Frontiers of Extreme Computing 2007 Applications and Algorithms Working Group

October 25, 2007

Horst Simon (chair), David Bailey, Rupak Biswas, George Carr, Phil Jones, Bob Lucas, David Koester, Nathan Price, Joshua Schrier, Mark Stalzer (co-chair), David Womble

Outline

- Sample Applications
- Exascale Applications Characteristics
- Zettascale Applications with Societal Impacts
  - Climate Modeling
  - Personalized Medicine and Phylogenomics
  - Fusion and Astrophysics
  - Decision Making
  - Nano/Material Science, Chemistry
- Adaptivity
- Requirements and Recommendations

Sample Applications

- Climate modeling.
- Biology: spread/analysis of diseases; phylogenetics.
- Nano/materials: simulate from first principles a transistor/quantum computer.
- Chemistry: beyond Born-Oppenheimer.
- Economic/behavioral models: Baghdad in faster than real-time.
- Cognitive models of brain: vision; mapping out functions.
- Medical applications: personalized medicine.
- Astrophysics.

Sample Applications, Cont.

- Engineering: crash modeling.
- Decision support.
- Calibration of large instruments (LHC, LIGO).
- HEP and nuclear physics.
- Real-time weather: storms, tornados, icing; disaster planning.
- Combustion.
- Nuclear energy: fission; fusion.
- Stockpile stewardship (ASC).
- Oil&gas exploration.
- Earth tomography.
- Math research.

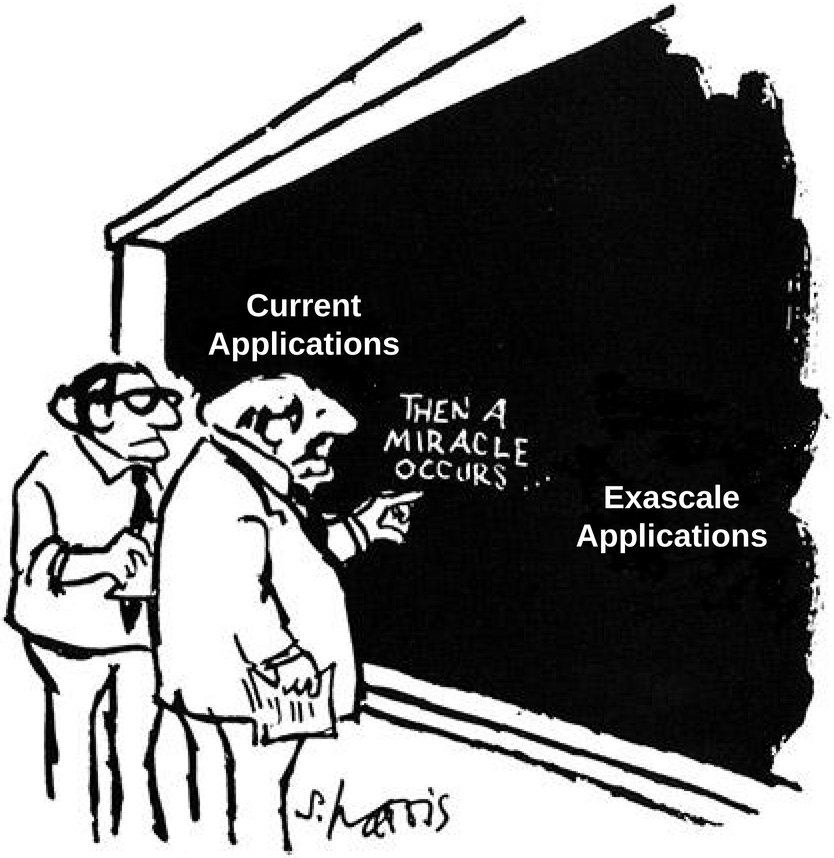

"I think you should be more explicit here in step two."

Exascale Application Characteristics

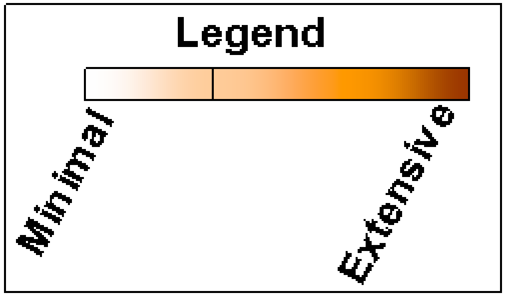

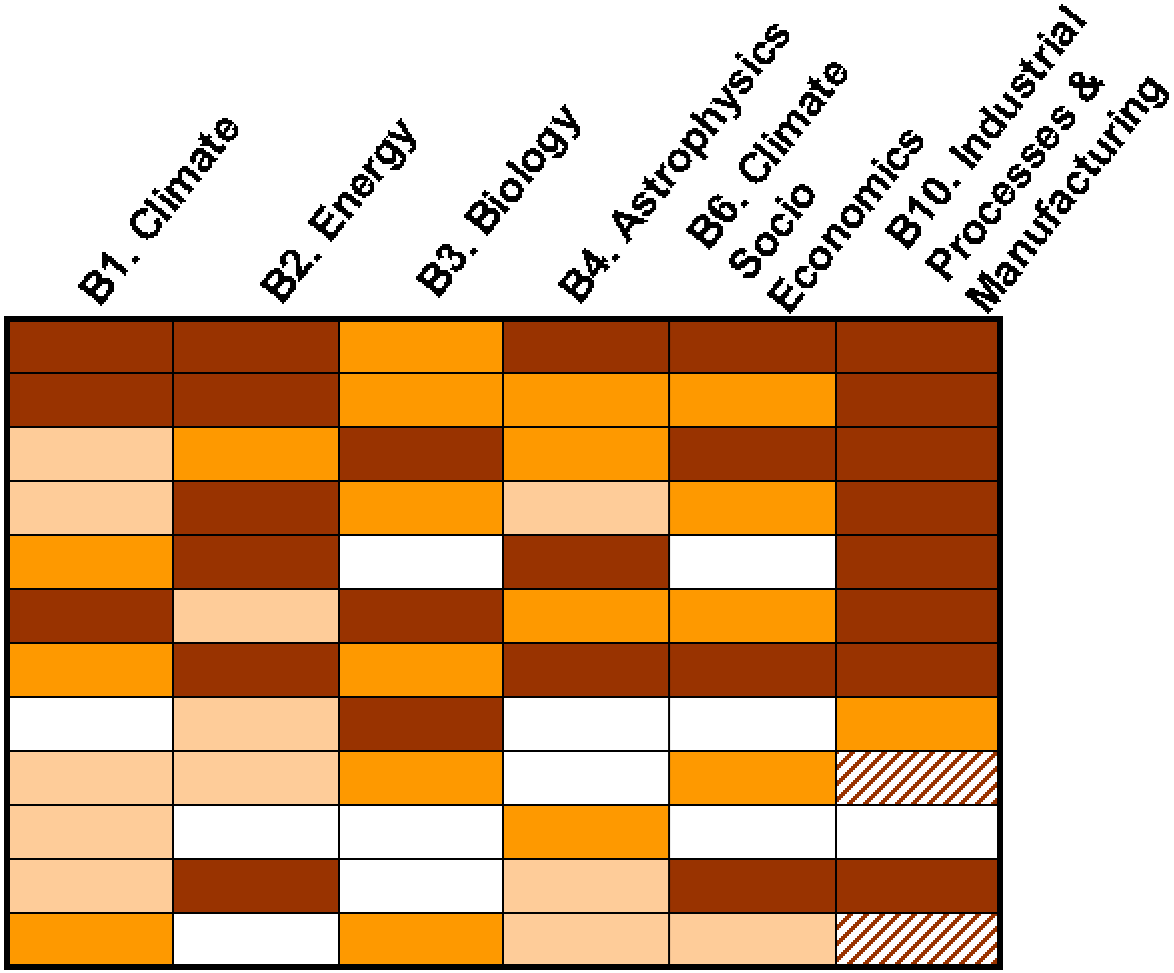

	B1. Climate	B2. Energy	B3. Biology	B4. Astrophysics	B6. Climate Socio Economics	B10. Industrial Processes & Manufacturing
Multi-Scale, Multi-Model	▓	▓	▒	▓	▓	▓
V&V & Uncertainty Quantification	▓	▓	▒	▒	▒	▓
Optimization	░	▒	▓	▒	▓	▓
Implicit Solvers	░	▓	▒	░	▒	▓
Adaptive Mesh Refinement	▒	▓		▓		▓
Data analysis	▓	░	▓	▒	▒	▓
Linear and nonlinear solvers	▒	▓	▒	▓	▓	▓
Graph algorithms		░	▓			▒
Agent-based modeling	░	░	▒		▒	▨
High-precision arithmetic	░			▒		
High-dim. parameter spaces	░	▓		░	▓	▓
Data assimilation	▒		▒	░	░	▨

Application characteristics from Group B7 Mathematics and Algorithms, David Bailey, et.al.

Climate Modeling

- Application and impact on society
  - Provide projections of future climate to policy makers to support decisions on energy, mitigation and adaptation strategies.
  - Require accurate models of both physical and biogeochemical climate system.
  - Require projections at regional spatial scales.
- Algorithms
  - Coupled multi-model application.
  - Wide variety of algorithms.
  - Largely explicit forward integration of PDEs.
  - Future needs: implicit time, efficient methods for large tracer counts, more efficient strategies for grids, discretizations and decomposition.

Computing Needs (10^{10}-10^{12})

Resolution (10^3-10^5)

- x10 each x&y, x10 timestep, x5-10 z
- Regional prediction (10km)
- Eddy resolving ocean (10km)

Completeness (10^2)

- Biogeochem (30-100 tracers, interactions)
- Ice sheets

Fidelity (10^2)

- Better cloud processes, dynamic land etc.

Increase length/number of ensembles (10^3)

- Run length (x100)
- Number of scenarios/ensembles (x10)
- Data assimilation (3-10x)

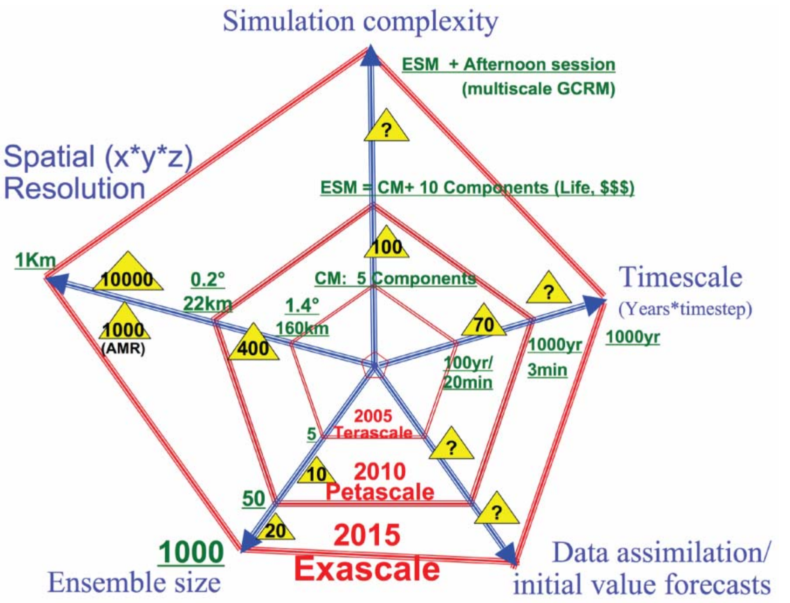

Storage/Data Requirements

- Currently 10 GB/year output
  - Similar multipliers (min x1000).
  - Staged migration to archival systems.
  - 2x to move initially, recall many times.
  - At 24 years/day => 10 GB/hour I/O rates.
  - Much higher data input rates if data assimilation required.
- Analysis
  - Requires large memory systems.
  - Need lots of intermediate disk (don’t want to analyze from archival systems).
  - Distributed data (IPCC multi-model analyses).

Personalized Medicine

- Personalized Medicine - Identifying genomic predictors of disease from individual human genomes.
- Resources: Possible sets of mutation sets of size 6 (e.g. typical for cancer) in the human genome ~ $10^{55.}$
  - Thus, we will need significant computational resources and efficient approximation algorithms to best deliver on personalized medicine (e.g. for cancer).
  - For even a simple binary classification, would thus need >> 185 individual genomes to begin to have hopes of separating signal from noise – i.e. to believe good patterns that are found (i.e. $10^{55} \sim 2^{185}$) (and this assumes simplistic yes/no model selection).
  - But there is also a decreasing probability of finding good patterns as number of classifications is done.

Phylogenomics

- Phylogenomics - learning about the origin and history of life.
- Resources: Phylogenomic calculations are NP-complete; approximate solutions get better with increased computational power, as more accurate but computationally intensive algorithms are used O(N,NlogN, $N^{2,}$ …).

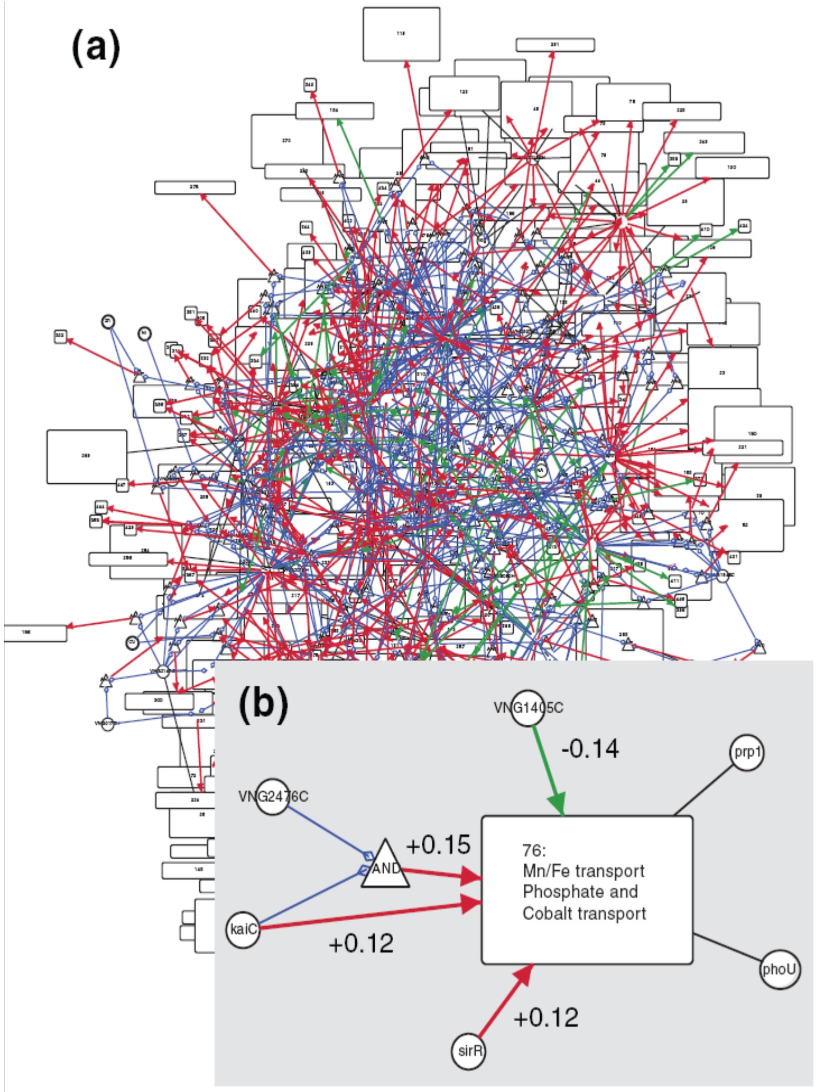

Fusion Code Characteristics

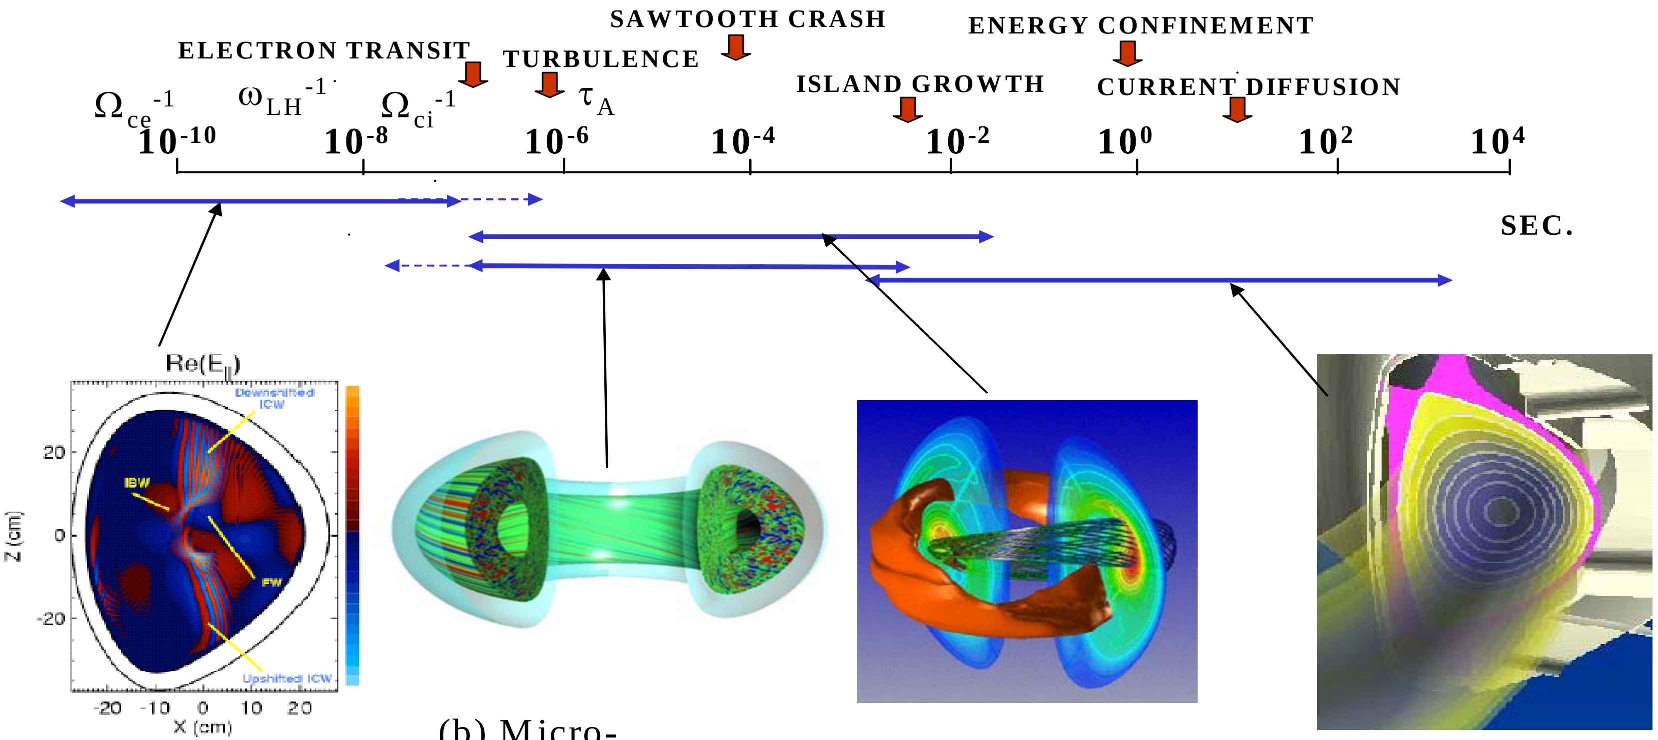

(a) RF codes

(b) Micro-turbulence codes

(c) Extended-MHD codes

(d) Transport Codes

- **Multi-physics, multi-scale computations.**
- **Numerous algorithms and data structures.**
- **Regular and irregular access computations.**
- **Adaptive mesh refinement.**
- **Advanced nonlinear solvers for stiff PDEs.**

Fusion Requirements

Tokamak turbulence (GTC) - For ITER experiment, etc:

- Grid size: 10,000 x 4000 x 256, or about 10^{10} gridpoints.
- Each grid cell contains 8 particles, for total of 8 x 10^{10}.
- 200,000 time steps required.
- Improved plasma model (increases by 10-100X).
- Total cost: 6 x 10^{21} flop = 1 hours on 1 Eflop/s system; 10 Pbyte main memory.

All-Orders Spectral Algorithm (AORSA) - to address absorption of RF electromagnetic waves in plasmas.

Present Day:

- 120,000 x 120,000 complex linear system requires 230 Gbyte memory, 1.3 hours on 1 Tflop/s.
- 300,000 x 300,000 linear system requires 8 hours.

Future (ITER scale):

- 6,000,000 x 6,000,000 system = 1 hour on 1 Eflop/s system; 1 Pbyte memory.

Supernova Physics

Supernova shock wave instability from 3D core collapse supernova simulation.

[K.L. Ma (viz) and J.M. Blondin]

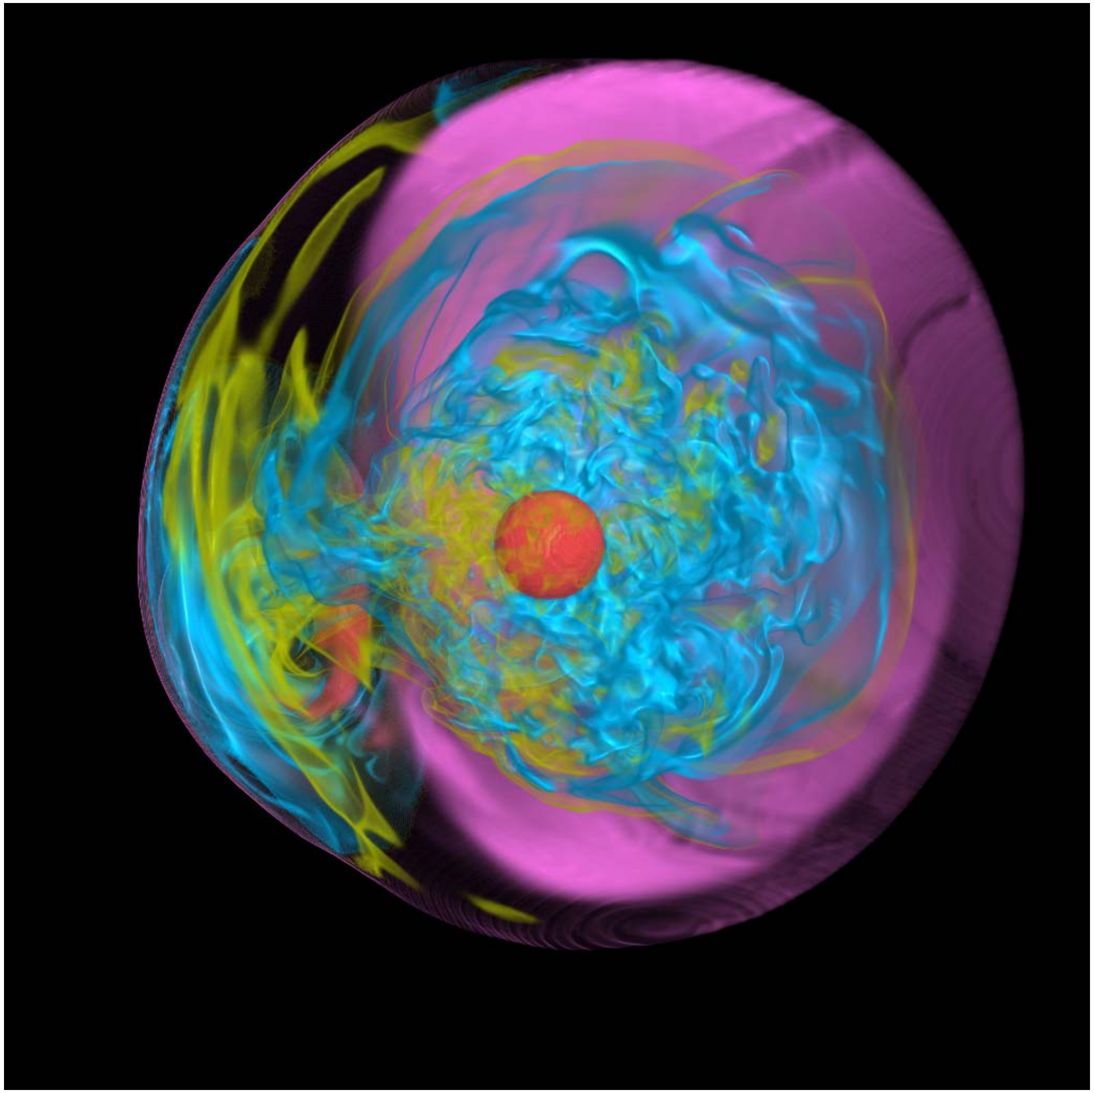

Astrophysics Code Characteristics

- Undergoing a transformation from a data-starved discipline to a data-swamped discipline.
- Large effort in experimental data analysis (microwave background and supernova)
- Typical simulation applications: supernova hydrodynamics, energy transport, black hole simulations.
- Multi-physics and multi-scale phenomena.
- Large dynamic range in time and length.
- Requires adaptive mesh refinement.
- Dense linear algebra.
- FFTs and spherical harmonic transforms.
- Operator splitting methods.

Astrophysics Requirements

Supernova simulation:

- 3-D model calculations will require 1M processor-hours per run, on 1 Pflop/s system, or 1000 hours per run on 1 Eflop/s system.

Analysis of cosmic microwave background data:

WMAP (now)	3×10^{21} flops, 16 Tbyte mem
PLANCK (2007)	2×10^{24} flops, 1.6 Pbyte mem
CMBpol (2015)	1×10^{27} flops, 1 Ebyte mem

Note: Microwave background data analysis, and also supernova data analysis, involves mountains of experimental data, not simulation data.

Decision Making Flow

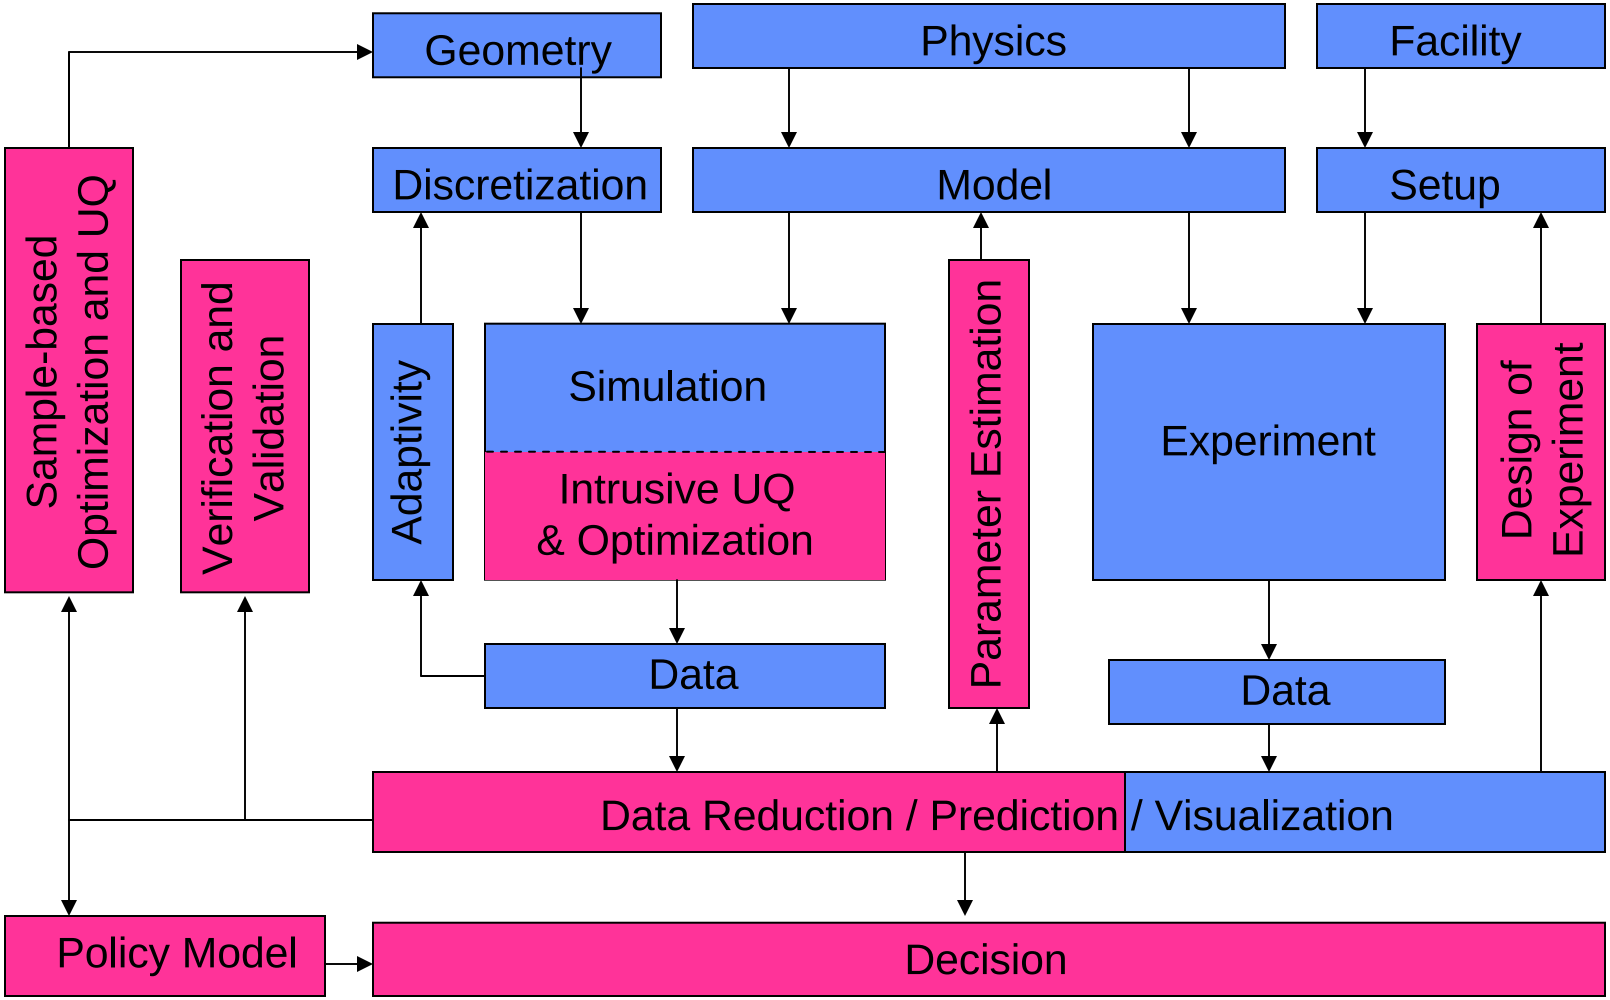

Decision Making

- Any code can be put into a decision making framework to ask
  - Do I believe the model/code? (i.e., V&V)
  - How good? What if? How can I make it better? How can I achieve this through policy? (i.e., design)
- Needs
  - Optimization and uncertainty quantification (includes model building, e.g., setting knobs).
    - sampling-based (need 100x or more computation, capacity computing okay).
    - Intrusive (need 10x computation, need same capability as the application).
  - Policy modeling
    - Looks like discrete event simulation.
    - Based on communication, not computation (need low latency communication and global synchronizations).
  - Data reduction / Prediction / Coupling to experiment
    - Large data sets – need large-scale storage.
    - Streaming data – need bandwidth and multiple threads.

Nano/Material Science, Chemistry

- Objective:
  - *Ab initio* million-atom electronic structure simulations.
- Impact:
  - Atomistic simulation and design of nanoscale-, molecular-, quantum-information processing devices 10^6 atoms, 1-10 ns timescales.
- Resources (time/mem/disk):
  - Local DFT: N-N^2-N^3 / N-N^2 / N-N^2
    - Current state of the art: $O(10^3)$ atom static, O(50) atoms dynamics.
    - QBOX: 207 Tflop/s.
    - Strategies to avoid orbital basis exist, still require science+algorithms.
  - Hybrid DFT: N-N^4 time (e.g., LinK)

- Driver for computer architecture:
  - FFT and associated communication bottleneck.
    - Primarily an issue in planewave methods, however, utilized in some O(N) localized orbital methods (but currently small prefactor).
  - Both dense AND sparse linear algebra problems.
  - Global arrays successful in NWChem.

Adaptivity

- Adaptivity will be required even in zettaflops regime
  - Some problems may not require more resolution beyond a certain point, but adaptivity will improve overall efficiency and time-to-solution.
  - Other problems will benefit from localized higher resolution to capture physics, chemistry, other characteristics of the problem.
  - Minimal floating-point ops, mostly conditionals and branching.
  - Merely a tool to increase efficiency and not a goal in itself.
- Adaptivity of two types:
  - hpr-refinement (structured, unstructured, Cartesian grids, overset grids).
  - Domain specific (e.g. add appropriate models at certain resolutions).
- Adaptivity can be done in parallel (has sufficient concurrency), but:
  - Will need synchronization at processor / domain boundaries unless requirements relaxed from mesh consistency and numerical accuracy.
  - Mesh and solution quality issues become more important in zettaflops regime because # of refinement levels expected to increase exponentially.

Adaptivity, Cont.

- Multi-objective dynamic load balancing needed for adaptivity
  - Balance computations (multi-disciplinary, multi-science, multi-physics).
  - Balance communications (inter-processor, processor-to-memory).
  - Will have to do load balancing at a higher level (since #DOF >> #processing units).
- Solution quality
  - Estimating error in the solution generally more expensive than generating solution.
  - Must have built-in sensitivity / uncertainty analyses to generate confidence in numerical solutions.
  - Holy grail is obtaining grid-independent solutions.

Exascale Performance Requirements

A balanced exaflops machine needs:

- Main memory: ~400 PB
- Working storage: 4 to 40 EB
- Archival storage: 40 to 1000 EB
- Bisection Bandwidth: 0.5 to 1 EB/s
    - Climate pushes local bandwidth
    - FFT pushes global bandwidth
- Balanced RandomAccess and HPL performance

Other Exascale Requirements

- Support legacy applications:
  - $2.5B (just DOE & NSF) MPI software base.
  - 20 year lead time for new programming models.
- Develop new petascale programming environments with exascale in mind:
  - A mere 100,000 threads is difficult too!
  - New algorithms opportunity to use new programming models.
- Fault tolerance (real system throughput):
  - Error checking and notification.
  - System level checkpoint/restart.
- Chip & network performance monitoring

Other Issues

- Extensions to Von Neumann model?
  - Better support for adaptive codes
  - Rick Stevens applications
- Seamless integration with distributed data-rich computing environment.