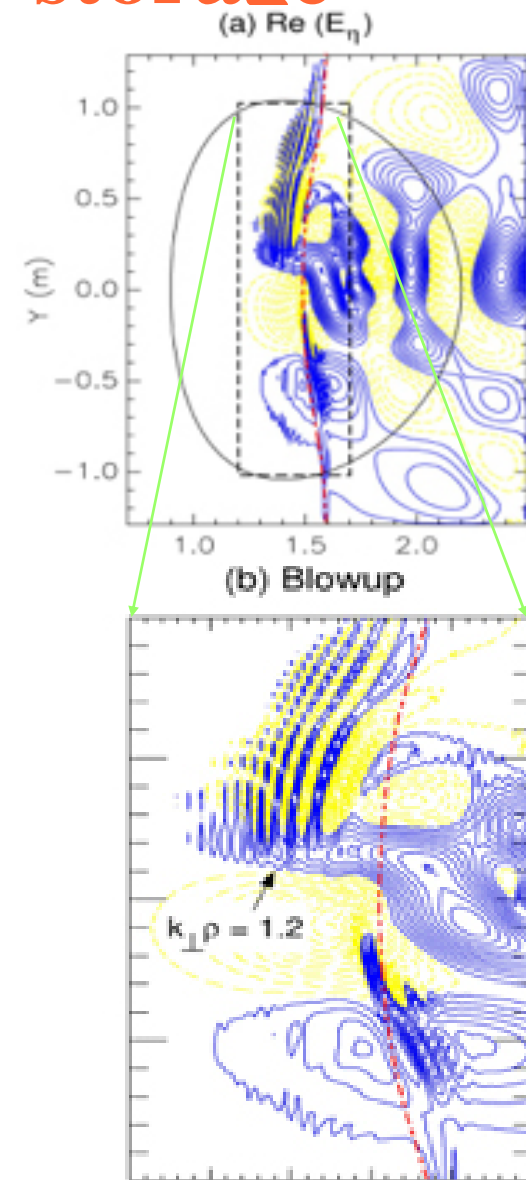
AORSA: new bases for storage

- **AORSA code**

- fully spectral harmonic Maxwell formulation for RF plasma heating
- large, dense systems with **780 GB** of matrix data


- **SciDAC accomplishment**

- replacement of Fourier formulation with physical (“configuration”) space
- 3D production runs are **27 × faster** (the linear systems are solved **100 × faster**)
- storage is only **26 GB**



“Moore’s Law” for MHD simulations

Magnetic Fusion Energy: “Effective speed” increases came from both faster hardware and improved algorithms



“Semi-implicit”:
All waves treated implicitly, but still stability-limited by transport

“Partially implicit”:
Fastest waves filtered, but still stability-limited by slower waves



Today's take on an old proverb...


“A few months in the laboratory can frequently save a few hours in the library.”

-Frank Westheimer,
Professor Emeritus of Chemistry, Harvard

“A few hours on the supercomputer can frequently save a few months in the laboratory.”



“SCaLeS” Report



Volume I (July 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

Volume II (September 2004)

- 11 chapters on applications
- 16 chapters on enabling technologies

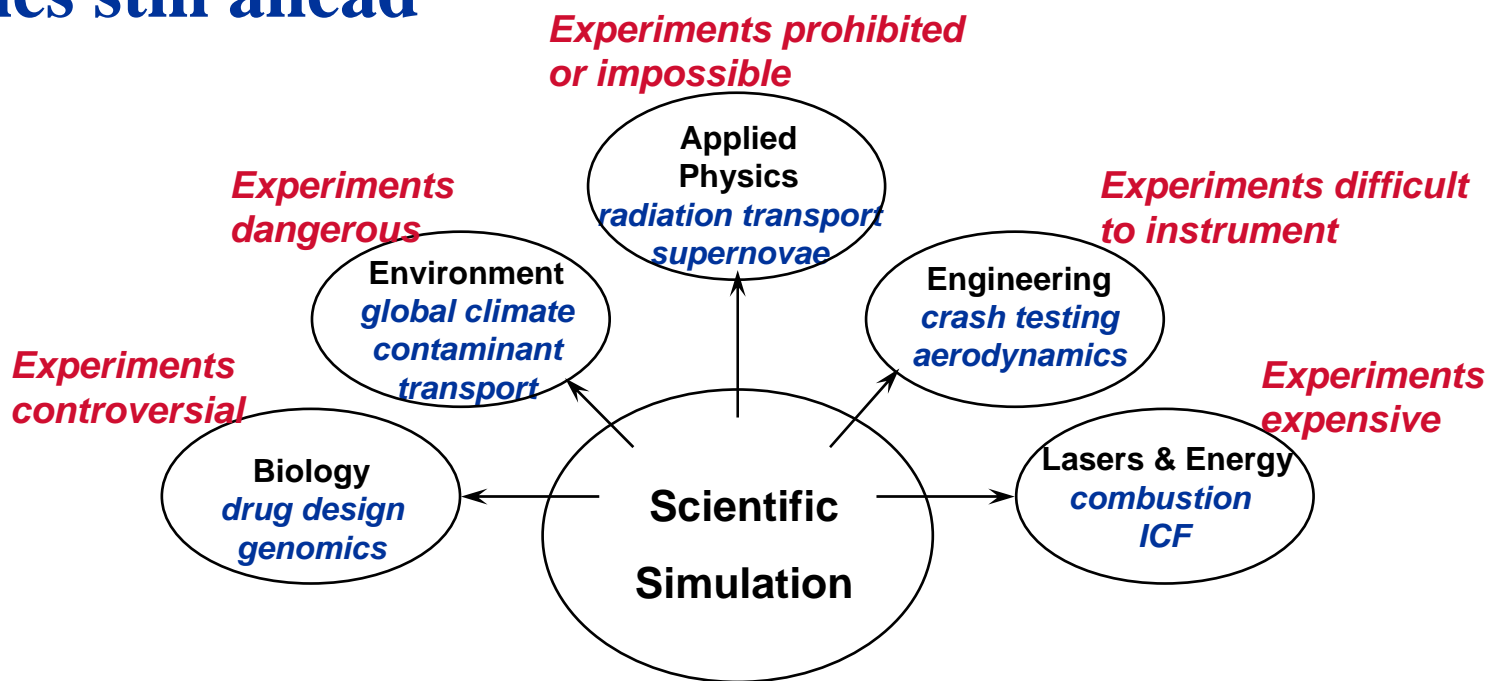
Climate for scientific simulation

- **Much federal pulse-taking ...**
 - Cyberinfrastructure (NSF, 2003)
 - SCaLeS (DOE, 2003)
 - HECRTF (Interagency, 2004)
 - Future of Supercomputing (NAS, 2005)
 - PITAC-2 (Interagency, 2005)
 - Science-based Engineering Simulation (NSF, 2005, to appear)
- **... but threatened decline in federal support**
 - NSF supercomputer centers, DARPA, ASCI
- **Labs leading in hardware and mission “buy-in”**
- **Professional societies, conferences, journals rising to the opportunities**
- **Universities slowing catching up with CS&E certificates and degrees**



Scientific “market share” must increase

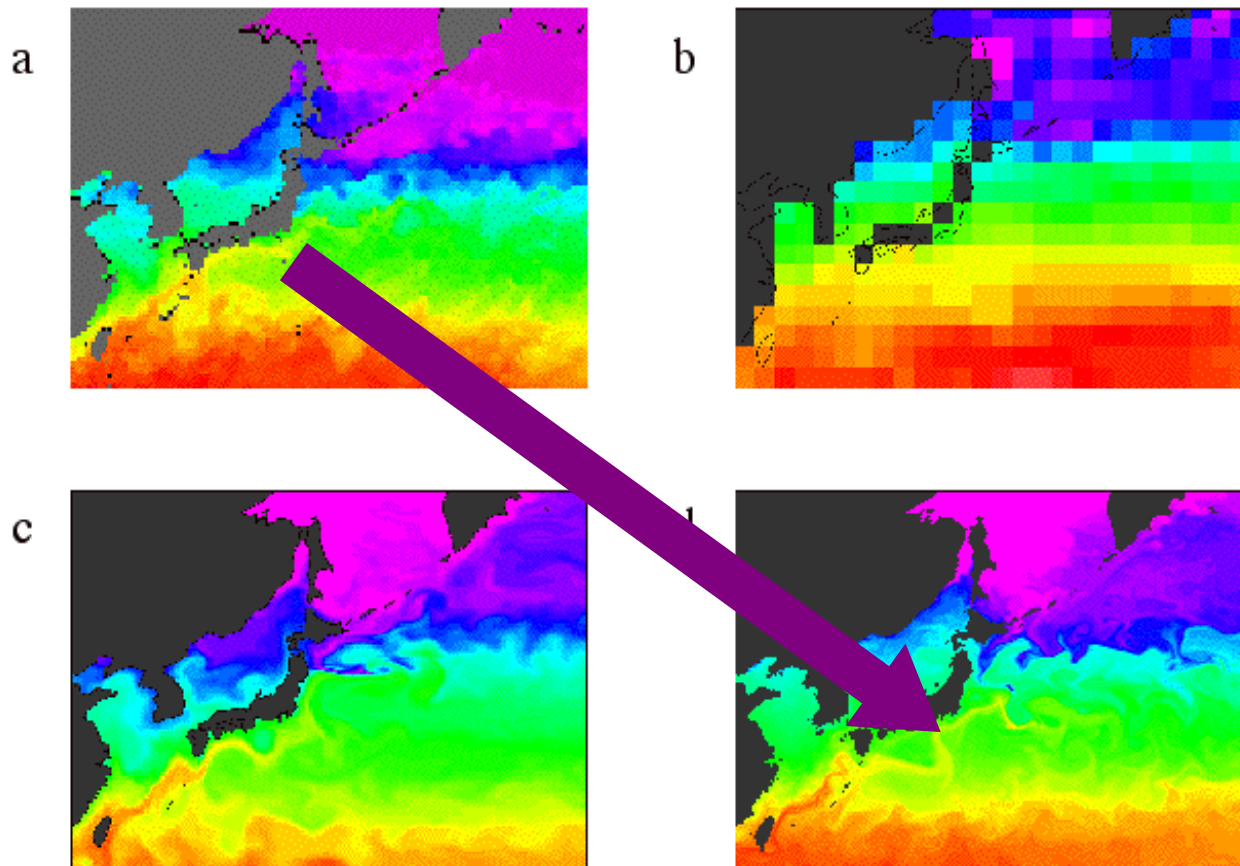
- **Cost curves crossing in every domain**
 - Accelerators to wind tunnels, drug design to reservoir development
- **Tools improving**
 - Capability pushing upwards, training barrier extending downwards
- **Best stories still ahead**



What would we do with 100-1000x more?

Example: predict future climates

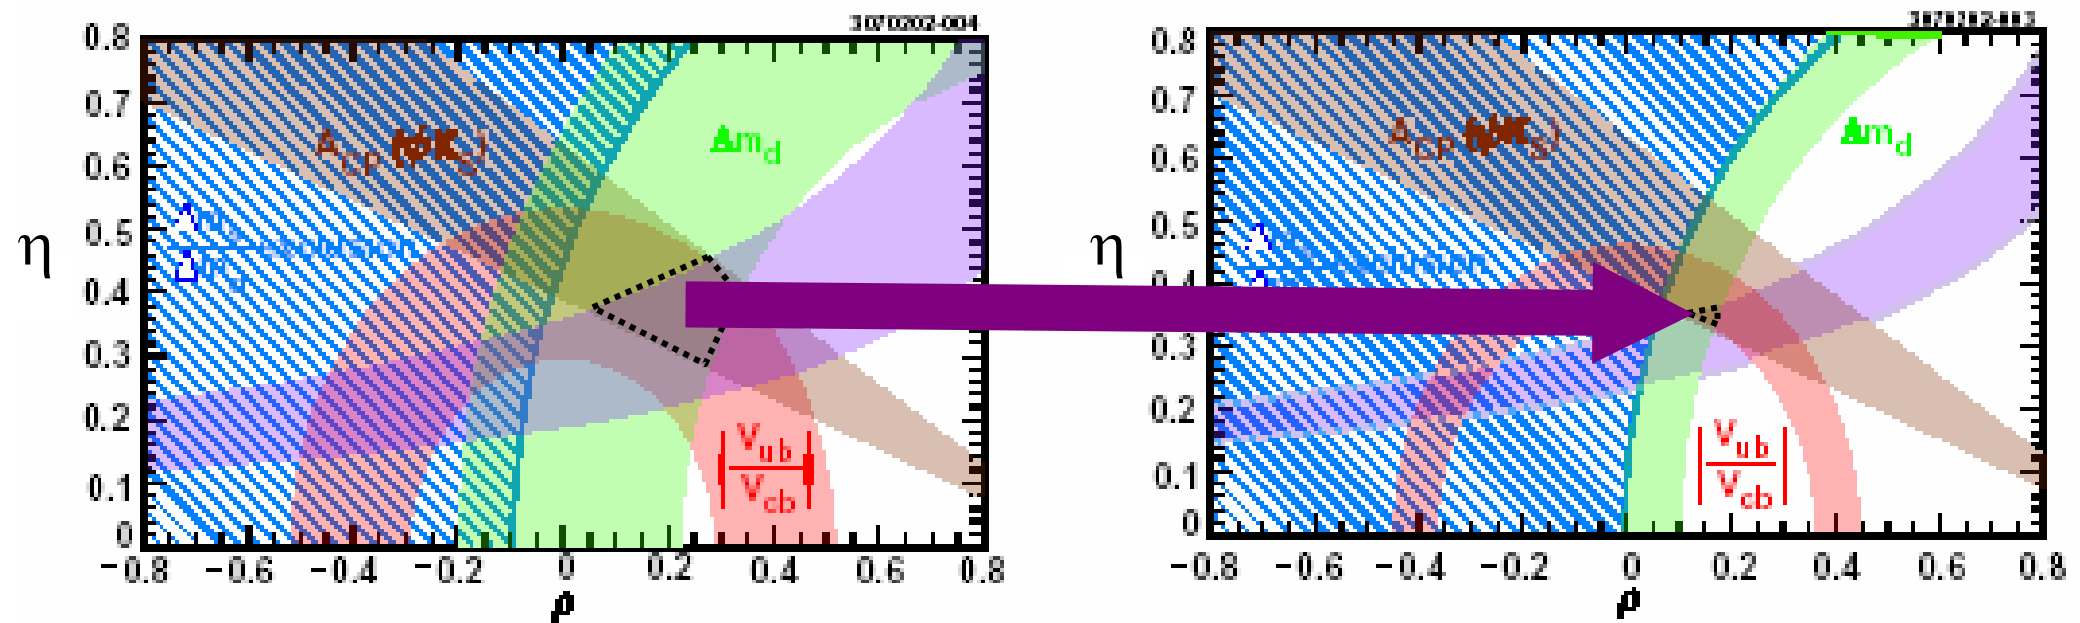
Resolution of Kuroshio Current: Simulations at various resolutions have demonstrated that, because equatorial meso-scale eddies have diameters $\sim 10\text{-}200$ km, the grid spacing must be < 10 km to adequately resolve the eddy spectrum. This is illustrated in four images of the sea-surface temperature.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



What would we do with 100-1000x more?

Example: probe the structure of particles

Constraints on the Standard Model parameters ρ and η . 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%.



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*




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”



Related URLs

- **TOPS project**

<http://www.tops-scidac.org>

- **SciDAC initiative**

<http://www.science.doe.gov/scidac>

- **SCaLeS report**

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



EOF

