



High Performance Computing and Networking

Bill Feiereisen

wfeiereisen@mail.arc.nasa.gov



Topics

- Relationship to Code Y needs
- Trends in high performance computing
 - state of the art and where it's going
- Two activities by Ames and its partners
- Proposed Code Y collaborations
 - near term
 - longer term research agenda.

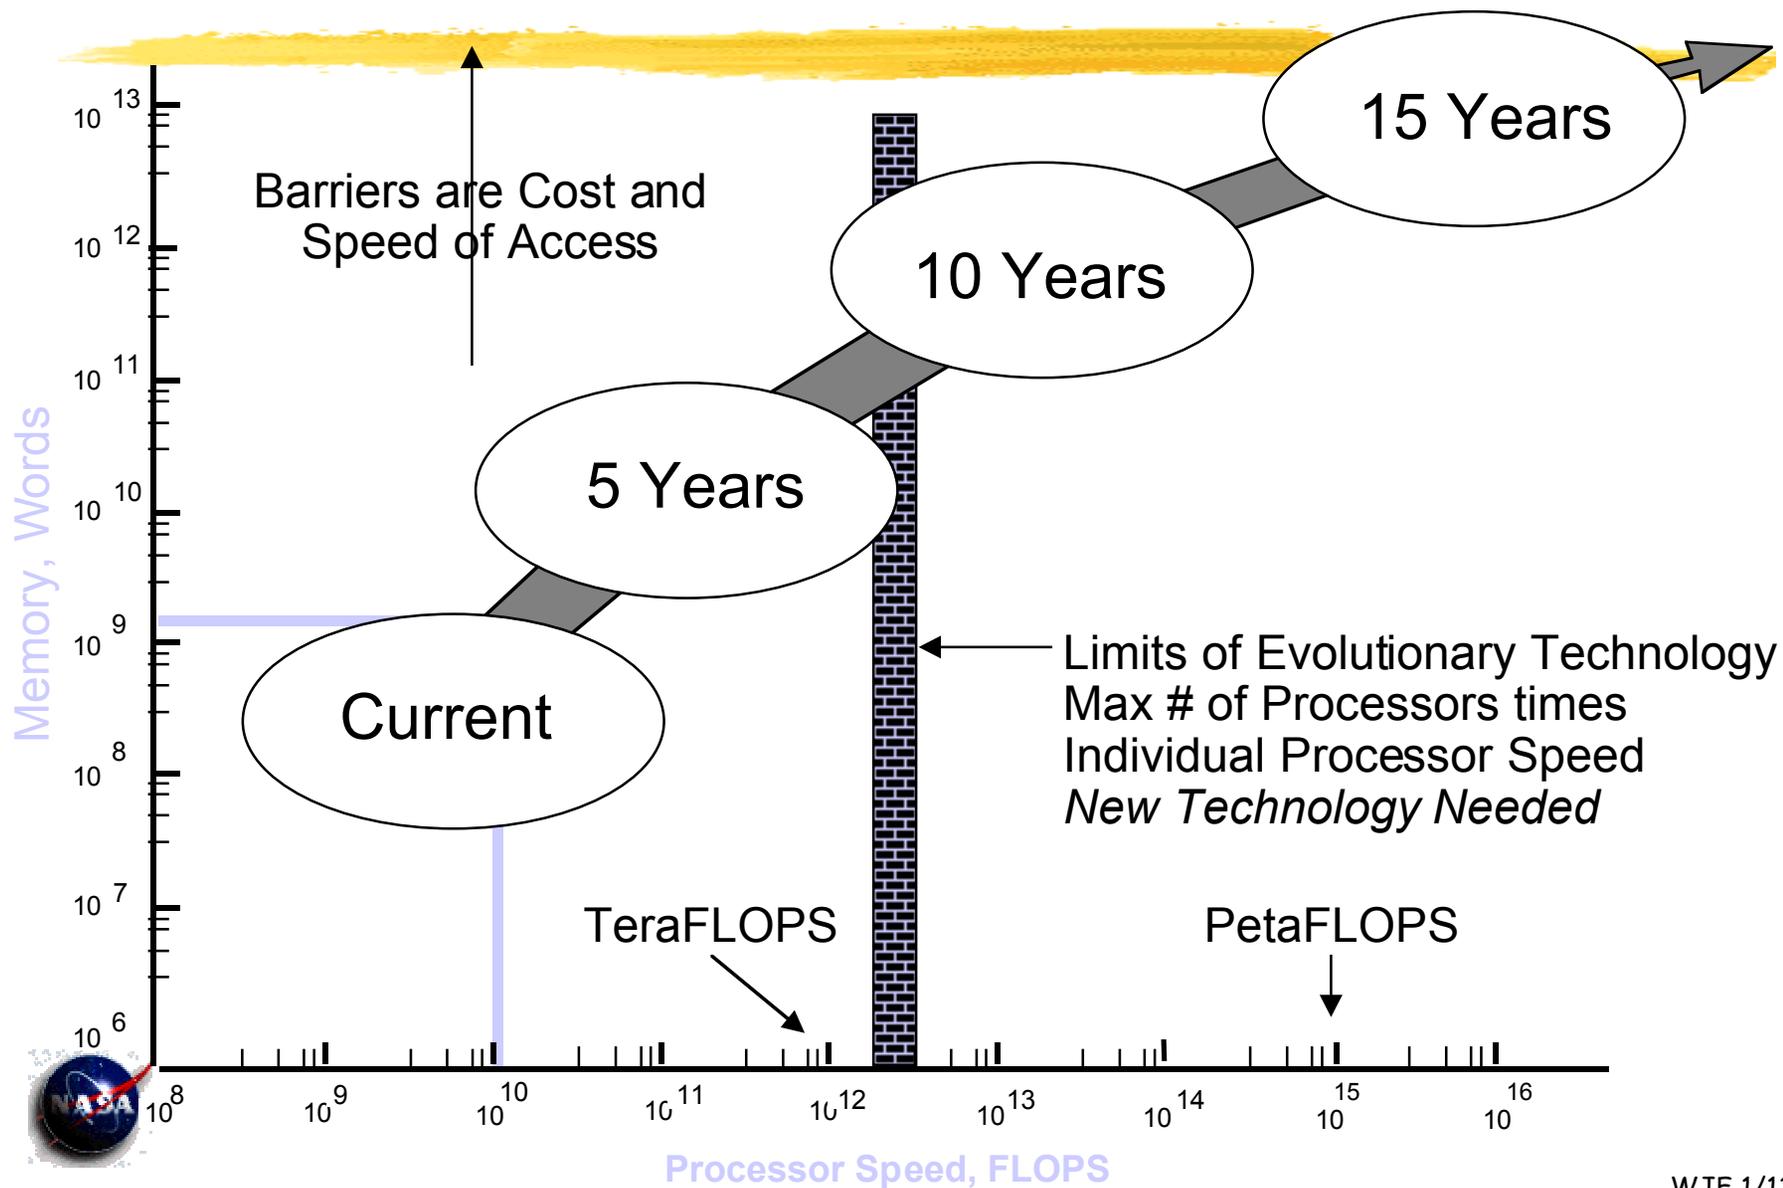


Technical Areas

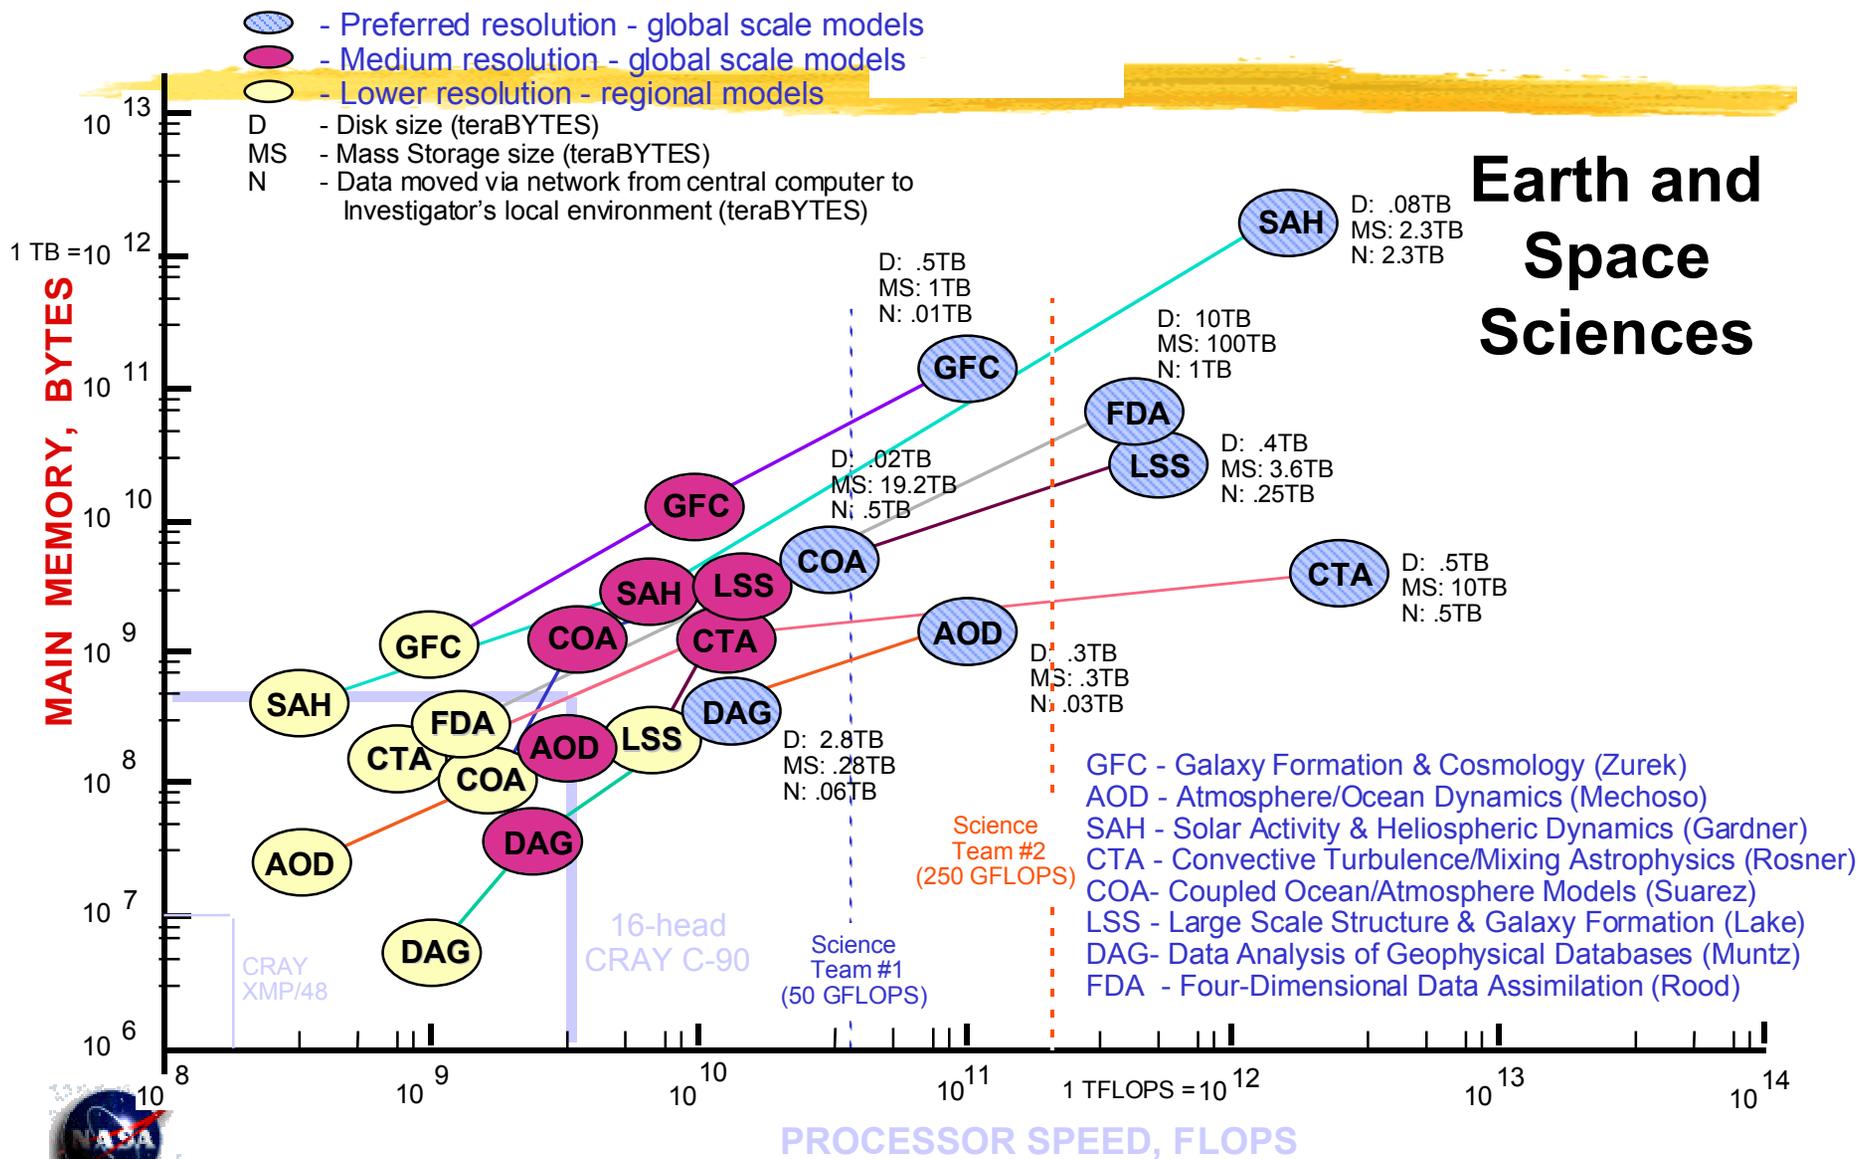
- Current and Emerging Computer Architectures
- Legacy Code Modernization, Applications expertise in the use of modern high performance computing architectures



Mission Demands on Computing



Main Memory vs CPU Requirements of ESS Grand Challenge Investigators (Science Team #1)

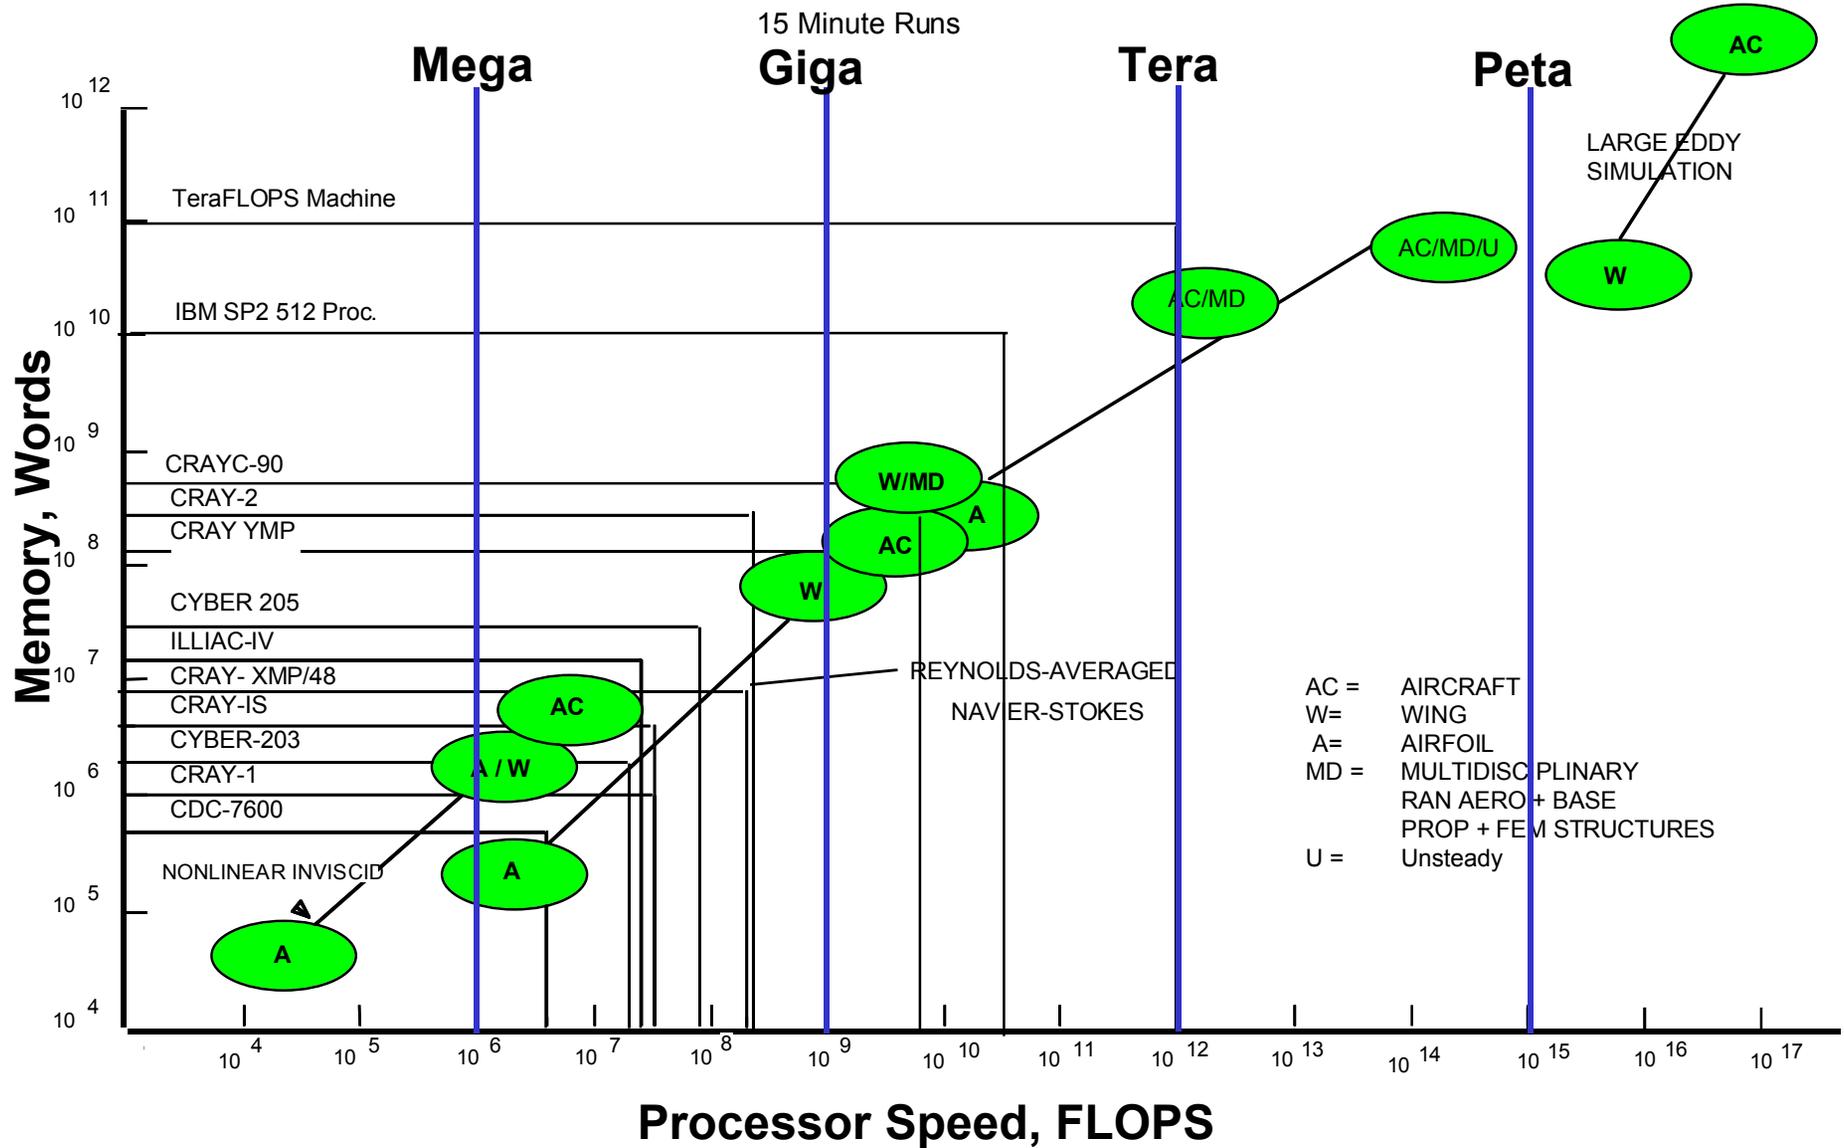


Earth and Space Sciences

- GFC - Galaxy Formation & Cosmology (Zurek)
- AOD - Atmosphere/Ocean Dynamics (Mechoso)
- SAH - Solar Activity & Heliospheric Dynamics (Gardner)
- CTA - Convective Turbulence/Mixing Astrophysics (Rosner)
- COA - Coupled Ocean/Atmosphere Models (Suarez)
- LSS - Large Scale Structure & Galaxy Formation (Lake)
- DAG - Data Analysis of Geophysical Databases (Muntz)
- FDA - Four-Dimensional Data Assimilation (Rood)



Aeronautics Modeling and Simulation



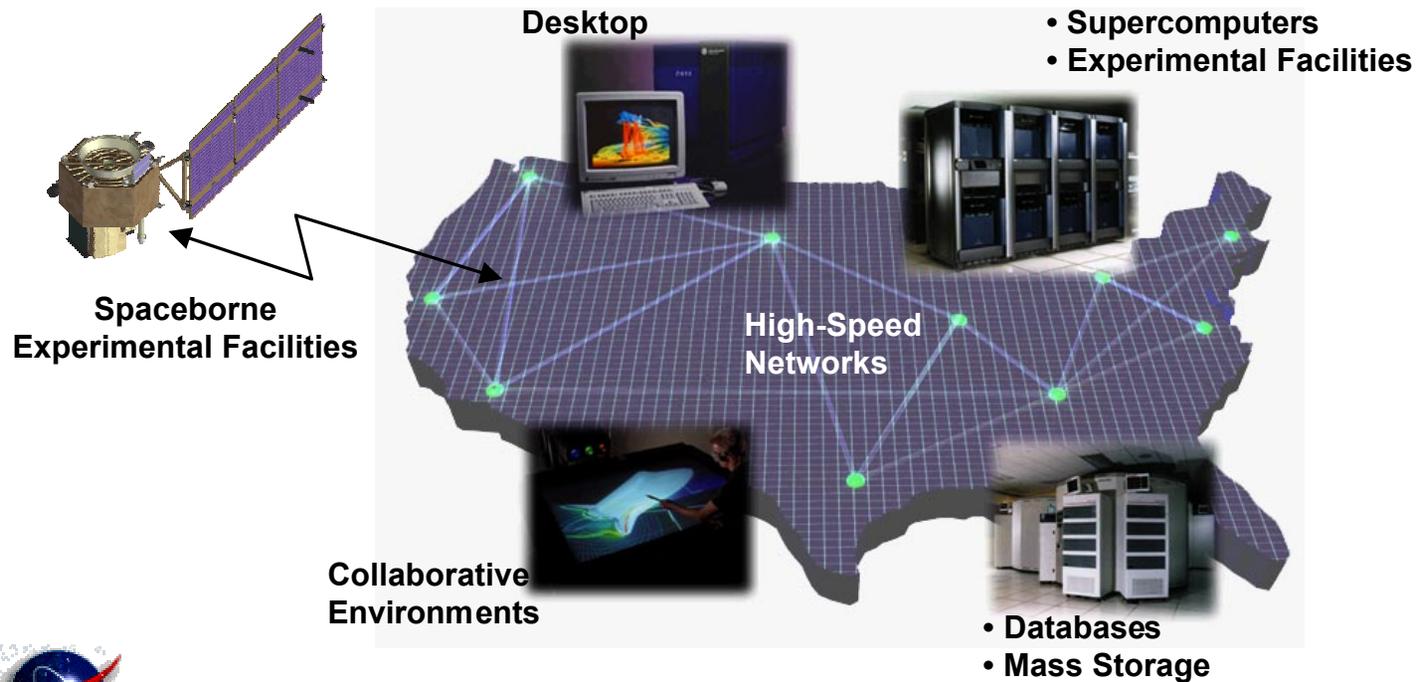
Trends in High Performance Computing and Networking

- A broadening from simulation of physical phenomena to knowledge extraction from information...regardless of source
- A change from special purpose low volume supercomputers to commodity based parallel machines
- The rise of extremely fast network connectivity
- The need to meld geographically dispersed colleagues together



State of the Art

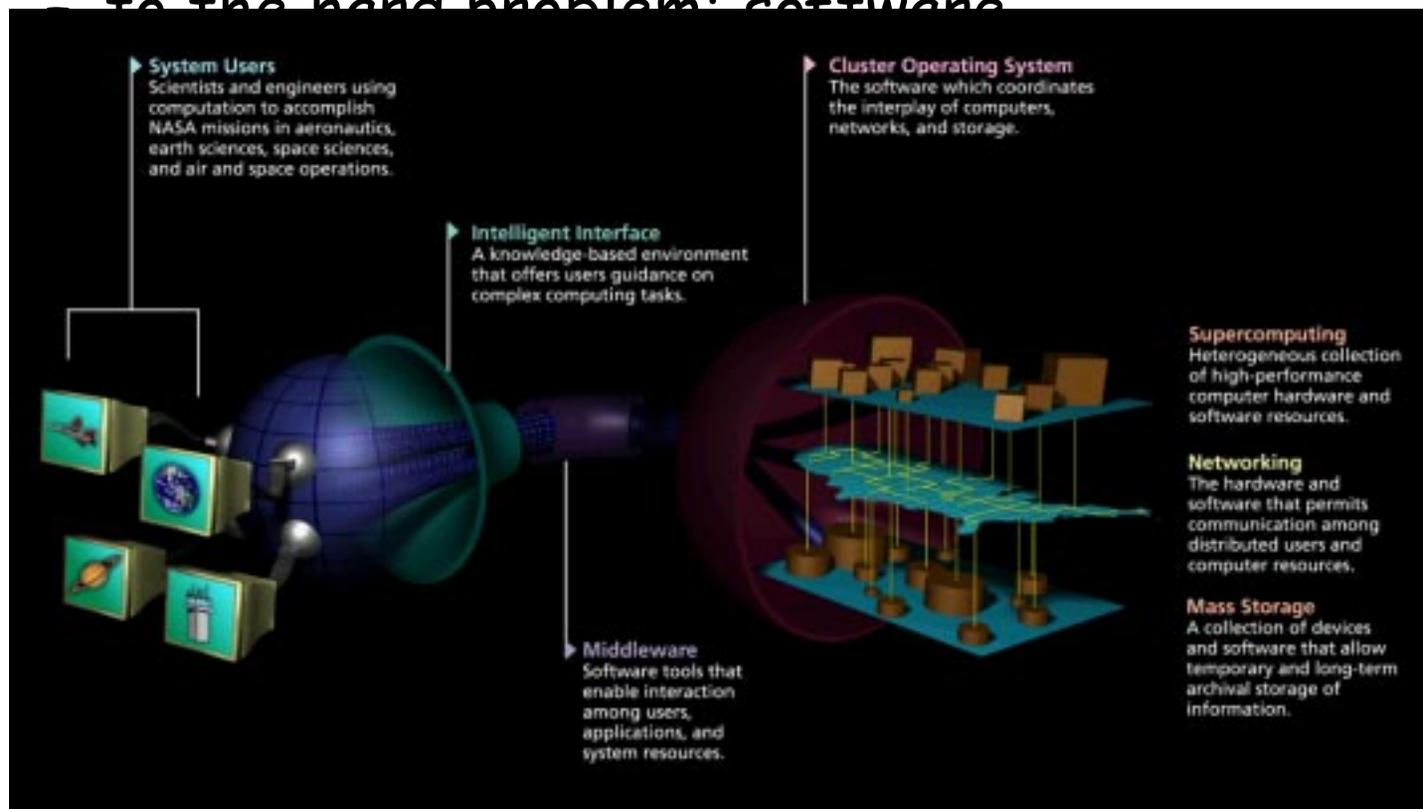
- Current: centrally managed data centers
- Future: distributed federation - A Grid



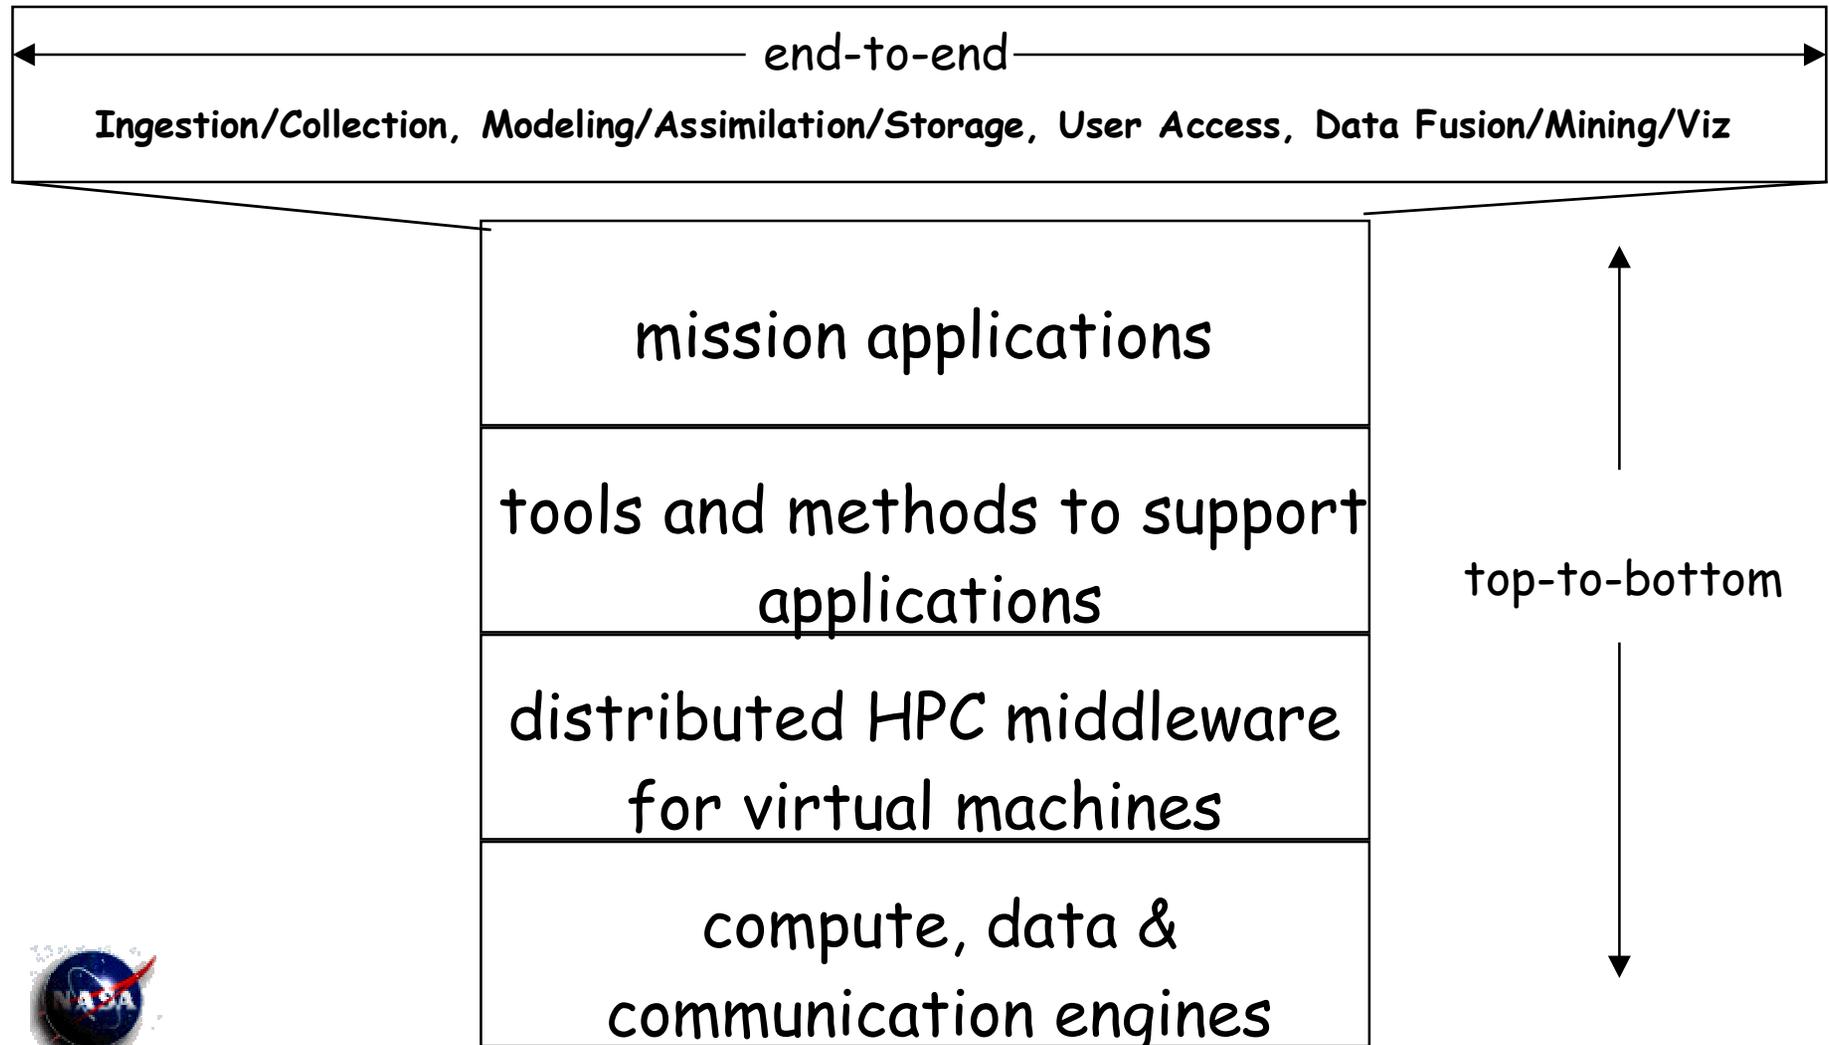
Software is the key

- Reorient investments
 - to trends

to the hard problem: software



HPC and Networking are Layered



HPC Processor Evolution

- Past & current: Special purpose low sales volume production processors
 - vector, multi threaded
- Current & future: High sales volume commodity based processors
 - distributed memory, tightly coupled, loosely coupled
 - shared memory, ccNUMA
- Future:
 - PetaFlops, New device technologies/New



Network and Storage

- Network Connectivity
 - Current: "best effort"
 - Future: "differentiated services"
 - reservable bandwidth, security.....
- Storage and Retrieval of Data
 - Current: centralized solutions that don't interoperate. "deep archive" - data is mostly stored but not retrieved
 - Future: distributed data archives as the basis for knowledge discovery



Information Power Grid

- The layered architecture of the grid

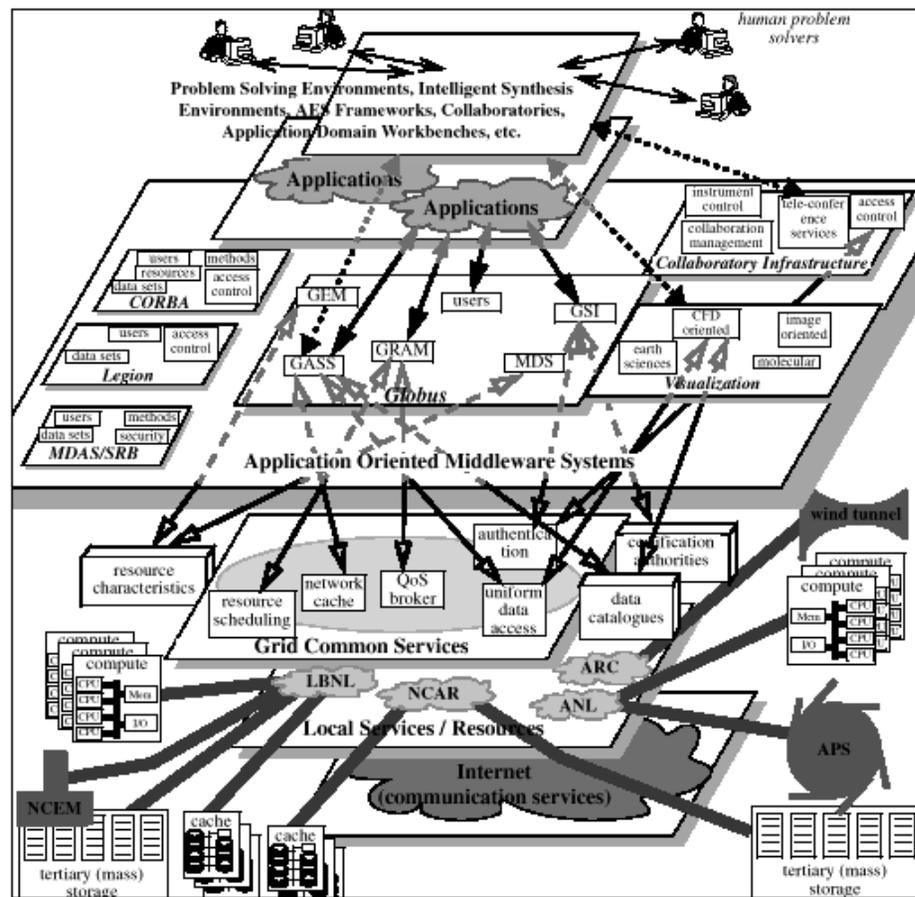
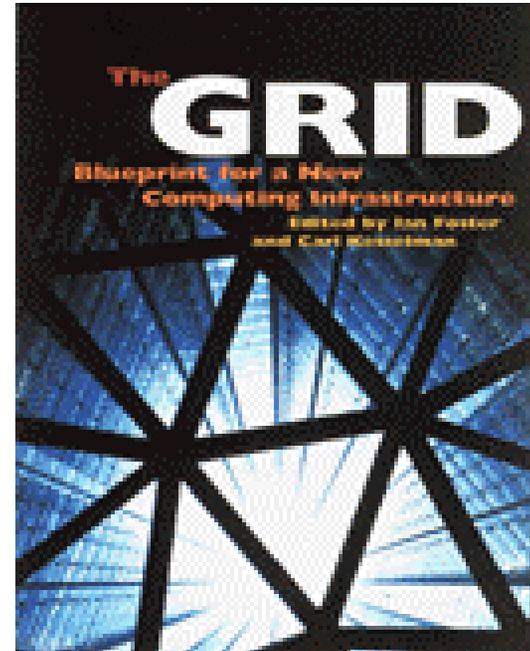


Figure 4 The layered architecture of the Grid.



Who are our partners?

- NASA Research Centers, Ames, Langley, Lewis...**GSFC, JPL & Code Y?**
- The NSF PACI consortia
 - The Alliance
 - NPACI
- DOE
- DOD/MSRC
- Other Academic Partners
- The Aerospace Industry



Key Applications Software Issues

- Software Engineering
- Verification and Validation
- Adaption and Portability to new Architectures
 - experience with the largest shared memory machines
 - distributed memory

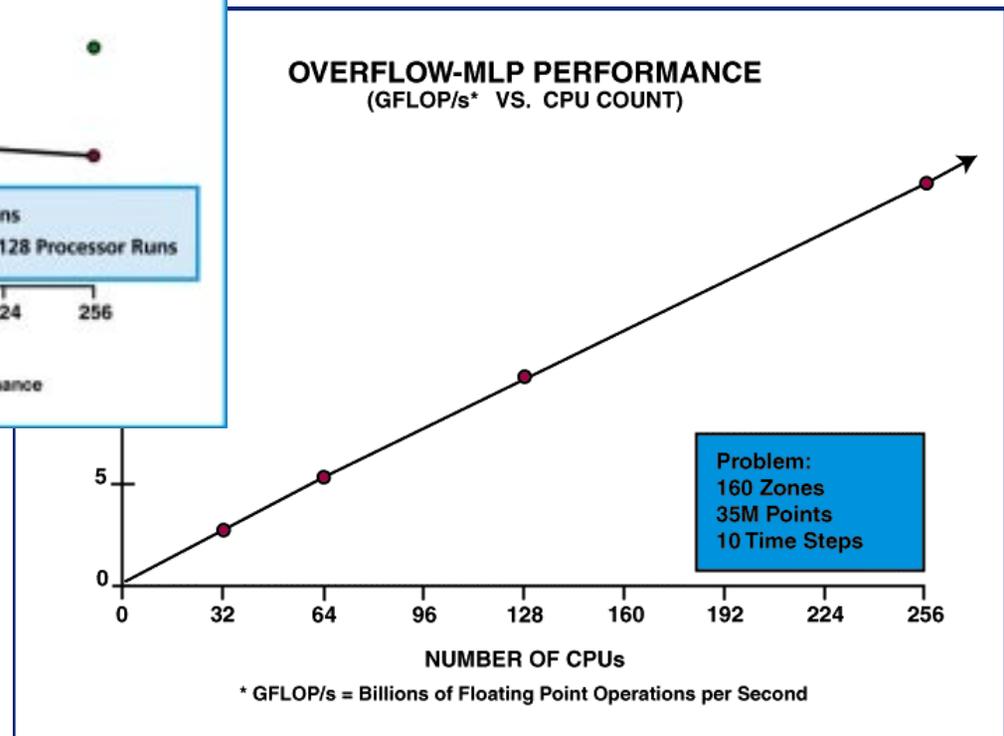
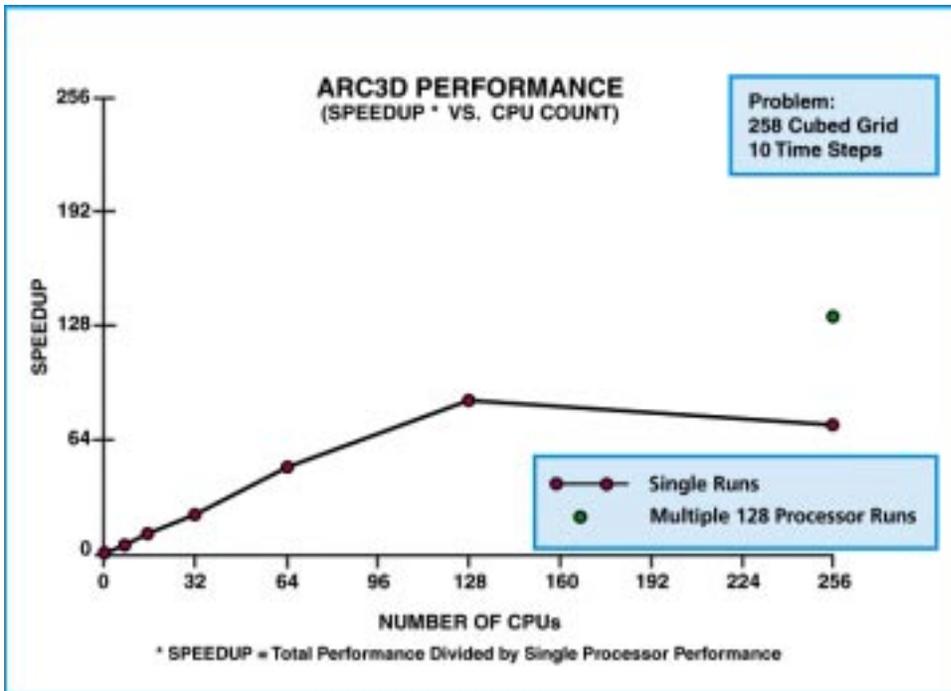


Adapting Applications to new Architectures

- Methods
 - expert knowledge, hand-coding
 - relying on vendor-supplied parallelizing compilers/tools
 - annotating/rewriting application using data- and task-parallel directives/languages
 - rewriting application codes using semi-custom building blocks or libraries
 - using interactive parallelization/tuning environments: Legacy Code Modernization

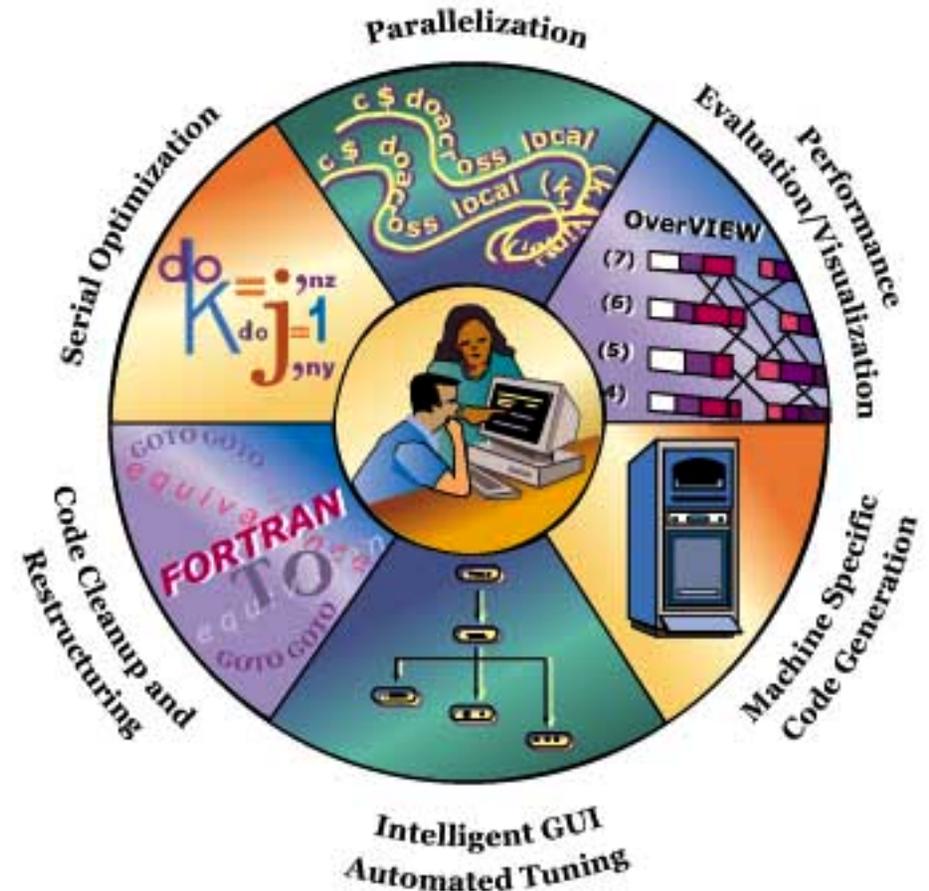


Expert Knowledge Captured

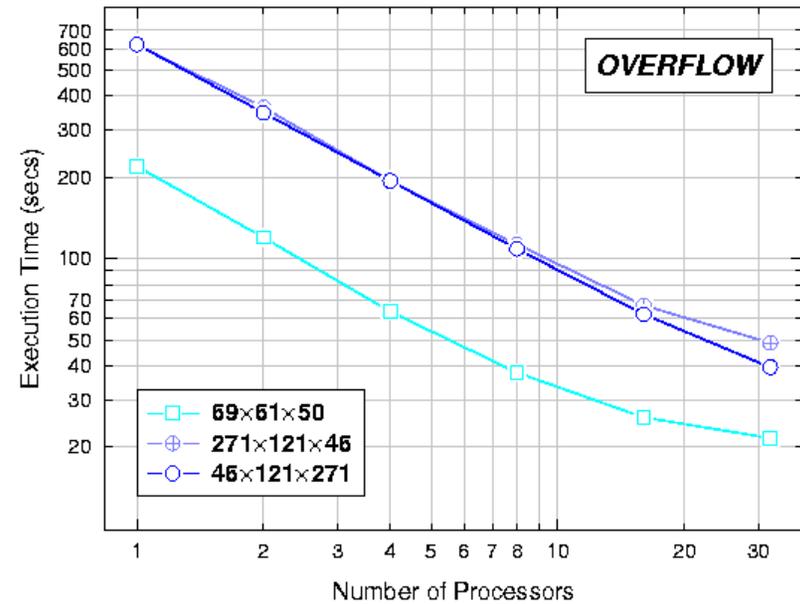
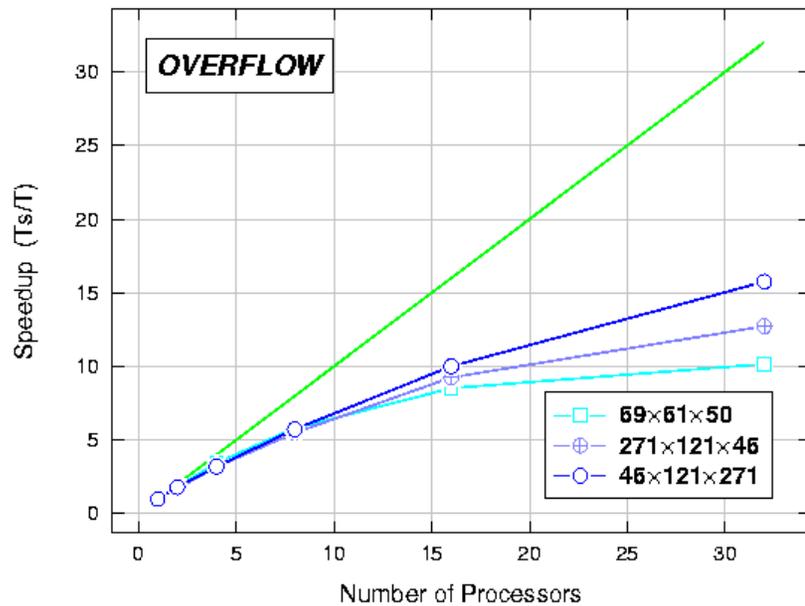


Integrated Programming Environment

- Parallelization
- Performance Evaluation/Visualization
- Machine Specific Code Generation Intelligent GUI
- Code Cleanup and Restructuring
- Serial Optimization



Computer Assisted Parallelization



Immediate and Long Term Collaboration

- Immediately
 - Full spectrum compute resources - Vector and Parallel
 - Parallel computing support of DAO
 - hardware expertise: 288 Silicon Graphics O2K, mass storage, networking, the next prototype for DAO data
 - software expertise: Code migration, tuning and scaling, data fusion, CapTools, Legacy Code - set up a team to assist DAO
 - Relationship with ASCI
 - parallelization expertise
 - a level beyond the 256p SGI --> 512 and 1024
- Longer term - intermediate products also
 - Distributed High Performance Computing
 - Information Power Grid - NCCS

