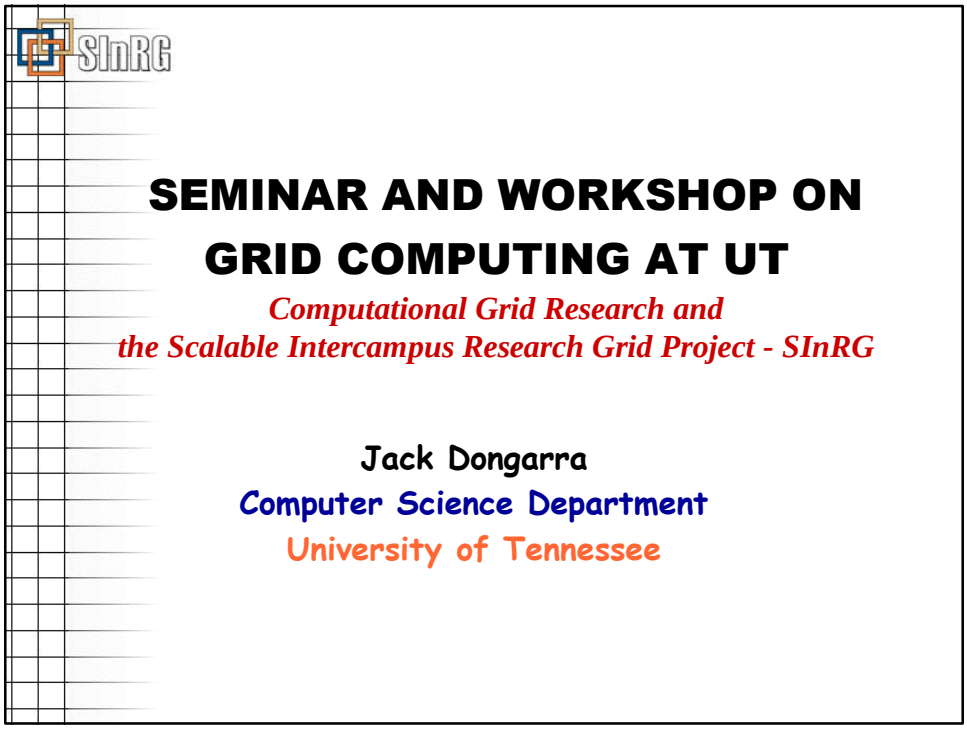

SEMINAR AND WORKSHOP ON GRID COMPUTING AT UT

Computational Grid Research and the Scalable Intercampus Research Grid Project - SInRG

Jack Dongarra
Computer Science Department
University of Tennessee

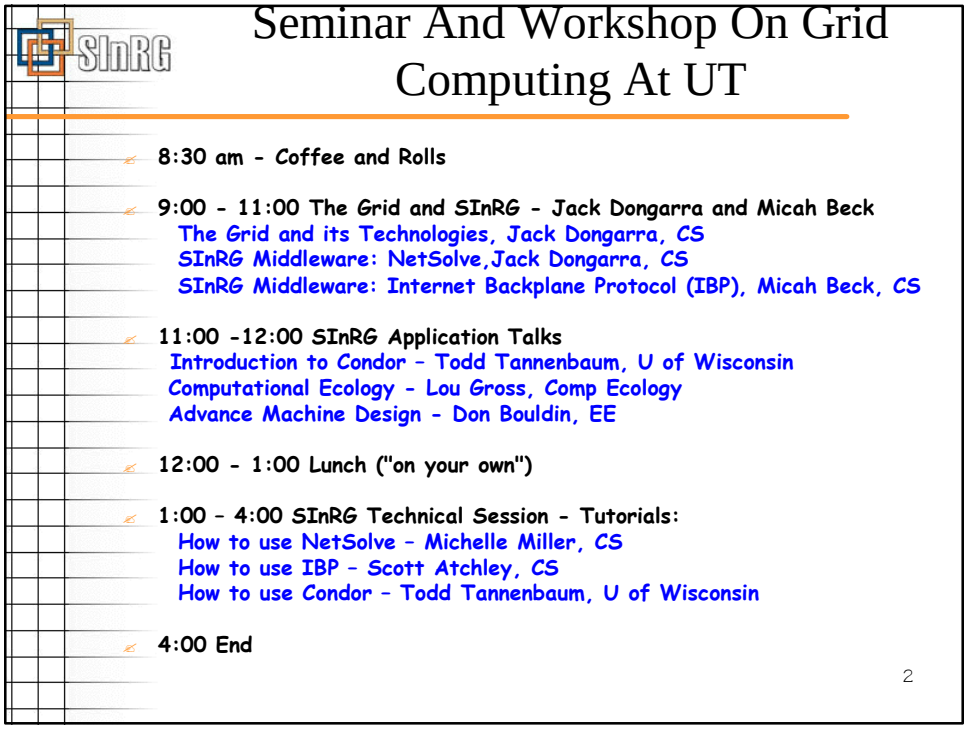

Seminar And Workshop On Grid Computing At UT

- 8:30 am - Coffee and Rolls
- 9:00 - 11:00 The Grid and SInRG - Jack Dongarra and Micah Beck
  - The Grid and its Technologies, Jack Dongarra, CS
  - SInRG Middleware: NetSolve,Jack Dongarra, CS
  - SInRG Middleware: Internet Backplane Protocol (IBP), Micah Beck, CS
- 11:00 -12:00 SInRG Application Talks
  - Introduction to Condor – Todd Tannenbaum, U of Wisconsin
  - Computational Ecology - Lou Gross, Comp Ecology
  - Advance Machine Design - Don Bouldin, EE
- 12:00 - 1:00 Lunch ("on your own")
- 1:00 – 4:00 SInRG Technical Session - Tutorials:
  - How to use NetSolve – Michelle Miller, CS
  - How to use IBP – Scott Atchley, CS
  - How to use Condor – Todd Tannenbaum, U of Wisconsin
- 4:00 End

2

What is Grid Computing?

Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations

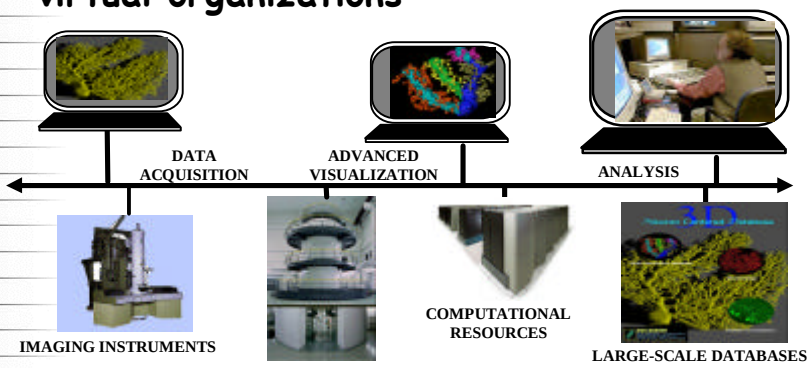

3

The Computational Grid is…

- …a distributed control infrastructure that allows applications to treat compute cycles as commodities.
- Power Grid analogy
  - Power producers: machines, software, networks, storage systems
  - Power consumers: user applications
- Applications draw power from the Grid the way appliances draw electricity from the power utility.
  - Seamless
  - High-performance
  - Ubiquitous
  - Dependable

4

Computational Grids and Electric Power Grids

- **Why the Computational Grid is like the Electric Power Grid**
  - Electric power is ubiquitous
  - Don't need to know the source of the power (transformer, generator) or the power company that serves it

- **Why the Computational Grid is different from the Electric Power Grid**
  - Wider spectrum of performance
  - Wider spectrum of services
  - Access governed by more complicated issues
    - » Security
    - » Performance
    - » Socio-political factors

5

An Emerging Grid Community

1995-2000

- **"Grid book" gave a comprehensive view of the state of the art**
- **Important infrastructure and middleware efforts initiated**
  - » Globus
  - » Legion
  - » Condor
  - » NetSolve, Ninf
  - » Storage Resource Broker
  - » Network Weather Service
  - » AppLeS, …

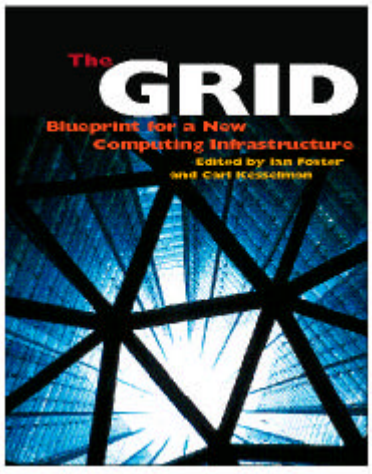

6

Grids are Hot

IPG NAS-NASA	http://nas.nasa.gov/~wej/home/IPG
Globus	http://www.globus.org/
Legion	http://www.cs.virgina.edu/~grimshaw/
AppLeS	http://www-cse.ucsd.edu/groups/hpcl/apples
NetSolve	http://www.cs.utk.edu/netsolve/
NINF	http://phase.etl.go.jp/ninf/
Condor	http://www.cs.wisc.edu/condor/
CUMULVS	http://www.epm.ornl.gov/cs/cumulvs.html
WebFlow	http://www.npac.syr.edu/users/gcf/
LoCI	http://loci.cs.utk.edu/

7

SInRG The Grid

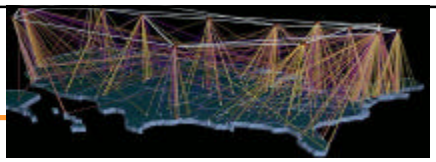

- To treat CPU cycles and software like commodities.
- Napster on steroids.
- Enable the coordinated use of geographically distributed resources – in the absence of central control and existing trust relationships.
- Computing power is produced much like utilities such as power and water are produced for consumers.
- Users will have access to "power" on demand
- "When the Network is as fast as the computer's internal links, the machine disintegrates across the Net into a set of special purpose appliances"
  - Gilder Technology Report June 2000

8

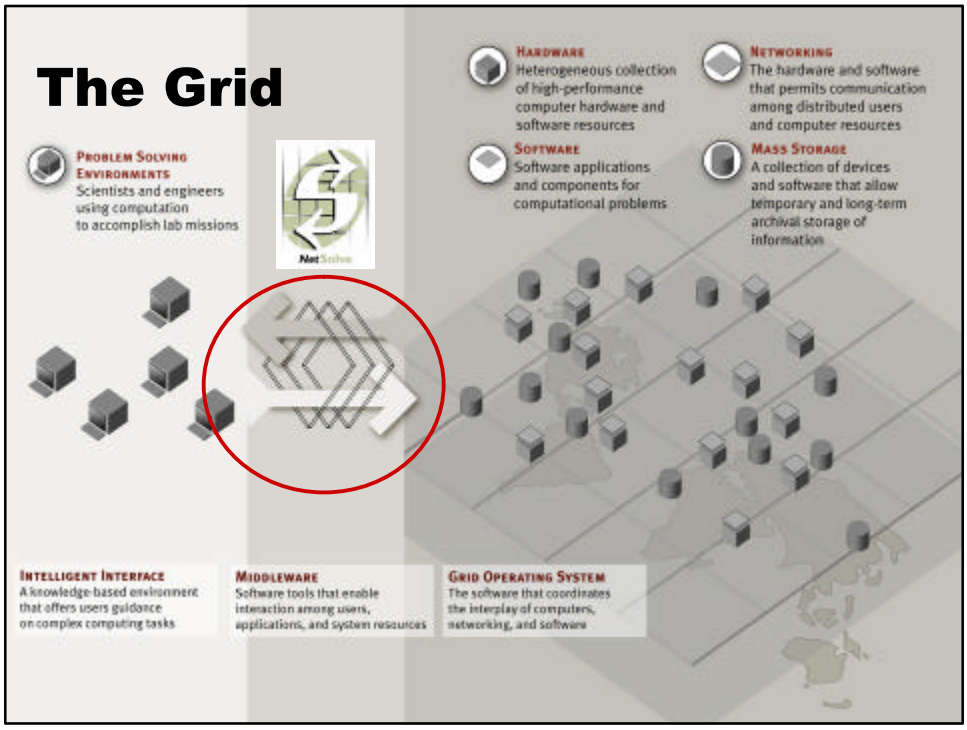
The Grid
Problem Solving Environments
Scientists and engineers using computation to accomplish lab missions
Hardware
Heterogeneous collection of high-performance computer hardware and software resources
Networking
The hardware and software that permits communication among distributed users and computer resources
Software
Software applications and components for computational problems
Mass Storage
A collection of devices and software that allow temporary and long-term archival storage of information
Intelligent Interface
A knowledge-based environment that offers users guidance on complex computing tasks
Middleware
Software tools that enable interaction among users, applications, and system resources
Grid Operating System
The software that coordinates the interplay of computers, networking, and software

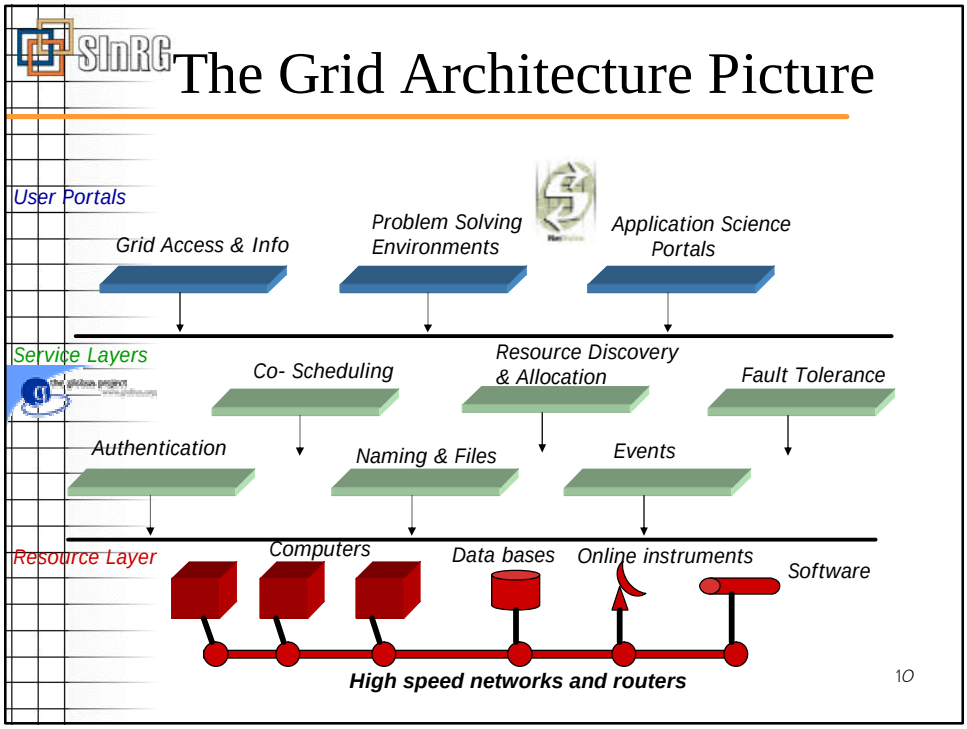
SInRG
The Grid Architecture Picture
User Portals
Grid Access & Info
Problem Solving Environments
Application Science Portals
Service Layers
Co- Scheduling
Resource Discovery & Allocation
Fault Tolerance
Authentication
Naming & Files
Events
Resource Layer
Computers
Data bases
Online instruments
Software
High speed networks and routers
10

Globus Grid Services

- The Globus toolkit provides a range of basic Grid services
  - Security, information, fault detection, communication, resource management, ...
- These services are simple and orthogonal
  - Can be used independently, mix and match
  - Programming model independent
- For each there are well-defined APIs
- Standards are used extensively
  - E.g., LDAP, GSS-API, X.509, ...
- You don't program in Globus, it's a set of tools like Unix

Basic Grid Building Blocks

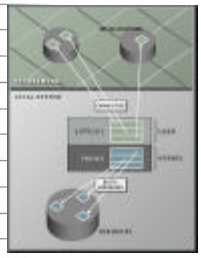

IBP – Internet Backplane Protocol is middleware for managing and using remote storage.

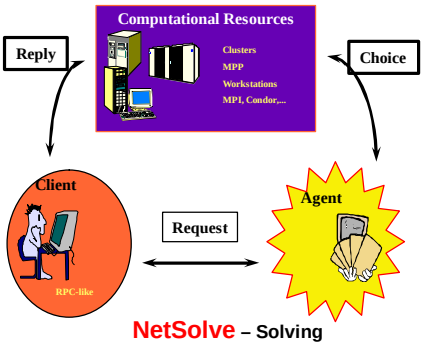

NetSolve – Solving computational problems remotely

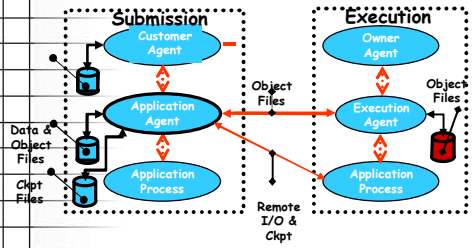

Condor – harnessing idle workstations for high-throughput computing

Maturation of Grid Computing

- Research focus moving from building of basic infrastructure and application demonstrations to
  - Middleware
  - Usable production environments
  - Application performance
  - Scalability ✍ Globalization
- Development, research, and integration happening outside of the original infrastructure groups
- Grids becoming a first-class tool for scientific communities
  - GriPhyN (Physics), BIRN (Neuroscience), NVO (Astronomy), Cactus (Physics), …

13

Broad Acceptance of Grids as a Critical Platform for Computing

- Widespread interest from government in developing computational Grid platforms

NSF's Cyberinfrastructure

NASA's Information Power Grid

DOE's Science Grid

14

Broad Acceptance of Grids as a Critical Platform for Computing

- Widespread interest from industry in developing computational Grid platforms
- IBM, Sun, Entropia, Avaki, Platform, …

On August 2, 2001, IBM announced a new corporate initiative to support and exploit Grid computing.

AP reported that IBM was investing $4 billion into building 50 computer server farms around the world.

AVAKI

ENTROPIA

SETI@home

- Use thousands of Internet-connected PCs to help in the search for extraterrestrial intelligence.
- Uses data collected with the Arecibo Radio Telescope, in Puerto Rico
- When their computer is idle or being wasted this software will download a 300 kilobyte chunk of data for analysis.

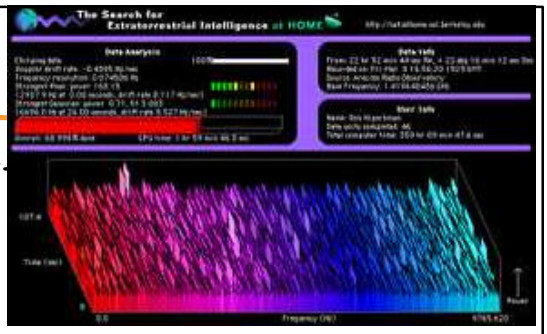

- The results of this analysis are sent back to the SETI team, combined with the crunched data from the many thousands of other SETI@home participants.

Grid Computing - from ET to Anthrax


CNN.com / SCI-TECH

PCs tapped to help fight anthrax

January 22, 2002 Posted: 12:10 PM EST (1710 GMT)

SAN JOSE, California (AP) -- A coalition of scientists and technology companies is asking people around the world to use their computers' extra processing power to help search for a cure for anthrax.

The project follows similar efforts to use "distributed computing" to hunt for extraterrestrial life and a cure for cancer. It is being launched Tuesday to help Oxford University researchers find ways to treat anthrax that can no longer be treated by antibiotics.

The project is based on the premise that the average personal computer uses between 15 percent and 18 percent of its processing power at any given time. It employs 'peer-to-peer' technology, in which millions of computers can share files over the Internet.

Participants download a screen-saver that runs whenever their computers have resources to spare, and uses that power to perform computations for the project. When the user connects to the Internet, the computer sends data back to a central hub and gets another assignment.

The company that designed the program, United Devices Inc. of Austin, Texas, promises that no personal information on participants' PCs can be compromised while they take part.

If the project attracts more than 100,000 participants, it can give researchers more computational power than the world's 10 best supercomputers combined, said United Devices spokesman Andy Prince.

With enough participants, the project would provide researchers 10 times more power than the world's best supercomputer, said Graham Richards, the Oxford professor leading the study.

"The screen-saver doesn't cost you anything, and at least you're taking part in something, adding your bit," he said.

Intel, Microsoft involved

Scientists have discovered that the anthrax toxin is made up of three proteins that

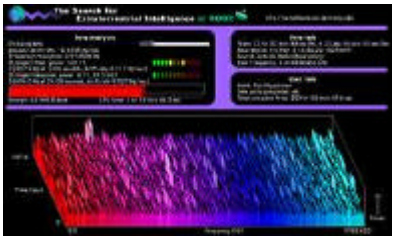

17

Grids Form the Basis of a National Information Infrastructure

SInRG

August 9, 2001: NSF Awarded $53,000,000 to SDSC/NPACI and NCSA/Alliance for TeraGrid

TeraGrid will provide in aggregate

- ***13.6 trillion calculations per second***
- ***Over 600 trillion bytes of immediately accessible data***
- ***40 gigabit per second network speed***
- ***Provide a new paradigm for data-oriented computing***
  - *Critical for disaster response, genomics, environmental modeling, etc.*

18

Drivers Wanted

- Where are new Grid researchers and developers being trained
- How many CS departments have faculty with a focus in Grid computing
- How can we increase the number of students with expertise and experience in Grid computing
  - Authors of the Grid Book will not live forever …

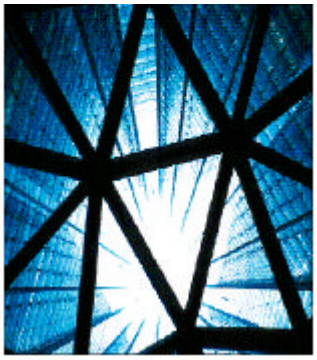

19

UTK's Grid Research Effort

- Create a Grid prototype on one campus and leverage locality of all resources to produce *vertical integration of research* elements:
  - Human collaborator (application scientist)
  - Application software
  - Grid middleware
  - Distributed, federated resource pool
- On site collaborations with researchers from other disciplines will help ensure that the research has broad and real impact.
- Interaction, validate research, test bed, try out ideas

20

The Scalable Intracampus Research Grid for Computer Science Research: SInRG

- **NSF Funded Computer Science CISE Infrastructure Project, additional support from Microsoft Research, Dell Computer, & Sun Microsystems**
- **Build a computational grid for Computer Science research that mirrors the underlying technologies and types of research collaboration that are taking place on the national technology grid**

21

Team of Investigators

- CS Grid iddleware Research (Gang of our)
  - **Jack Dongarra**
  - **Micah Beck**
  - **Rich Wolski**
  - **Jim Plank**
- CS aculty
  - **Michael Berry**
  - **Jens Gregor**
  - **Michael Langston**
  - **Michael Thomason**
  - **Bob Ward**
- Domain Application Collaborators
  - **Don Bouldin (EE)**
  - **Peter Cummings (ChE)**
  - **Lou Gross (CME)**
  - **Tom Hallam (CME)**
  - **Gary Smith (Radiology)**
  - **Christian Halloy (JICS)**
  - **DeWitt Latimer (DII)**

22

Resources: Grid Service Cluster

- **Computation**
  - used to run Grid controlware
  - Committed dynamically to augment other CPUs on Grid
- **Storage**
  - State management
    - » data caching
    - » migration and fault-tolerance
- **Network**
  - allows dynamic reconfigutation of resources

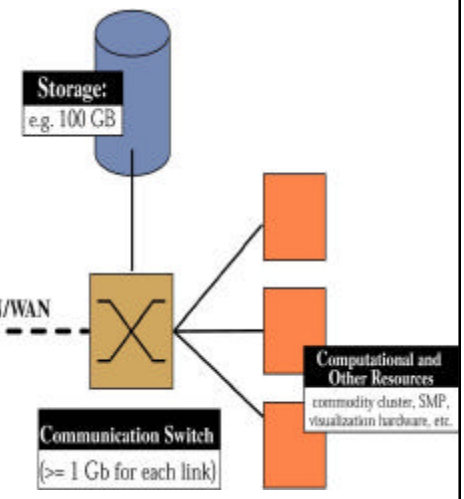

23

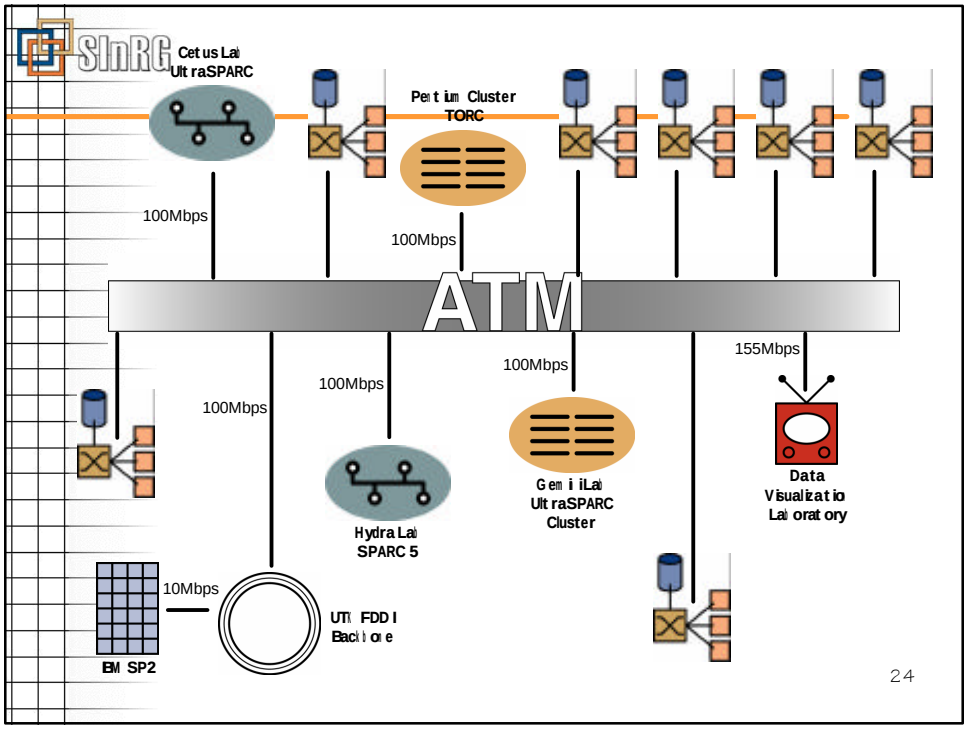

24

University of Tennessee Deployment: Scalable Intracampus Research Grid SInRG

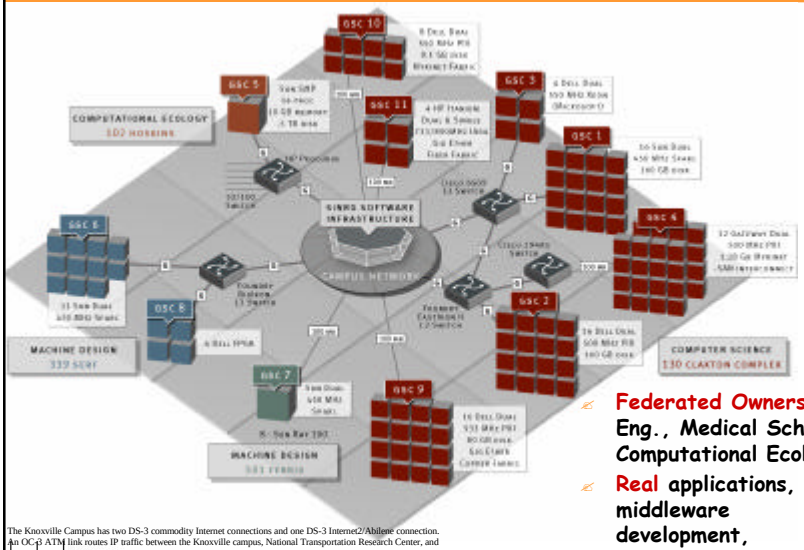

The Knoxville Campus has two DS-3 commodity Internet connections and one DS-3 Internet2/Abilene connection. An OC-3 ATM link routes IP traffic between the Knoxville campus, National Transportation Research Center, and Oak Ridge National Laboratory. UT participates in several national networking initiatives including Internet2 (I2), Abilene, the federal Next Generation Internet (NGI) initiative, Southern Universities Research Association (SURA) Regional Information Infrastructure (RII), and Southern Crossroads (SoX).

The UT campus consists of a meshed ATM OC-12 being migrated over to switched Gigabit by early 2002.

- **Federated Ownership**: CS, Chem Eng., Medical School, Computational Ecology, El. Eng.
- **Real** applications, middleware development, logistical networking

25

Middleware Research

- **Challenge:** Provide a solid, integrated, high-performance computing platform despite
  - Widely varying load conditions
  - Non-dedicated, heterogeneous and federated resource pool
  - Faults, power failures, football games, etc.
- Integration is key and non-trivial
  - NetSolve, NWS, IBP must be able to work together and with other middleware not-invented-here

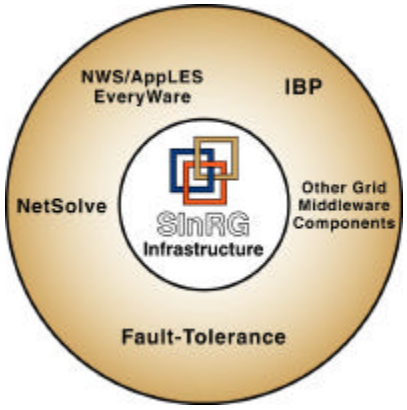

26

Comprehensive Research Approach

- NetSolve (Dongarra)
  - programming abstractions and resource control
- Internet Backplane Protocol (*IBP* – Beck, Plank)
  - distributed storage management
- Network Weather Service (Wolski)
  - dynamic performance prediction
- EveryWare (Wolski)
  - toolkit for leveraging *multiple Grid infrastructures* and resources to build adaptive programs
- G-Commerce (Wolski, Plank)
  - Provably stable market-economies for the Grid that support dynamic resource allocation
- Fault-tolerance (Plank, Dongarra)
  - process robustness and migration

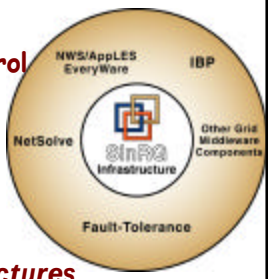

27

Application Collaborations in SInRG

- All developing apps targeted for SInRG
- Apps will drive CS research
- Model of national grid community
- Apps committing personnel and equipment

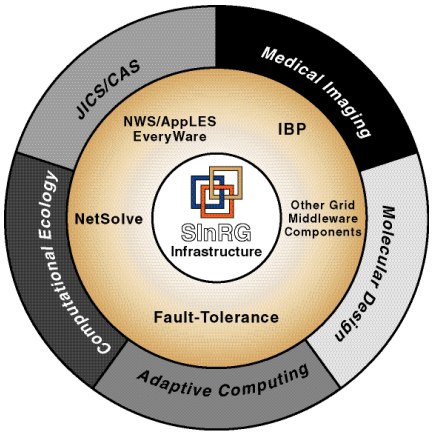

- All collaborative work supports the Grid research agenda.

28

Applications Research

- Advanced Machine Design (Bouldin, Langston)
  - Adaptive, reconfigurable computers
- Medical Imaging (Smith, Gregor, Thomason)
  - High-performance image reconstruction
- Computational Ecology (Gross, Hallam, Berry)
  - Large-scale, coupled simulations of diverse ecosystems
- Molecular Design (Cummings, Ward)
  - Computational chemistry

29

SInRG

- SInRG provides a testbed
  - CS grid middleware
  - Computational Science applications
- Many hosts, co-existing in a loose confederation tied together with high-speed links.
- Users have the illusion of a very powerful computer on the desk.
- Spectrum of users

30

Summary

- SInRG constitutes a novel Grid research approach
  - Empirical
  - Vertically integrated and collaborative
  - Both technology and applications driven
  - *A real* research project
- UT Grid research efforts are drawing national international attention
  - Burgeoning user communities for software artifacts
  - Research and infrastructure funding
  - Persistent installations

31

Motivation for NetSolve

Design an easy-to-use tool to provide efficient and uniform access to a variety of scientific packages on UNIX and Window's platforms

- Client-Server Design
- Non-hierarchical system
- Load Balancing and Fault Tolerance
- Heterogeneous Environment Supported
- Multiple and simple client interfaces
- Built on standard components

32

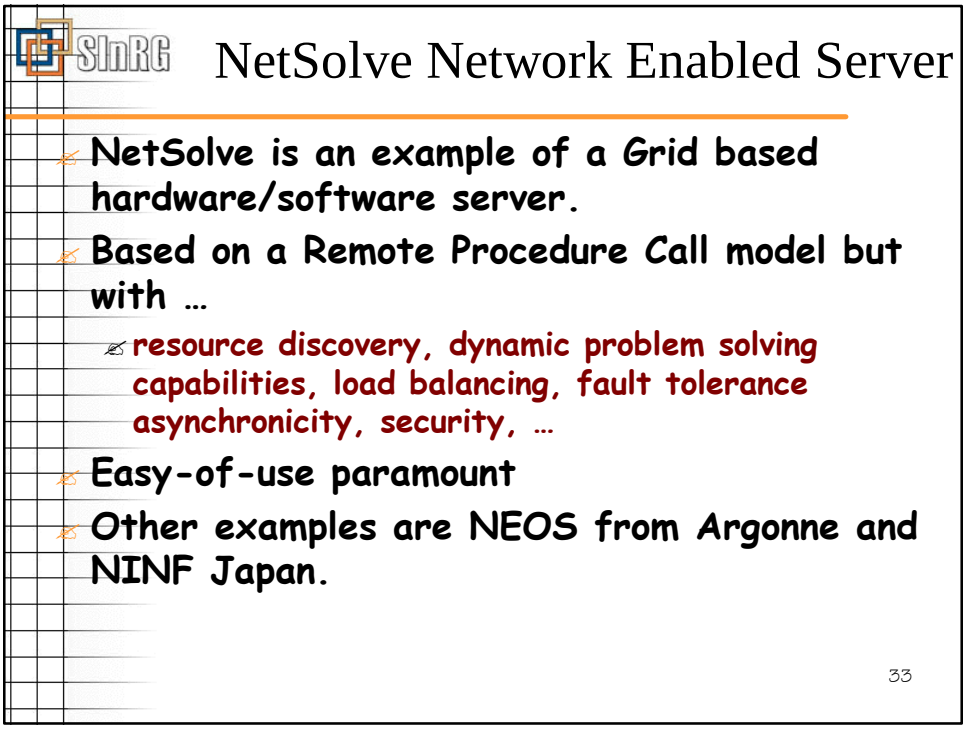
SInRG
NetSolve Network Enabled Server
NetSolve is an example of a Grid based hardware/software server.
Based on a Remote Procedure Call model but with …
resource discovery, dynamic problem solving capabilities, load balancing, fault tolerance asynchronicity, security, …
Easy-of-use paramount
Other examples are NEOS from Argonne and NINF Japan.
33

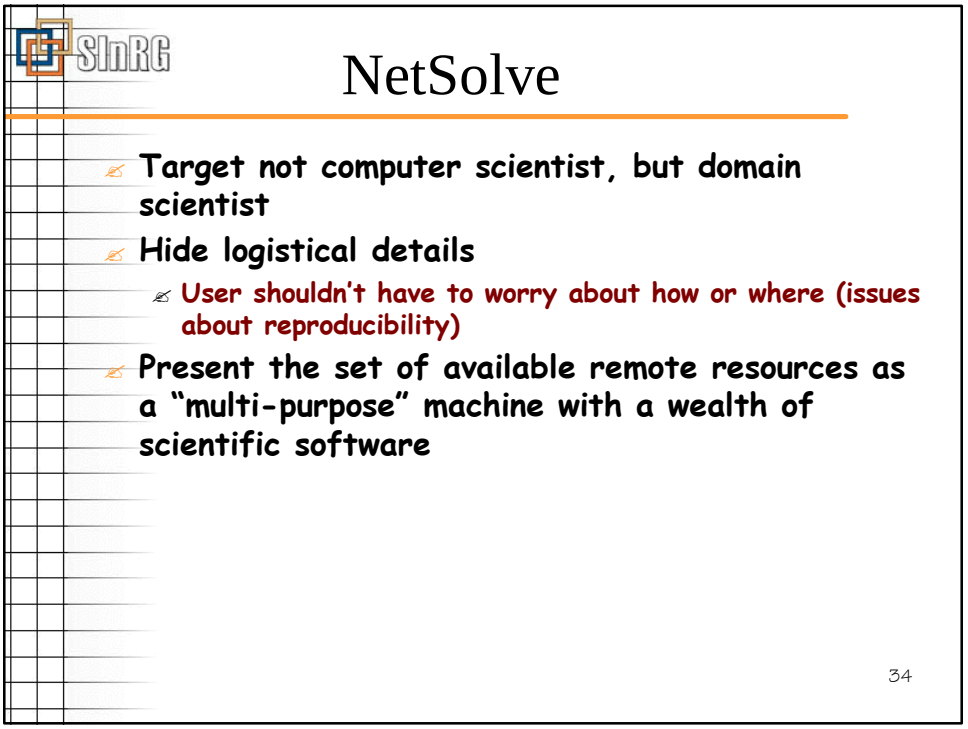
SInRG
NetSolve
Target not computer scientist, but domain scientist
Hide logistical details
User shouldn't have to worry about how or where (issues about reproducibility)
Present the set of available remote resources as a "multi-purpose" machine with a wealth of scientific software
34

NetSolve: The Big Picture

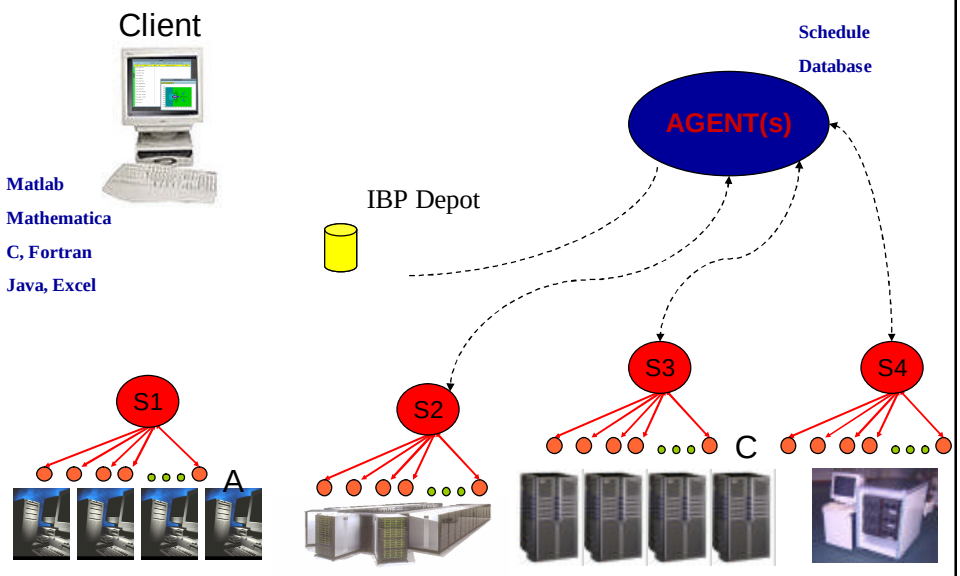

No knowledge of the grid required, RPC like.

NetSolve: The Big Picture

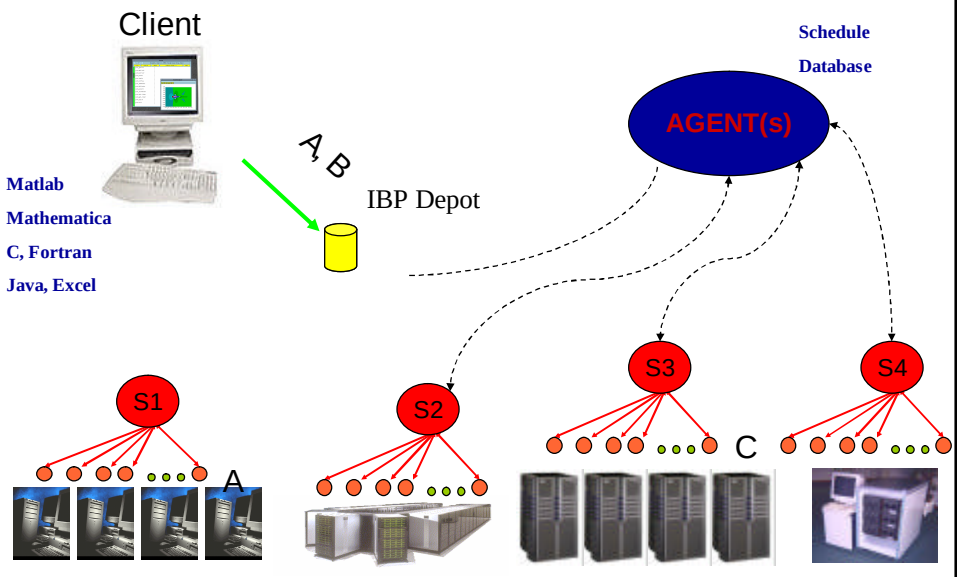

No knowledge of the grid required, RPC like.

NetSolve: The Big Picture

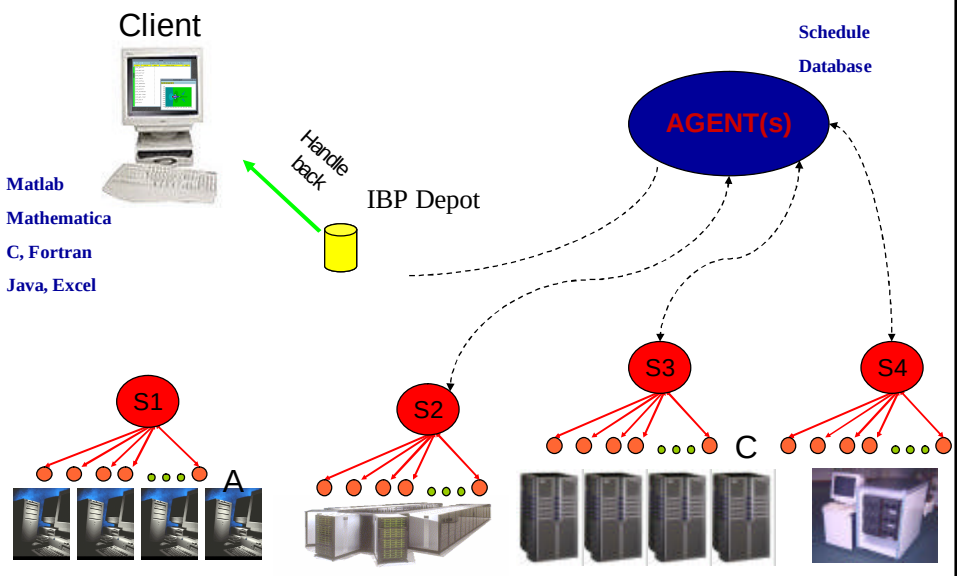

No knowledge of the grid required, RPC like.

NetSolve: The Big Picture

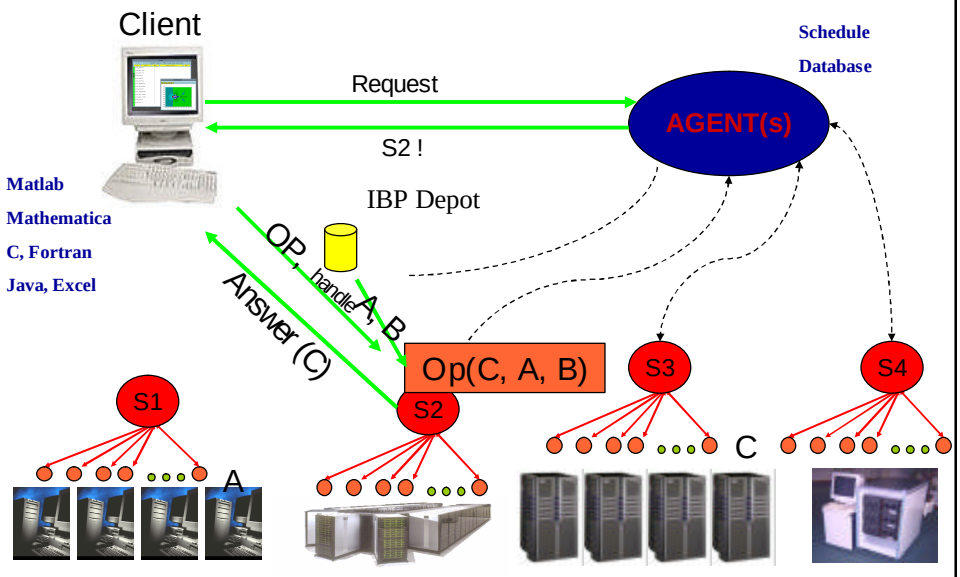

No knowledge of the grid required, RPC like.

Basic Usage Scenarios

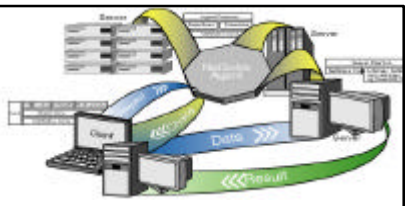

- Grid based numerical library routines
  - User doesn't have to have software library on their machine, LAPAC , SuperLU, ScaLAPAC , PETSc, A TEC, ARPAC
- Task farming applications
  - "Pleasantly parallel" execution
  - eg Parameter studies
- Remote application execution
  - Complete applications with user specifying input parameters and receiving output
- "Blue Collar" Grid Based Computing
  - Does not require deep knowledge of network programming
  - Level of expressiveness right for many users
  - User can set things up, no "su" required
  - In use today, up to 200 servers in 9 countries
- Can plug into Globus, Condor, NINF, …

39

NetSolve Agent

- Name server for the system.
- Information Service
  - client users and administrators can query the hardware and software services available.
- Resource scheduler
  - maintains both static and dynamic information regarding the NetSolve server components to use for the allocation of resources

40

NetSolve Agent

- Resource Scheduling (cont'd):
  - CPU Performance (LINPAC).
  - Network bandwidth, latency.
  - Server workload.
  - Problem size/algorithm complexity.
  - Calculates a "Time to Compute." for each appropriate server.
  - Notifies client of most appropriate server.

41

NetSolve Client

- Function Based Interface.
- Client program embeds call from NetSolve's API to access additional resources.
- Interface available to C, Fortran, Matlab, and Mathematica.
- Opaque networking interactions.
- NetSolve can be invoked using a variety of methods: blocking, non-blocking, task farms, …

42

NetSolve Client

- Intuitive and easy to use.
- Matlab Matrix multiply e.g.:
  - A matmul(B, C)

A = netsolve('matmul', B, C);

- Possible parallelisms hidden.

43

NetSolve Client

i. Client makes request to agent.

ii. Agent returns list of servers.

iii. Client tries first one to solve problem.

44

NPACI Alpha Project - MCell: 3-D Monte-Carlo Simulation of Neuro-Transmitter Release in Between Cells

- UCSD (F. Berman, H. Casanova, M. Ellisman), Salk Institute (T. Bartol), CMU (J. Stiles), UTK (Dongarra, M. Miller, R. Wolski)
- Study how neurotransmitters diffuse and activate receptors in synapses
- blue unbounded, red singly bounded, green doubly bounded closed, yellow doubly bounded open

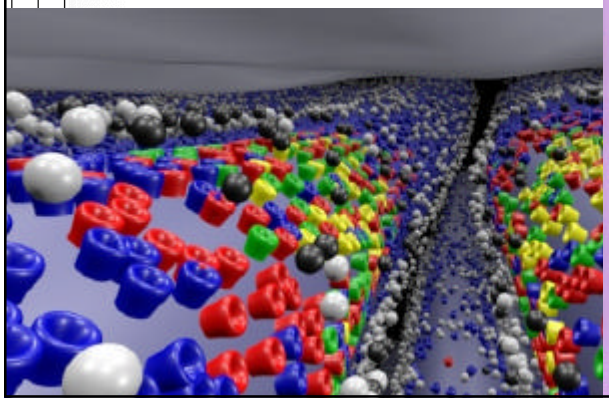

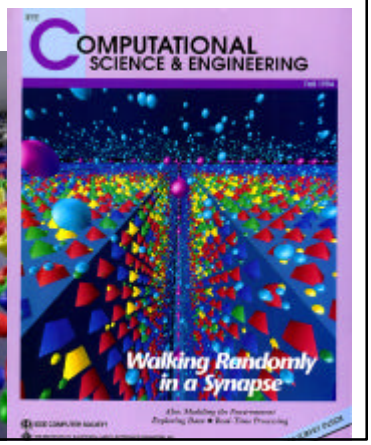

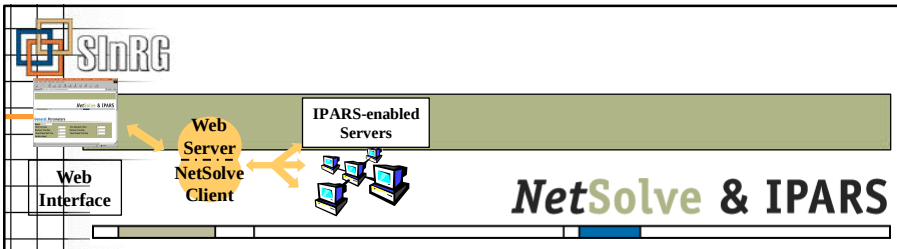

- Integrated Parallel Accurate Reservoir Simulator.
  - Mary Wheeler's group, UT-Austin
- Reservoir and Environmental Simulation.
  - models black oil, waterflood, compositions
  - 3D transient flow of multiple phase
- Integrates Existing Simulators.
- Framework simplified development
  - Provides solvers, handling for wells, table lookup.
  - Provides pre/postprocessor, visualization.
- Full IPARS access without Installation.
- IPARS Interfaces:
  - C, FORTRAN, Matlab, Mathematica, and Web.

46

Netsolve and SCIRun

SCIRun torso defibrillator application – Chris Johnson, U of Utah

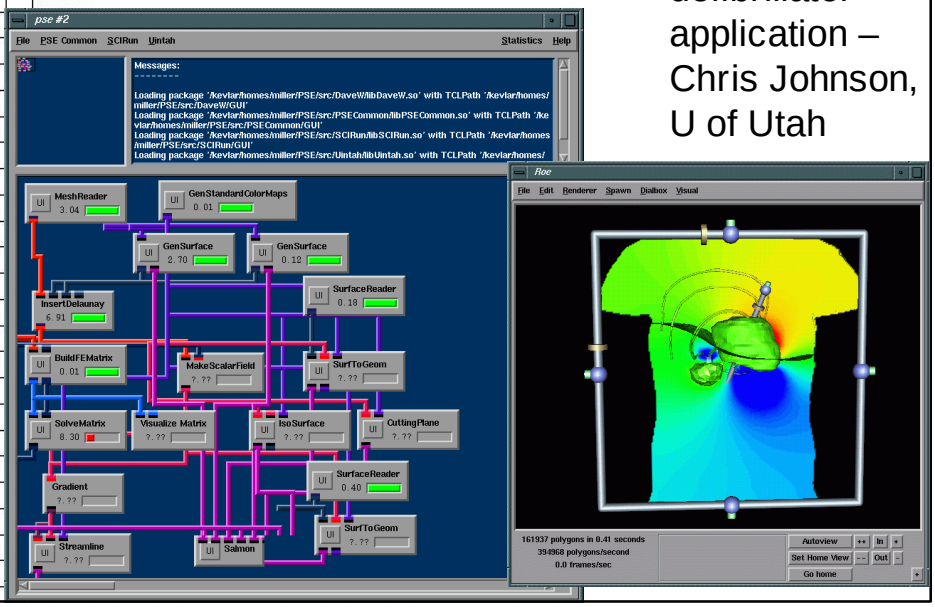

NetSolve: A Plug into the Grid

Grid middleware

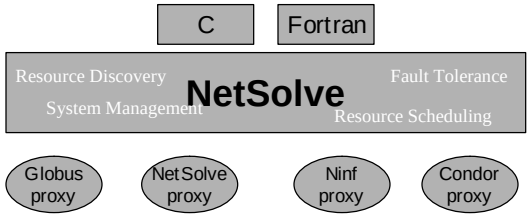

48

NetSolve: A Plug into the Grid

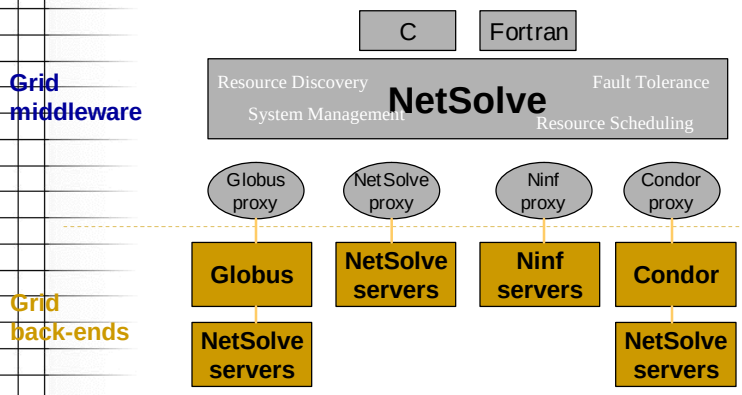

49

NetSolve: A Plug into the Grid

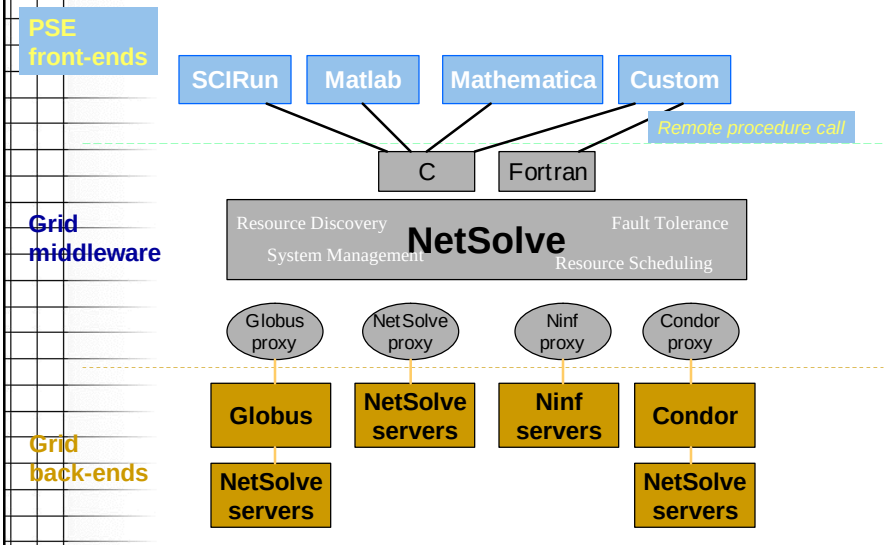

50

DEMO

- Monitor
- Simple calls
- Non blocking calls
- Sparse matrix software
- Graphic Mandelbrot
- Graphic quaduature

51

SEMINAR AND WORKSHOP ON GRID COMPUTING AT UT

- 8:30 am - Coffee and Rolls
- 9:00 - 11:00 The Grid and SInRG - Jack Dongarra and Micah Beck
  - The Grid and its Technologies
  - SInRG Middleware: NetSolve
  - SInRG Middleware: Internet Backplane Protocol (IBP)
- 11:00 -12:00 SInRG Application Talks
  - Introduction to Condor – Todd Tannenbaum, U of Wisconsin
  - Computational Ecology - Lou Gross
  - Advance Machine Design - Don Bouldin
- 12:00 - 1:00 Lunch ("on your own")
- 1:00 – 4:00 SInRG Technical Session - Tutorials:
  - How to use NetSolve – Michelle Miller
  - How to use IBP – Scott Atchley
  - How to use Condor – Todd Tannenbaum, U of Wisconsin
- 4:00 End

52