The Virtual Grid Application Development Software (VGrADS) Project

Ken Kennedy
Center for High Performance Software
Rice University

http://vgrads.rice.edu/

VGrADS
Virtual Grid Application Development Software Project

The VGrADS Team

- **VGrADS is an NSF-funded Information Technology Research project**

UCSB
Rich Wolski

UCSD
Fran Berman
Andrew Chien
Henri Casanova

RICE
Keith Cooper
Ken Kennedy
Charles Koelbel
Richard Tapia
Linda Torczon

THE UNIVERSITY of TENNESSEE
Jack Dongarra

ISI
Carl Kesselman

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL
Dan Reed

Lennart Johnsson

- **Plus *many* graduate students, postdocs, and technical staff!**

VGrADS
Virtual Grid Application Development Software Project

The VGrADS Vision: National Distributed Problem Solving

- **Where We Want To Be**
  - —Transparent Grid computing
    - Submit job
    - Find & schedule resources
    - Execute efficiently
- **Where We Are**
  - —Low-level hand programming
  - —Programmer needs to manage
    - Heterogeneous resources
    - Computation and data movement scheduling
    - Fault tolerance and performance adaptation
- **What Do We Need?**
  - —A more abstract view of the Grid
    - Each developer sees a *scalable* "virtual grid"
  - —Simplified programming models built on the abstract view
    - Permit the application developer to *focus on the problem*

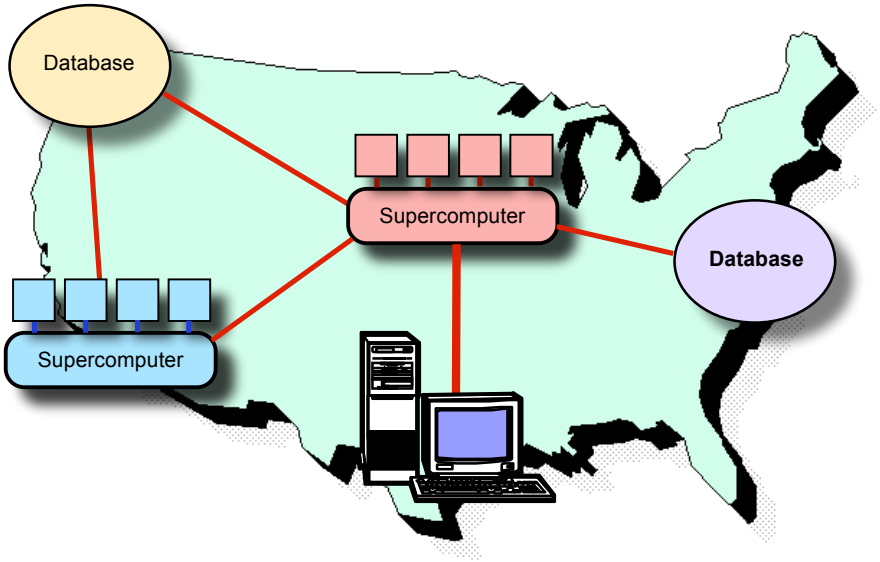

Abstraction: Virtual Grid Execution System (vgES)

- A Virtual Grid (VG) takes
  - Shared heterogeneous resources
  - Scalable information service
- and provides
  - An hierarchy of application-defined aggregations (e.g. ClusterOf) with constraints (e.g. processor type) and rankings

- Virtual Grid Execution System (vgES) implements VG
  - VG Definition Language (vgDL)
  - VG Find And Bind (vgFAB)
  - VG Monitor (vgMON)
  - VG Application Launch (VgLAUNCH+DVCW)
  - VG Resource Info (vgAgent)

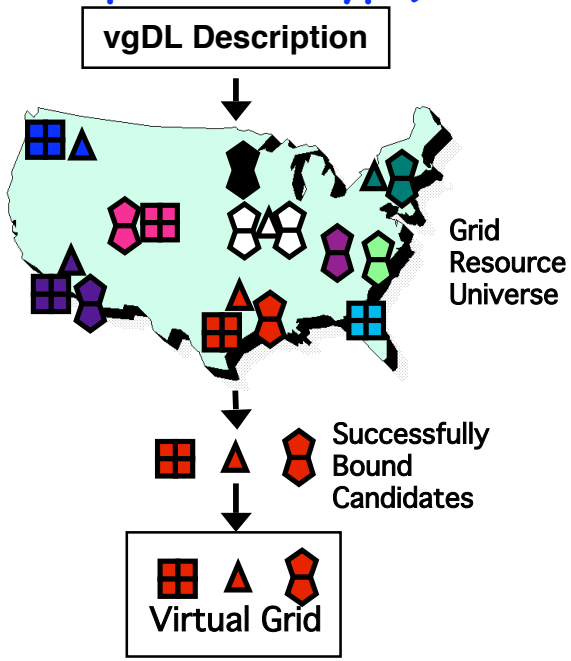

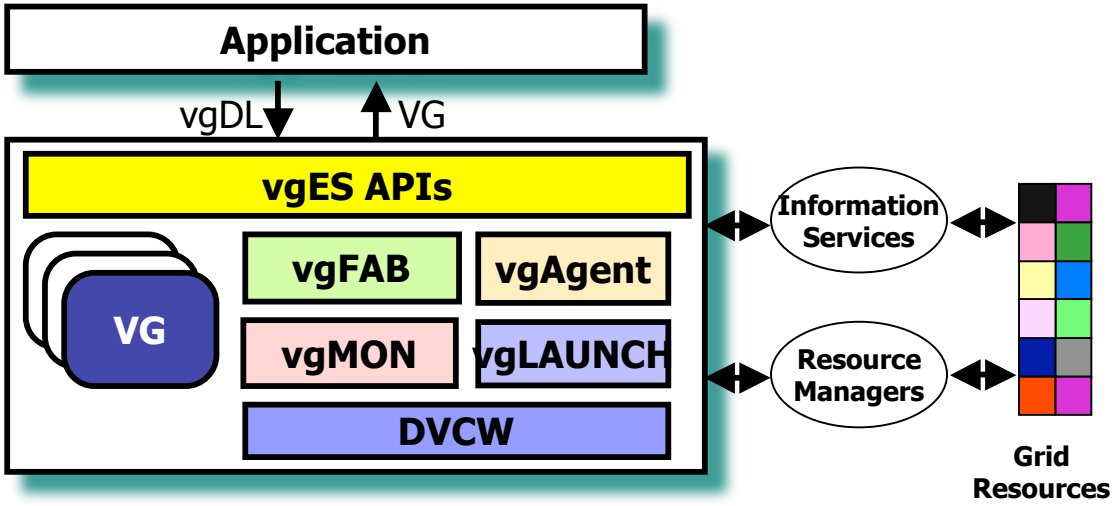

Tools: Scheduling and Fault Tolerance Methods

VGrADS is studying a range of tools for grid programming tasks, including

- **Scheduling of workflow computations**
  - Off-line look-ahead scheduling dramatically improves in makespan (total time)
  - Accurate performance models significantly affect quality of scheduling
  - Queue wait prediction allows scheduling into batch queues
- **Fault tolerance**
  - Diskless checkpointing for linear algebra computations (application-specific)
  - Temporal reasoning for fault prediction
  - Optimal checkpoint frequency for iterative applications

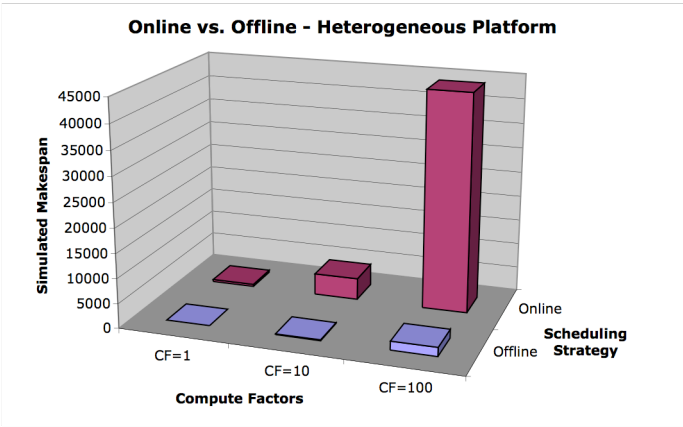

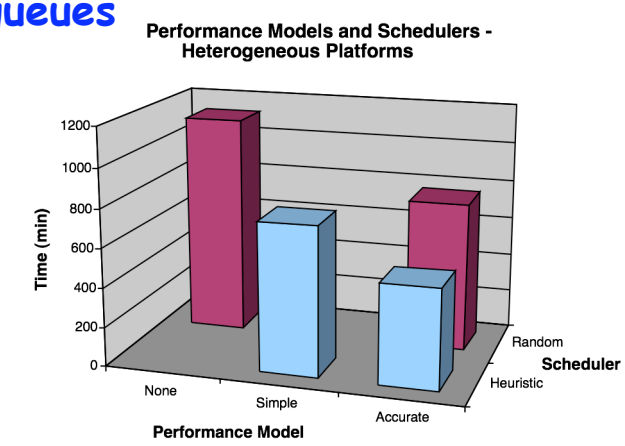

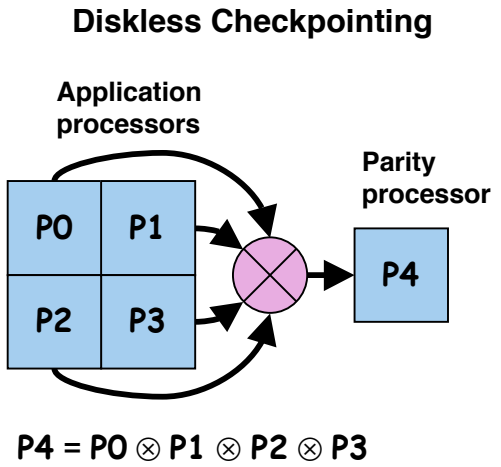

VGrADS
Virtual Grid Application Development Software Project

VGrADS Application Collaborations

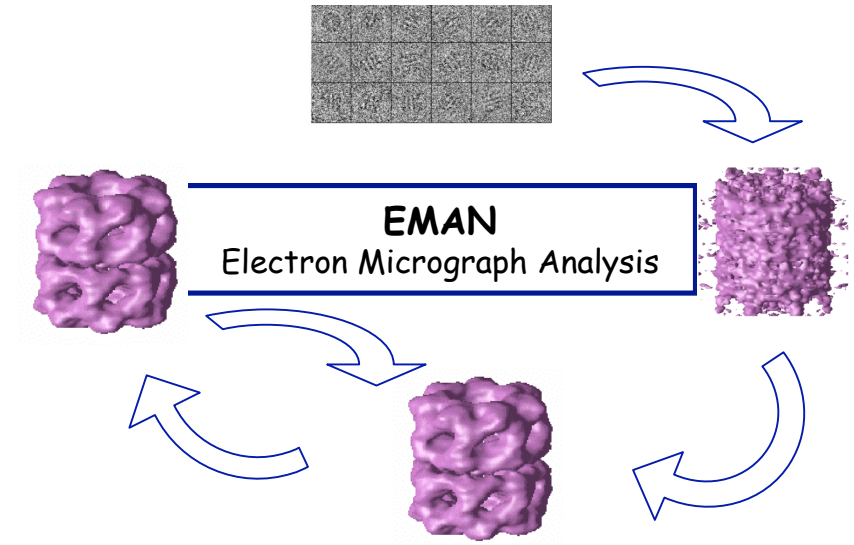

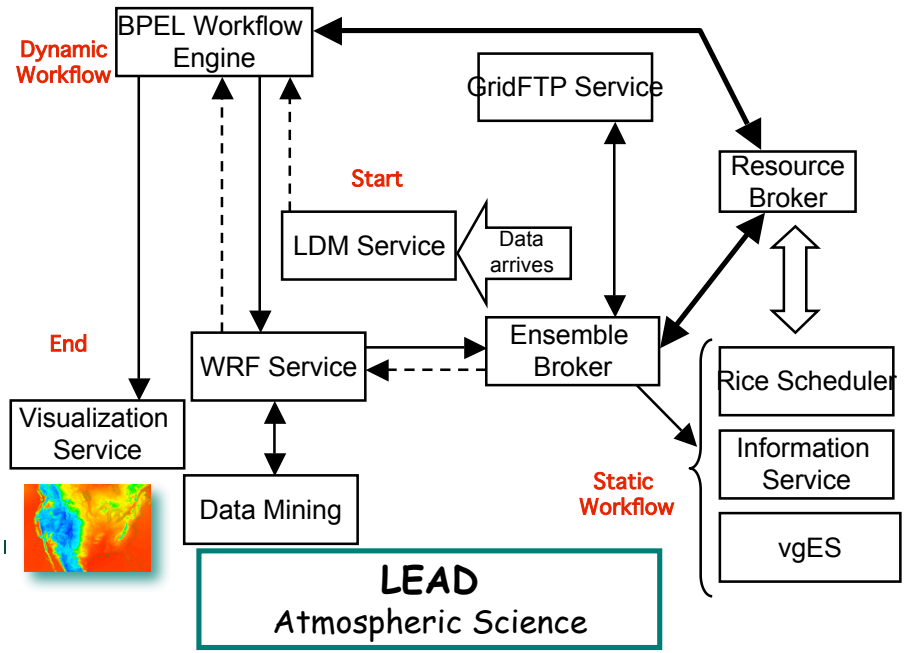

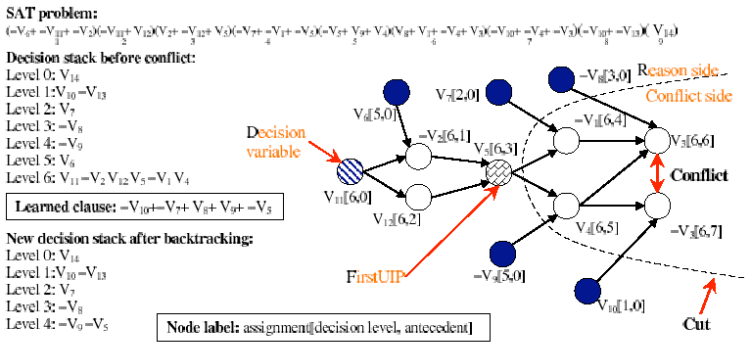

Figure 1: Example of conflict analysis with learning and non-chronological backtracking

GridSAT
Boolean Satisfiability

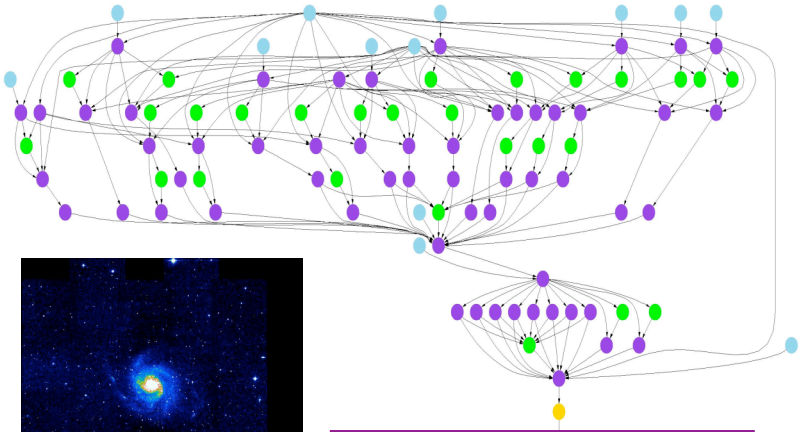

Montage
Astronomy

VGrADS
Virtual Grid Application Development Software Project

VGrADS Demos at SC|05

- vgES / vgMON (UCSD)
  - Runs EMAN application under vgES
  - Track and visualize progress with vgMON
- Batch queue scheduling
  - Schedules EMAN onto resources fronted by batch queues
  - Allows running across clusters
- GridSolve
  - Submits linear algebra problems for solution on the grid ala NetSolve
  - Uses vgES for
    - Integrated performance information
    - Integrated monitoring
    - Fault prediction
    - Integrating the software and resource information repositories