Enormous Computational Demands in Flow and Quantum Physics

Steve Jones and Masoud Aryanpour

Sept 17 2007

Center for Integrated Turbulence Simulations
Flow Physics and Computational Engineering
Stanford University

Computational Infrastructure

WCR Cluster

- Dell Frontend and Compute nodes
  - 212 PE 1950 (Clovertown)

- Clustercorp Rocks Operating System
  - Cisco-OFED, Panasas, PGI Rolls

- Cisco DDR Infiniband
  - 288 Port Switch
  - OFED 1.2
  - Open MPI and OSU MVAPICH

- Panasas Disk
  - 2 10TB Shelfs

- APC Hot-Aisle Containment System

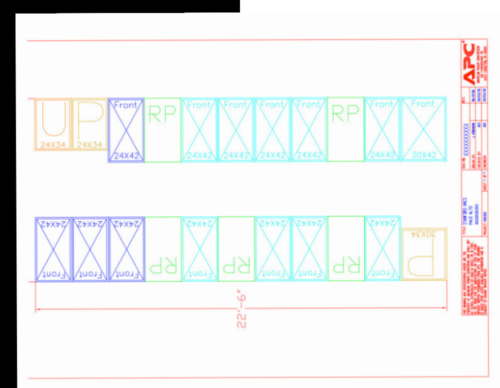

Computational Infrastructure

WCR Cluster

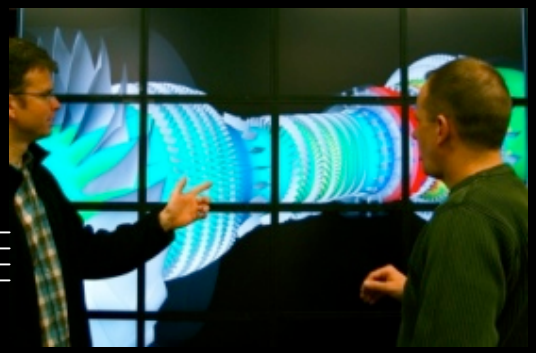

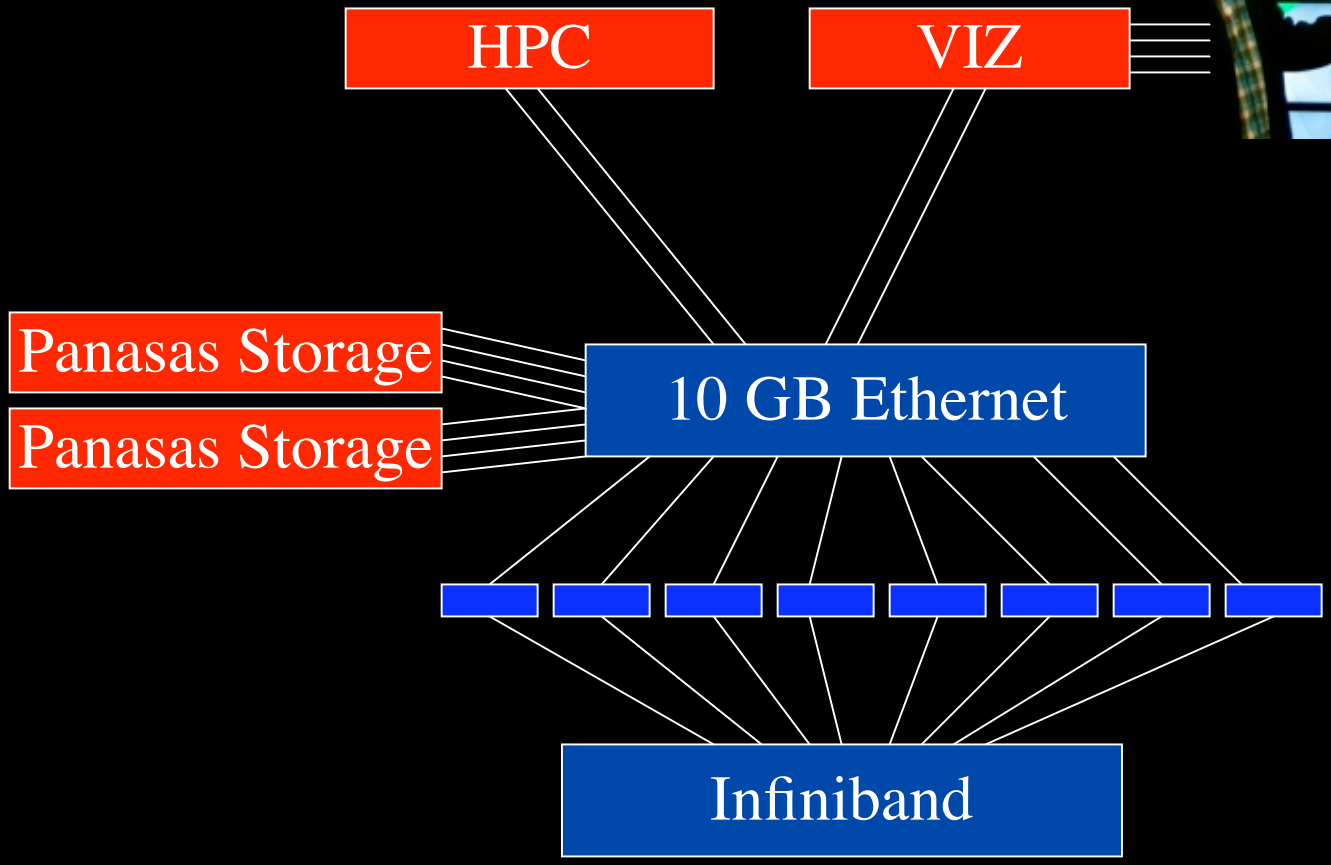

RANS/LES Codes

SUmb (Stanford University multiblock) and CDP (named for the late Charles David Pierce) are both massively parallel flow solvers developed at Stanford under the sponsorship of the Department of Energy's ASC program.

CDP is an unstructured LES code with multi-physics capabilities for computing high-fidelity turbulent reacting multiphase flows. Its low dissipation numerics are critical when important flow structure persists for a relatively long time, such as the trailing vortices generated by helicopter blades.

SUmb is a block-structured RANS code, and is currently being used to study compressible flow related to jet engines, aircraft, helicopters, and space-going vehicles.

Both codes have been run on thousands of processors and can support scalable simulation on large parallel computers such as BlueGene/L, which includes as many as 130,000 processors.

Integrated RANS/LES Computation of a Jet Engine

Simulation of Pratt & Whitney Engine

- 3 sub-domains
  - Compressor (including fan) - SUmb RANS
  - Combustor - LES CDP
  - Turbine - SUmb RANS

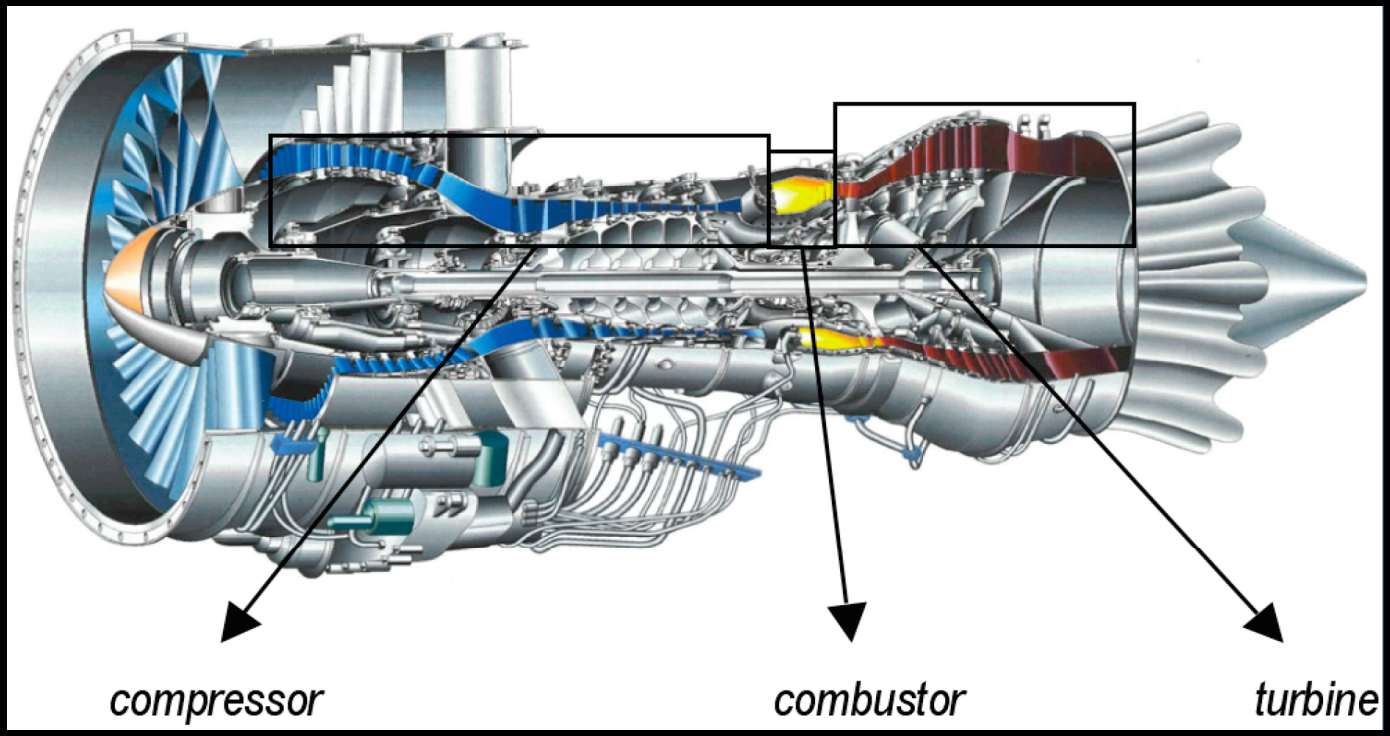

Integrated RANS/LES Computation of a Jet Engine

Coarse Simulation on Current Cluster

- 15 million cell
- 500 CPUs
- 20,000 time steps
- 10 days computing time

Fine Simulation for Red Storm

- 100 million cell
- 3000 CPUs

Simulating Airborne Dispersion of Chemical Agents

General Information

- Model created with Matlab
- Fluent for CFD Analysis (Steady-State)
- 5.2km (1 million cells, fine)

Simulating Airborne Dispersion of Chemical Agents

Simulation Data

- Grid size - 1 million cell (preliminary investigation)
- 15m elevation above street level
- 50m (coarse), 10m (medium), 2m (fine)

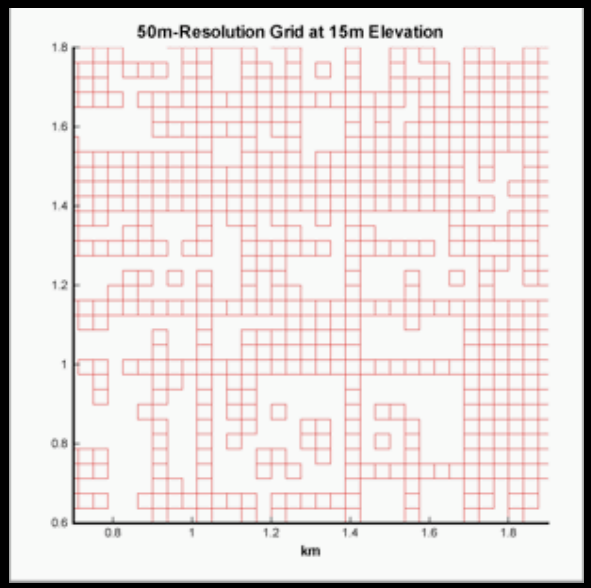

Simulating Airborne Dispersion of Chemical Agents

Computational Demands (preliminary)

- 1km vertical cut-plane
- Steady-State
- 128 Hours/CPU

Real Simulation Demands

- 1 Billion Cells +
- 600 proc x 48 hours = 2 min Simulation
  (30 minutes for useful data)
- TD simulation using CDP

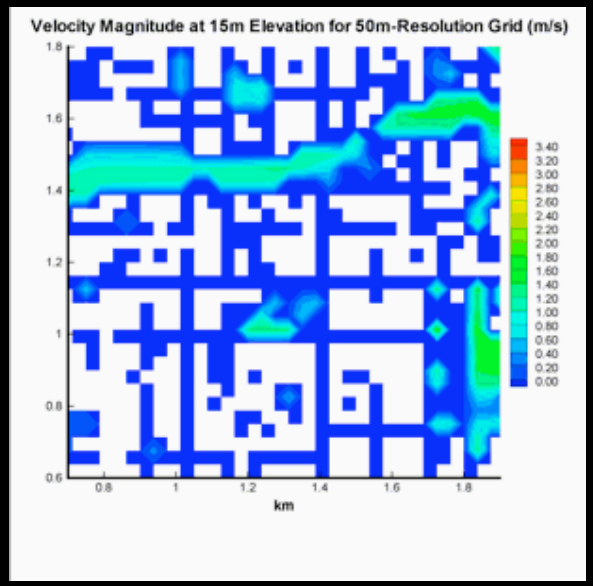

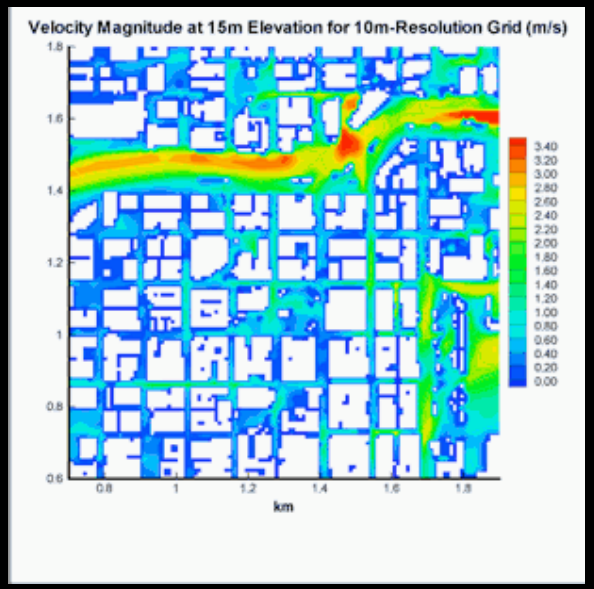

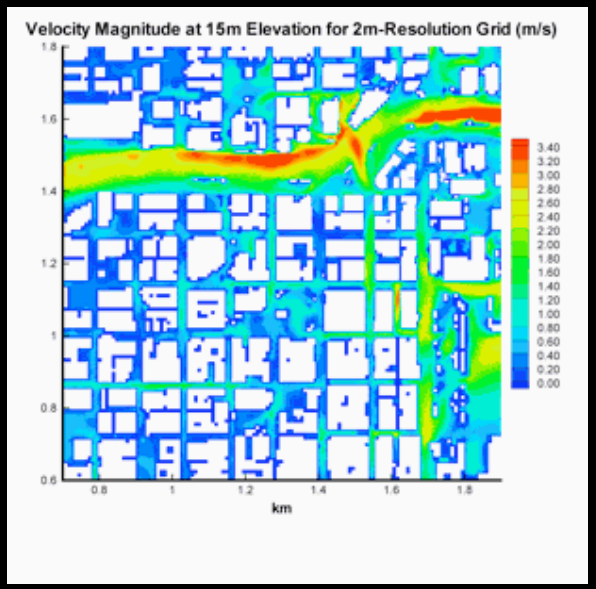

Question: planetary origin

- Nearly all models of planetary formation assume turbulent mixing and transport within the disk
- The goal is to simulate a turbulent disk at a high Reynolds number and observe the corresponding energy growth (Re_λ~200)

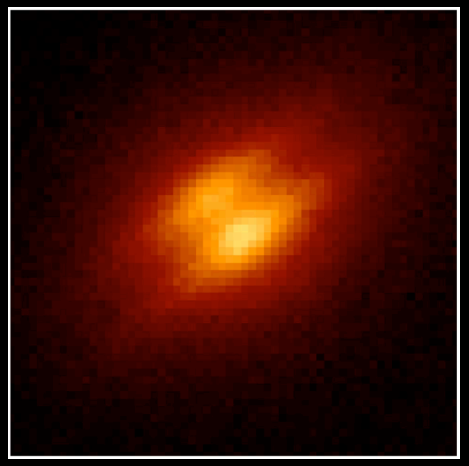

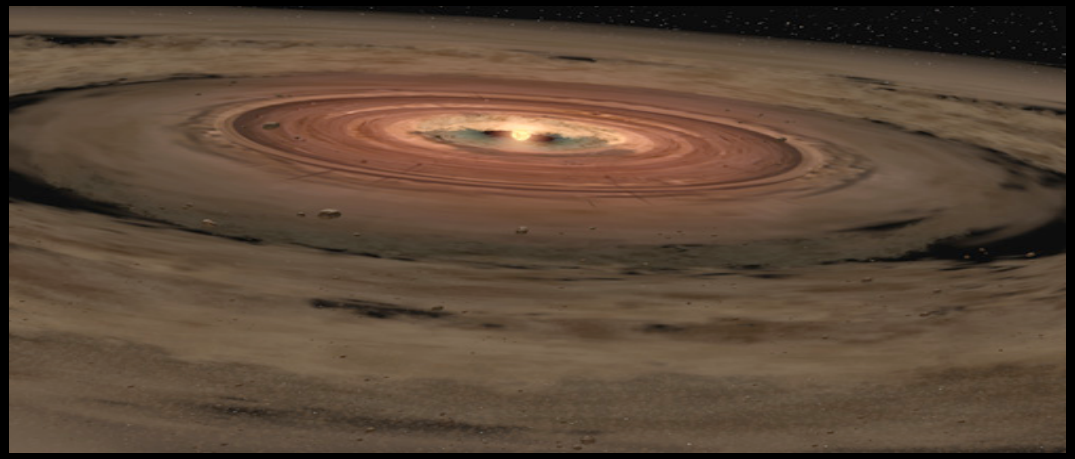

Computations

Simulations

- Grid size: 8192x2048x2048 ~ 34 billion grid points
- Ran on 65536 CPUs (half of the machine)
- 7 days of computing time
- FFT transforms require large data (full volume) communication multiple times per time-step

Data processing

- Each saved data field: 384 GB
- 14 data sets
- Inverse FFT on entire dataset must be performed to visualize solution in physical space

A week on BlueGene/L

- The machine remained stable for the entire week and 60% parallel efficiency was achieved (spectral code with multiple volume data communication in each timestep)

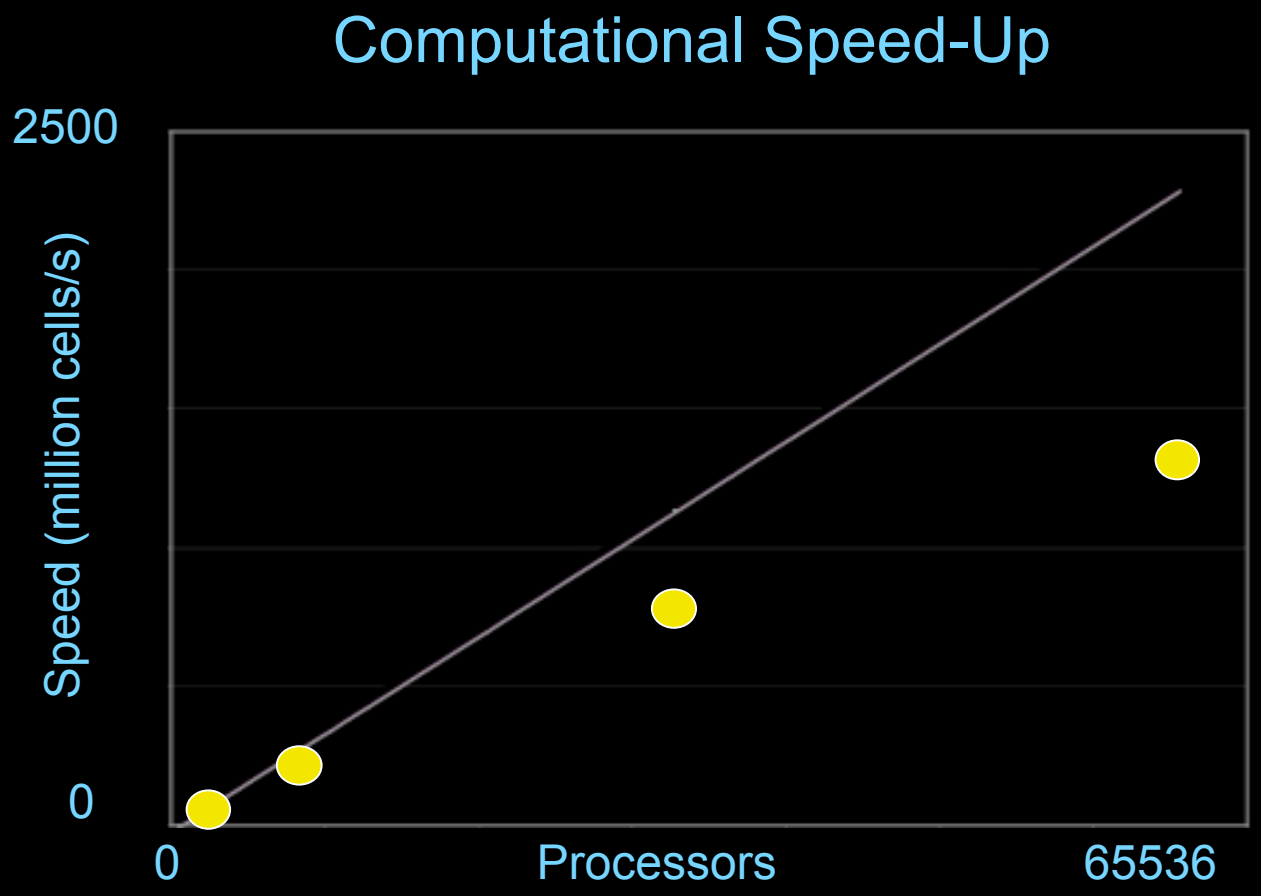

Growth of Ab-Initio Calculations

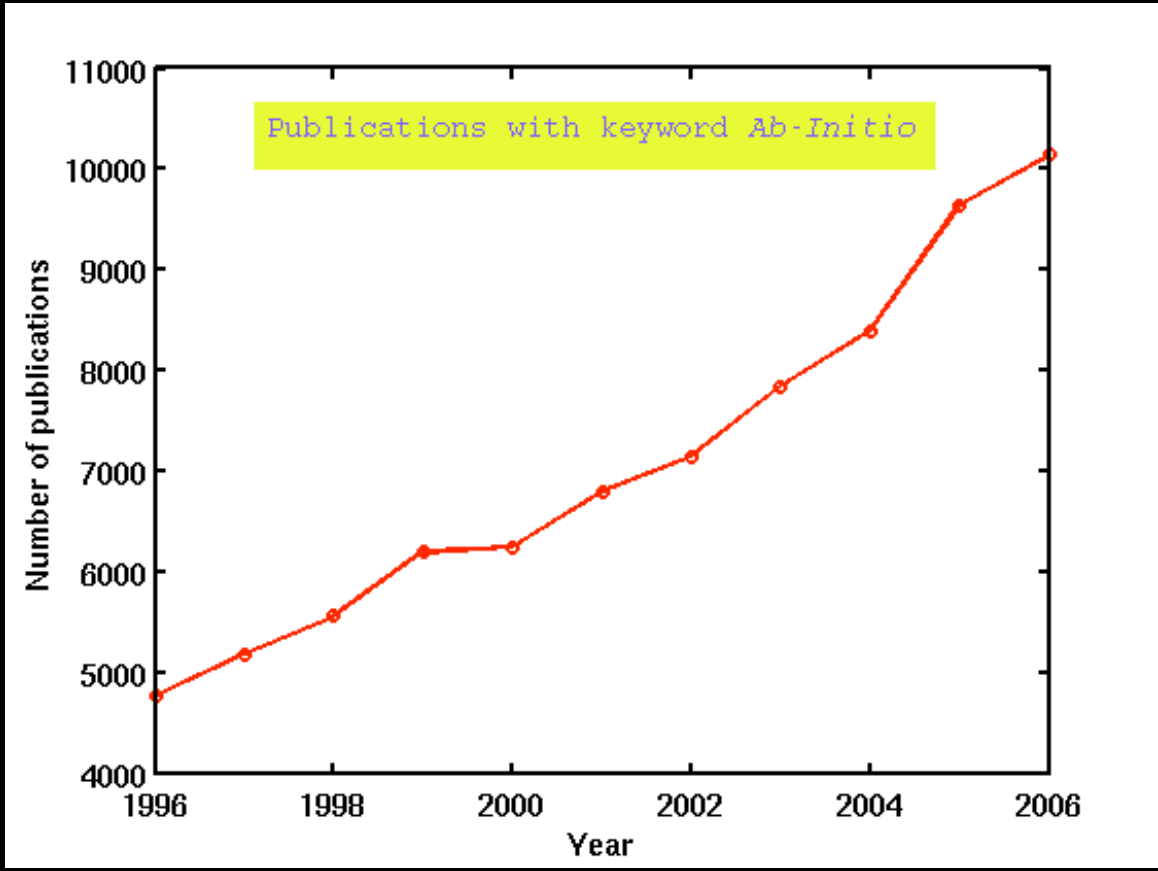

Ab-Initio refers to quantum computations

Power of Quantum Calculations in Science

Quantum calculation of materials

Structure

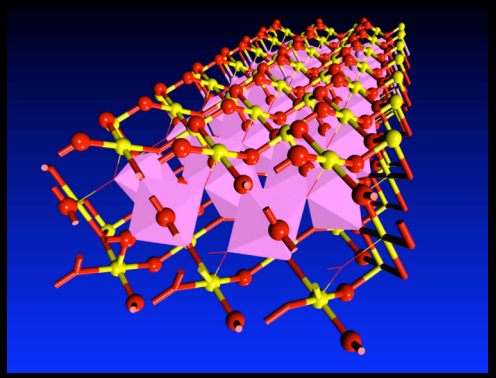

Energetics

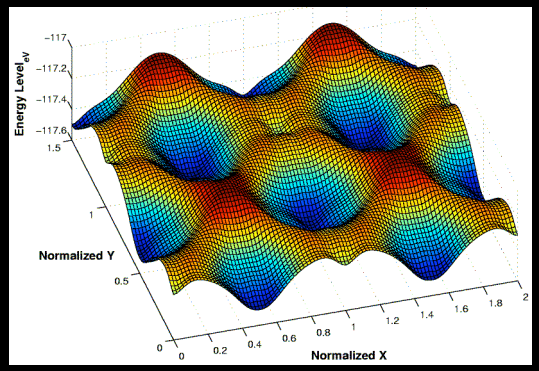

Chemical kinetics

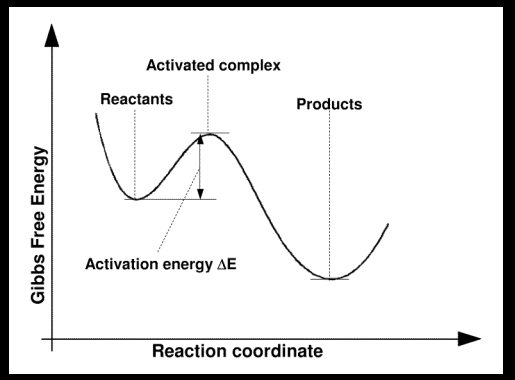

Mechanical properties
- Density
- Elasticity
- Stability

Electrical properties
- Conductivity
- Permittivity
- Electron-transfer

Chemical properties
- Reaction pathways
- Rate constants
- Catalytic activities

Simulation of Electrons & Protons

Sample system: O_2 on $Pt_{10,5}$

- 17 atoms
- Determine electronic structure
- Simulate only valence electrons: 286 out of 1186

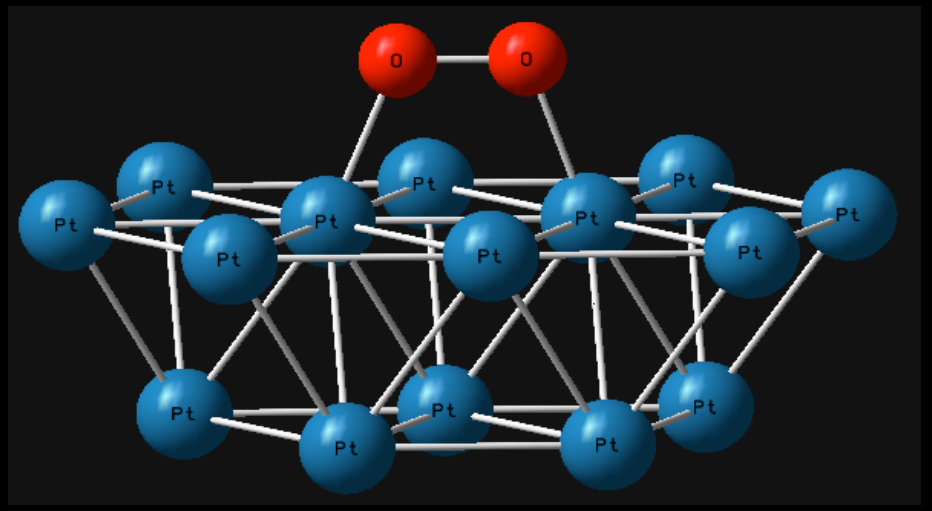

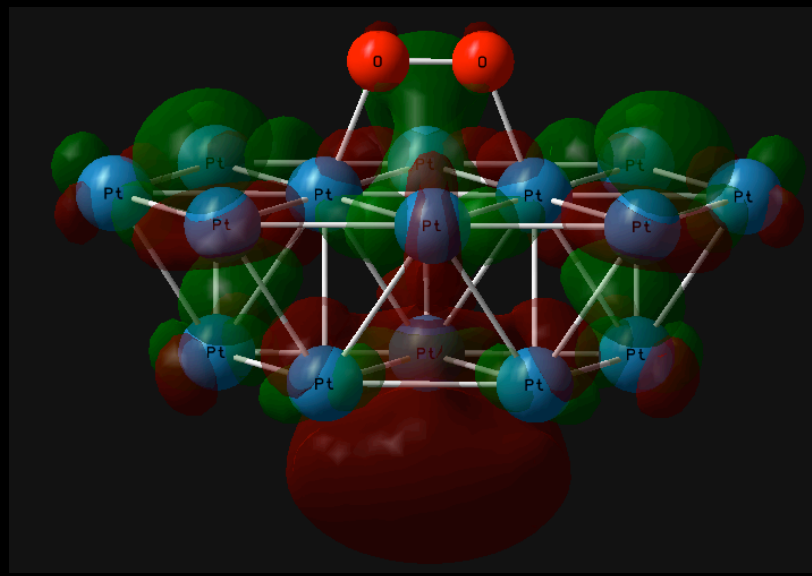

1-orbital, 2-electron structure

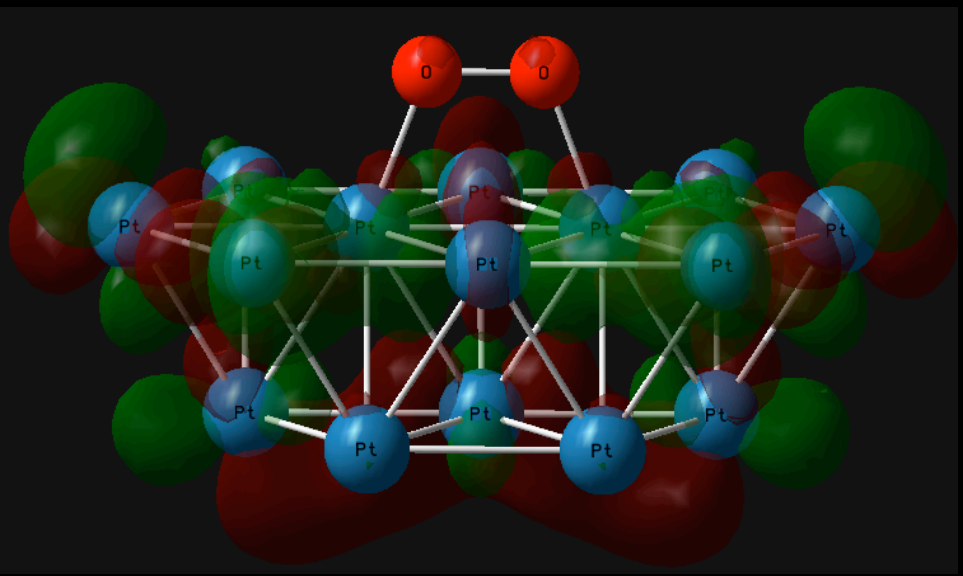

1-orbital, 2-electron structure

Number of Electrons in Elements

Computational approach

- Treating nuclei as points
- On average 10-15 valence electrons per atom
- Inner electrons frozen in computations

Basis Set Requirements

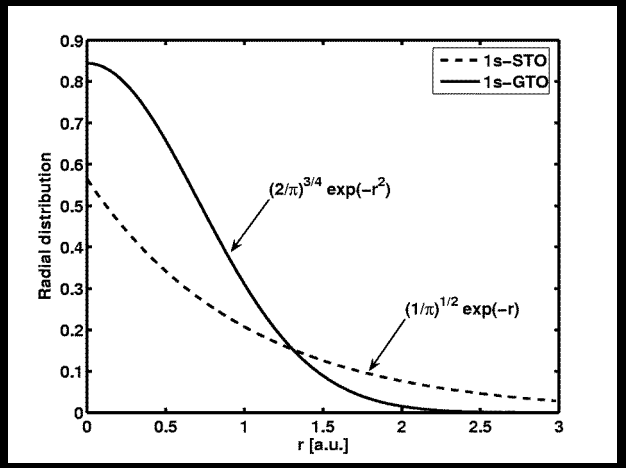

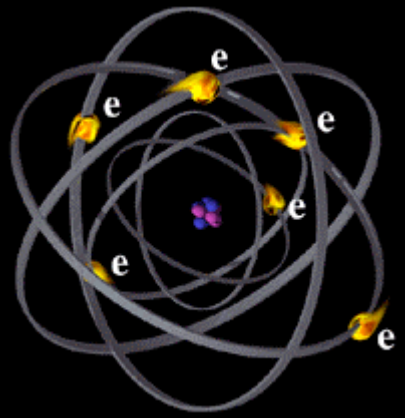

basis functions → one orbital → maximum two electrons

$N_b \sim 3*N_e/2$ ← $N_e/2$ = min. orbitals ← N_e = valence electrons

Computational cost depends on basis set and quantum method

Computational Scaling with Accuracy

Method	Scaling	Error in energy (kcal/mol)
HF	N^2 to N^4	12.4
MP2	N^5	5.5
CISD	N^8	3.9
CISD(T)	N^{10}	3.1

- **GAUSSIAN** simulation only up to 10-30 atoms using orbital methods due to high computational cost
- Maximum 1-4 CPUs used

N^3 Scaling Using Density Functional Theory (DFT)

Basic idea of DFT:

- Replace orbitals with charge density in quantum equation

Benefits & drawbacks

- Simpler equations
- More parallelizable algorithms
- Moderate accuracy
- Accuracy almost independent of basis set and method

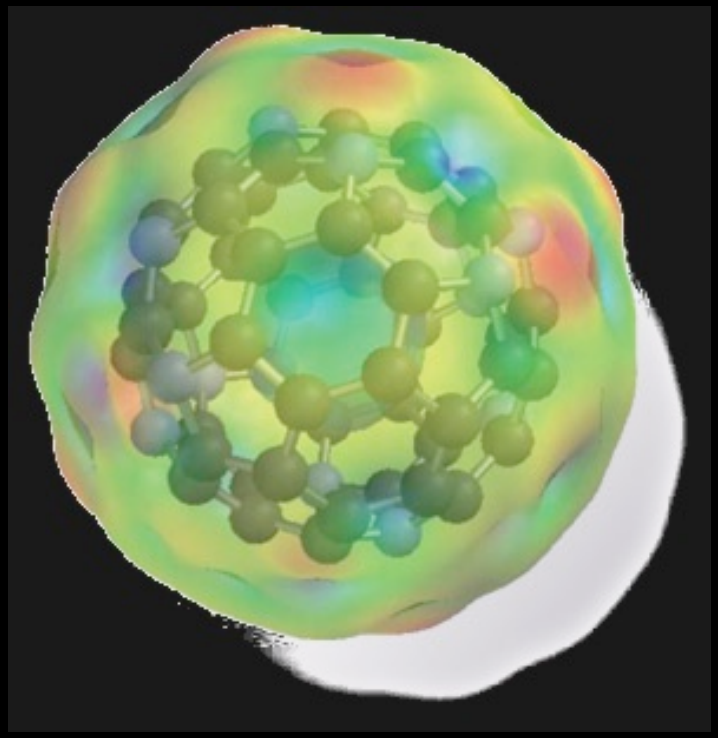

Density Functional methods scale as N^3 using **VASP**

Comparison of Ethernet Switch & Infiniband Network

- Source of data: **VASP** forum
- Case : 8 x TiO_2, 24 atoms
- Cluster: Opteron 275
- Compiler: PGI 6.0.8/MPICH 1.2.6

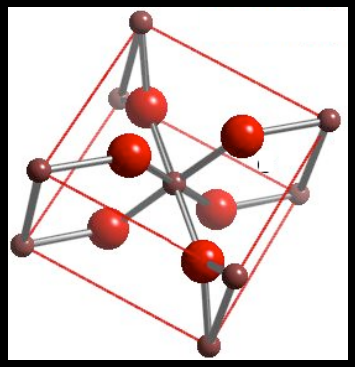

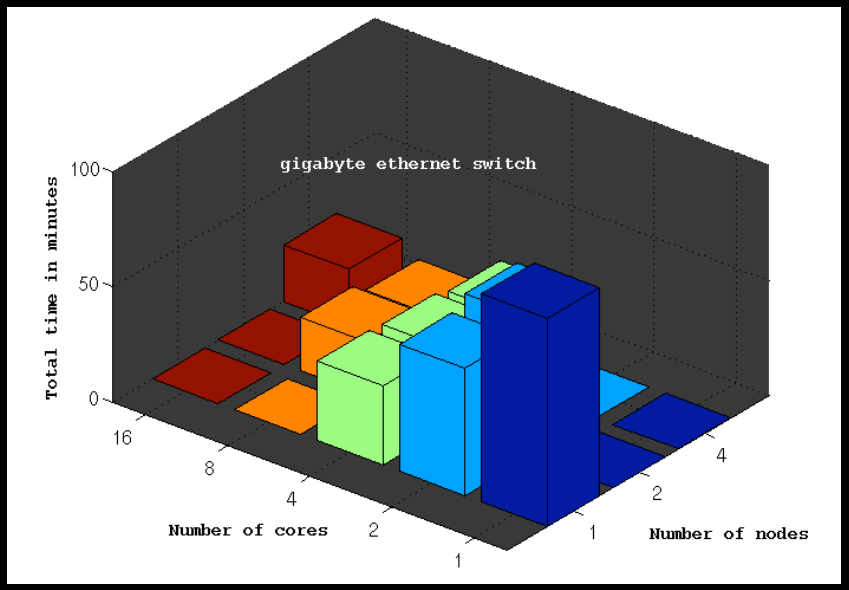

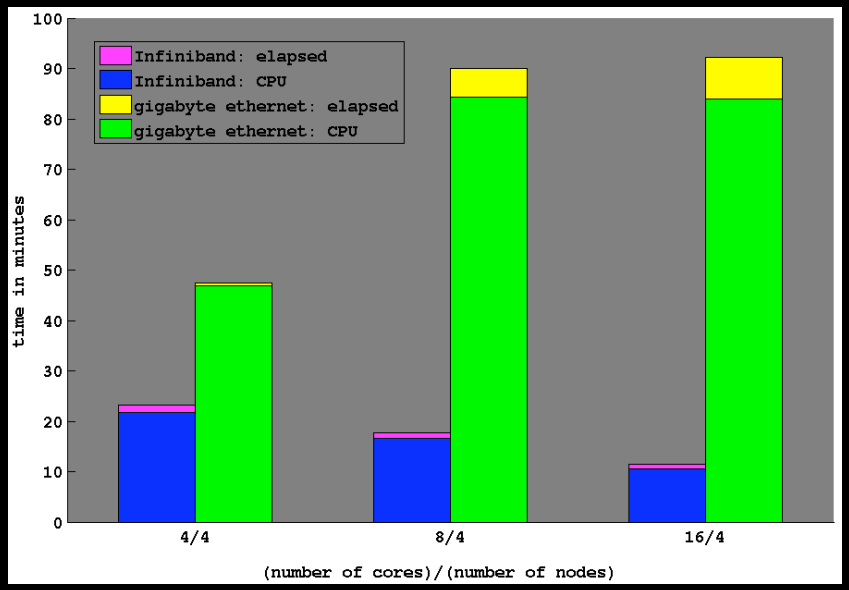

- NFS efficiency crucial for parallel performance of **VASP**
- Faster CPUs require faster interconnections

Benchmarking for **VASP**

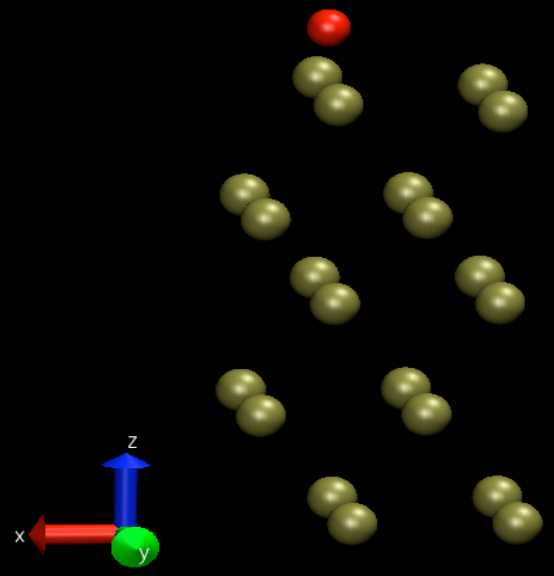

Case : O on Pt_{4x5}

- Periodic Boundary Conditions
- Unit cell: 5.6 x 5.6 x 31.8 Angstrom
- Grid points: 42 x 42 x 240
- Number of Plane Waves: 423360
- Cluster: WCR at Stanford

- Need for efficient MPI process cleanup
- Need for standardization of job launchers

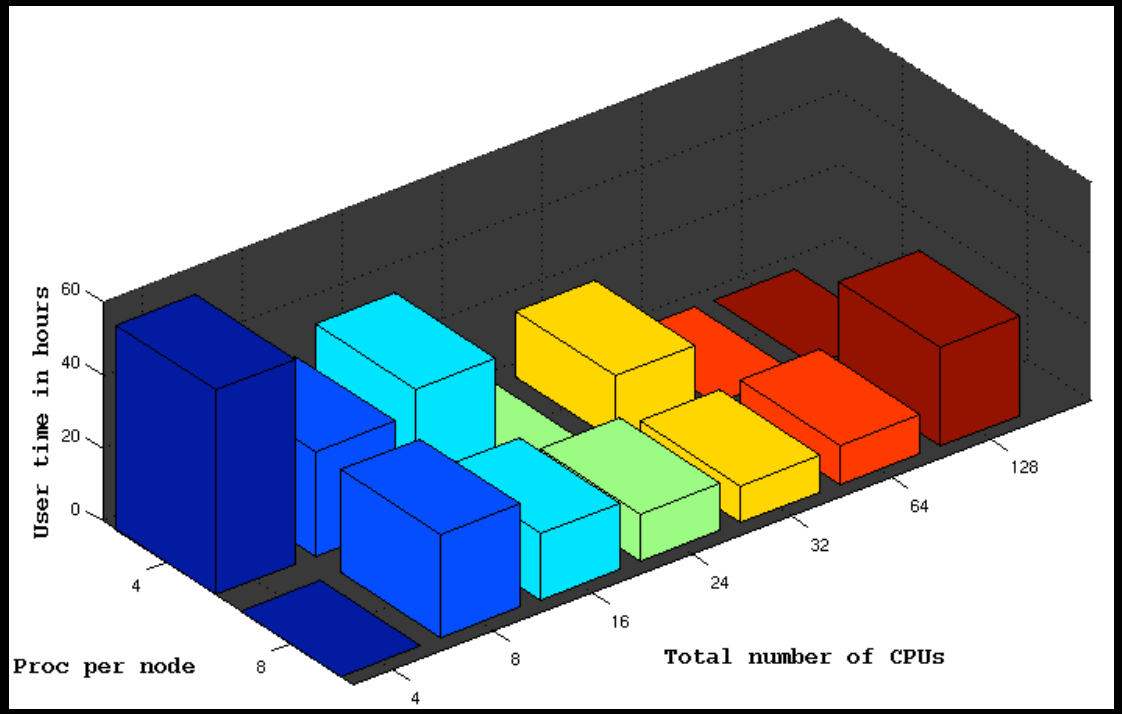

From Current Quantum Computations to Real Systems

	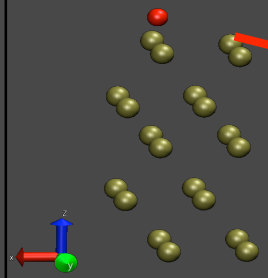	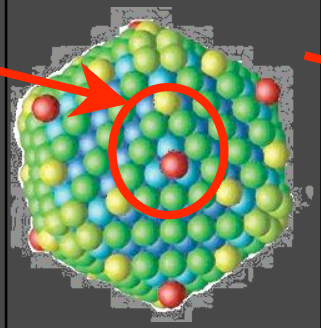	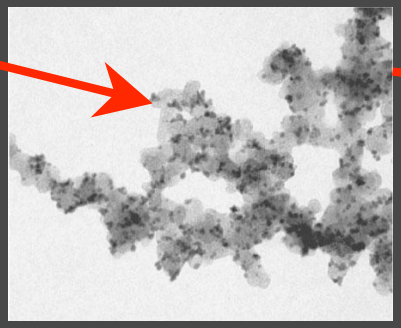	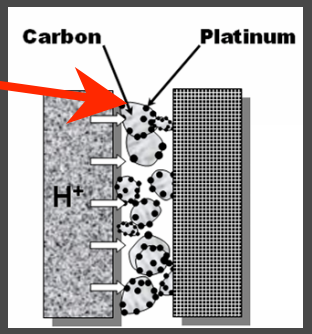
system	**Comput. cell**	**Pt particle**	**Carbon particles**	**Catalyst layer**
length	***6 A***	**3 nm**	**Each ~30 nm**	**~ 1mm**
atoms	**20 Pt**	**~1000 Pt**	**Each ~10000 C**	**?**

- Billions of atoms present in real systems
- Huge gap between computational models and real systems
- Much more realistic computations required to improve predictions

Impact of Fuel Cells Improvement on Energy Industry

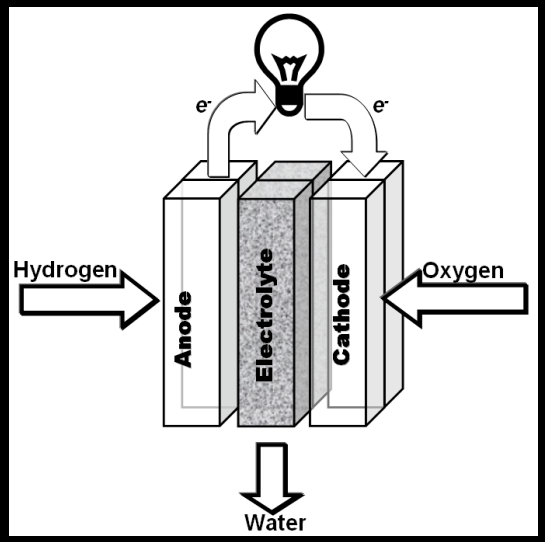

Cathode

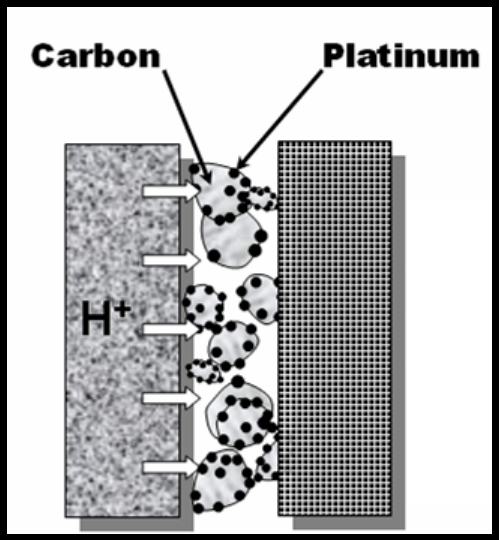

Commercial fuel cells

- much higher efficiencies
- more expensive

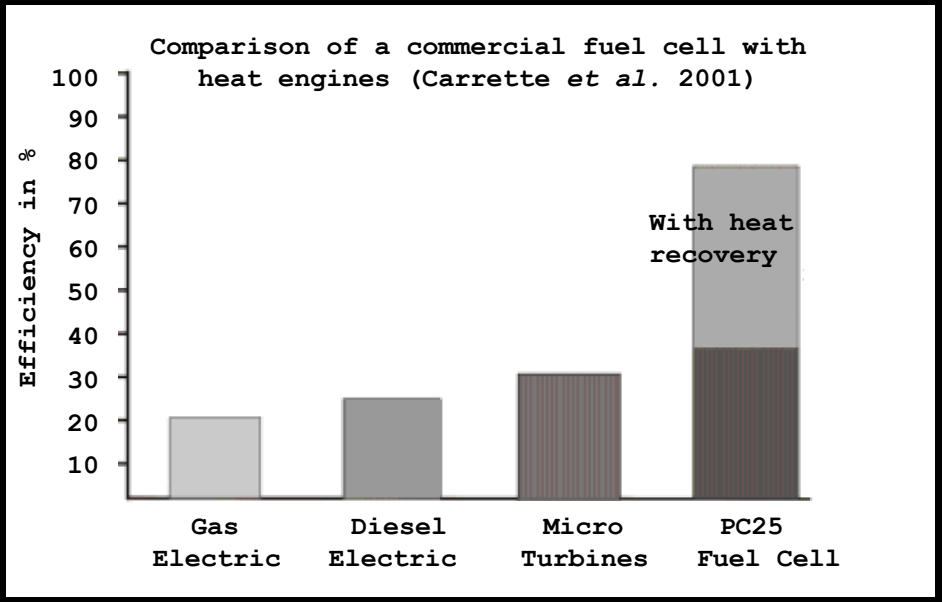

Conclusions

Computational Demands are not getting smaller

- Processor count per simulation - 10k +
- Grid size - 1 Billion +
- Size of dataset - 10TB +

What are the IB Software and Hardware Vendors doing to support this?

ISVs are not in line with HW Vendors

- Example - Versions of MPI or Driver Support differ greatly

Better Tools for Managing Fabrics

- Network Analysis and Monitoring (that actually does something useful)

Contact Information

Steve Jones stevejones@stanford.edu
Masoud Aryanpour aryanpour@stanford.edu