The Virtual Grid Application Development Software (VGrADS) Project

Chuck Koelbel

Rice University

Dan Nurmi

University of California, Santa Barbara

http://vgrads.rice.edu/

VGrADS
Virtual Grid Application Development Software Project

In Memoriam

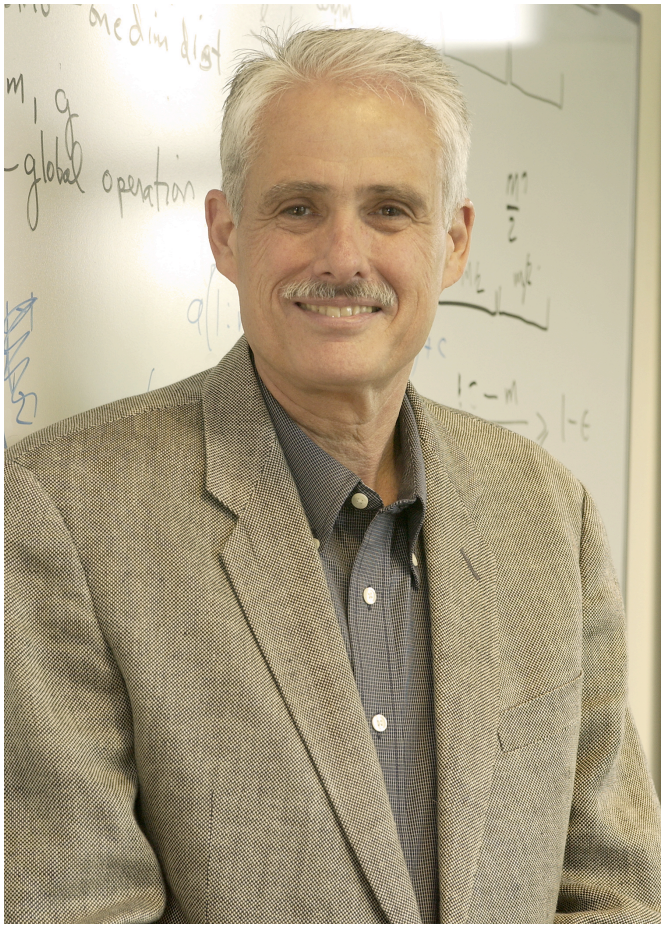

Ken Kennedy
1945-2007

The VGrADS Team

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

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Dan Reed

Keith Cooper
Ken Kennedy
Charles Koelbel
Richard Tapia
Linda Torczon

Jack Dongarra

Carl Kesselman

Fran Berman
Andrew Chien
Henri Casanova

Rich Wolski

Lennart Johnsson

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

VGrADS
Virtual Grid Application Development Software Project

Vision: Global Distributed Problem Solving

- Where We Want To Be
  - Transparent computing
    - In an increasingly distributed space
    - Applications to cloud computing
    - Applications to HPC
- Where We Are
  - Low-level hand programming
  - Programmer must manage:
    - Heterogeneous resources
    - Scheduling of computation and data movement
    - Fault tolerance and performance adaptation
- How Do We Get from Here to There?
  - Separate application development from resource management
    - Through an abstraction called the Virtual Grid
  - Provide tools to bridge the gap
    - Scheduling, resource management, distributed launch, simple programming models, fault tolerance, grid economies

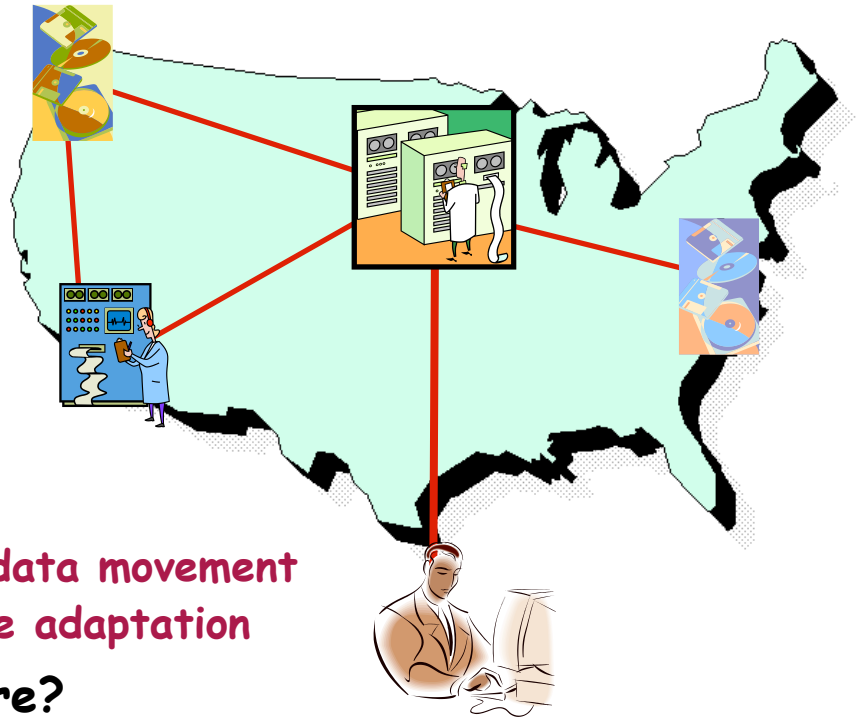

Vision: Global Distributed Problem Solving

- Where We Want To Be
  - Transparent computing
    - In an increasingly distributed space
    - Applications to cloud computing
    - Applications to HPC
- Where We Are
  - Low-level hand programming
  - Programmer must manage:
    - Heterogeneous resources
    - Scheduling of computation and data movement
    - Fault tolerance and performance adaptation
- How Do We Get from Here to There?
  - Separate application development from resource management
    - Through an abstraction called the Virtual Grid
  - Provide tools to bridge the gap
    - Scheduling, resource management, distributed launch, simple programming models, fault tolerance, grid economies

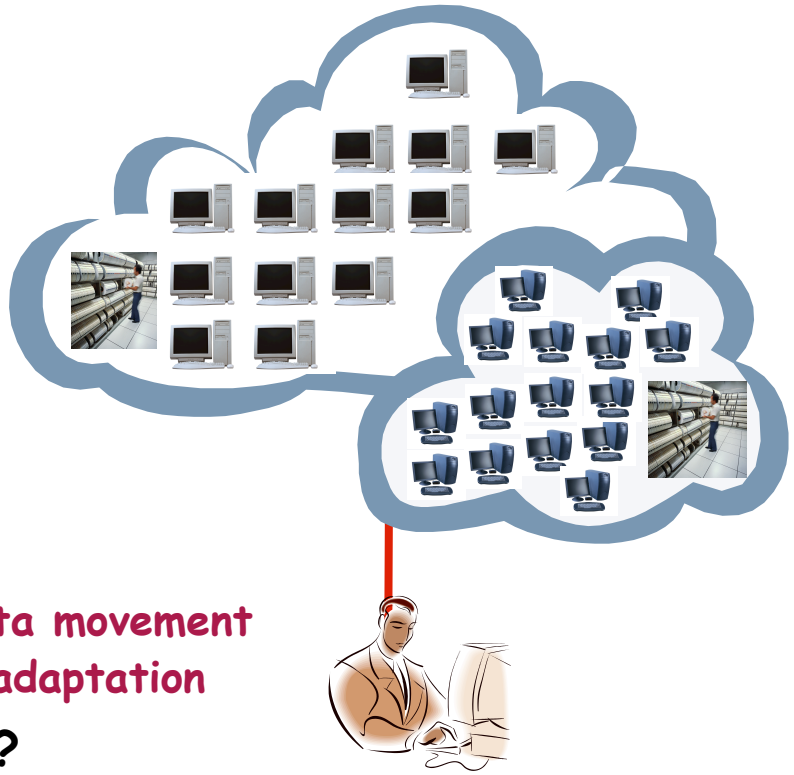

Vision: Global Distributed Problem Solving

- Where We Want To Be
  - Transparent computing
    - In an increasingly distributed space
    - Applications to cloud computing
    - Applications to HPC
- Where We Are
  - Low-level hand programming
  - Programmer must manage:
    - Heterogeneous resources
    - Scheduling of computation and data movement
    - Fault tolerance and performance adaptation
- How Do We Get from Here to There?
  - Separate application development from resource management
    - Through an abstraction called the Virtual Grid
  - Provide tools to bridge the gap
    - Scheduling, resource management, distributed launch, simple programming models, fault tolerance, grid economies

VGrADS Big Ideas

- Virtualization of Resources
  - Application specifies required resources in Virtual Grid Definition language (vgDL)
    - Give me a tight bag of as many Opterons as possible
  - Virtual Grid Execution System (vgES) produces specific virtual grid matching specification
    - May involve adding layer of capabilities
    - vgES also reports properties of the virtual grid to application
- Generic In-Advance Scheduling of Application Workflows
  - Application includes performance models for all workflow nodes
    - Performance models automatically constructed
  - Software schedules applications onto virtual Grid, minimizing total makespan
    - Including both computation and data movement times

Virtual Grids (VGs)

- A Virtual Grid (VG) takes
  - Shared heterogeneous resources
  - Scalable information service
- and provides
  - A hierarchy of application-defined aggregations with constraints and rankings
  - E.g. Cluster of Opteron nodes
- Virtual Grid Execution System (vgES) implements VG
  - Brings order out of chaos
- Can be implemented on multiple platforms
  - Grids
  - Compute clouds
  - HPC
  - Manycore chips

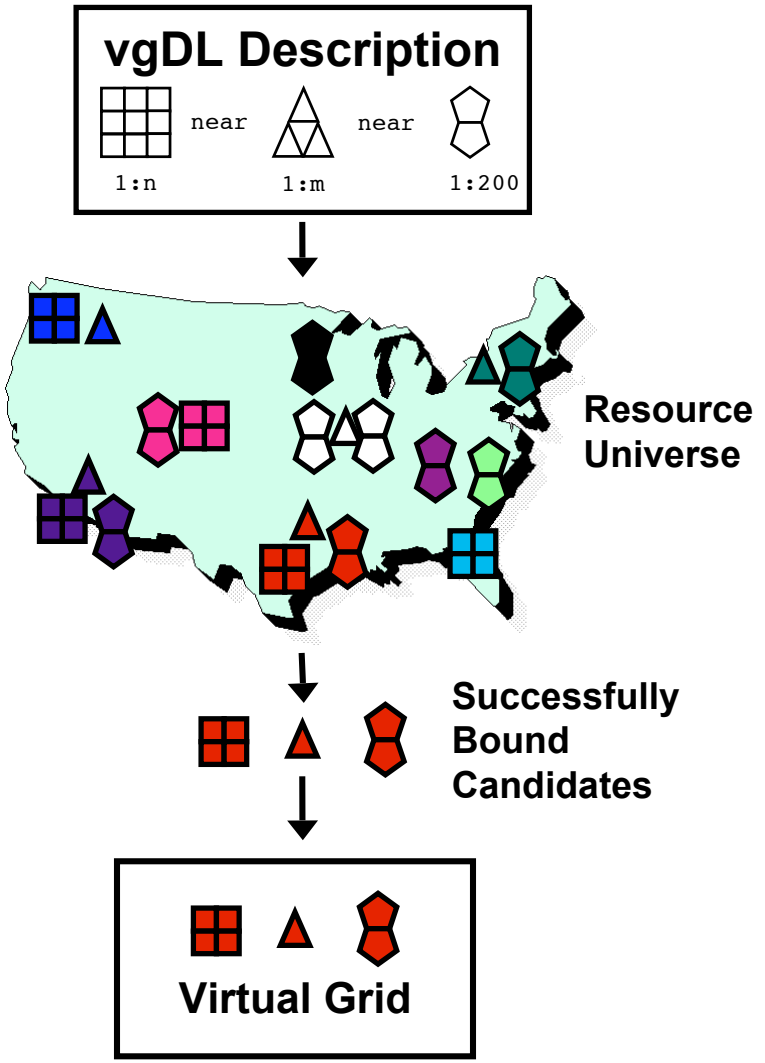

Virtual Grids (VGs)

- A Virtual Grid (VG) takes
  - Shared heterogeneous resources
  - Scalable information service
- and provides
  - A hierarchy of application-defined aggregations with constraints and rankings
  - E.g. Cluster of Opteron nodes
- Virtual Grid Execution System (vgES) implements VG
  - Brings order out of chaos
- Can be implemented on multiple platforms
  - Grids
  - Compute clouds
  - HPC
  - Manycore chips

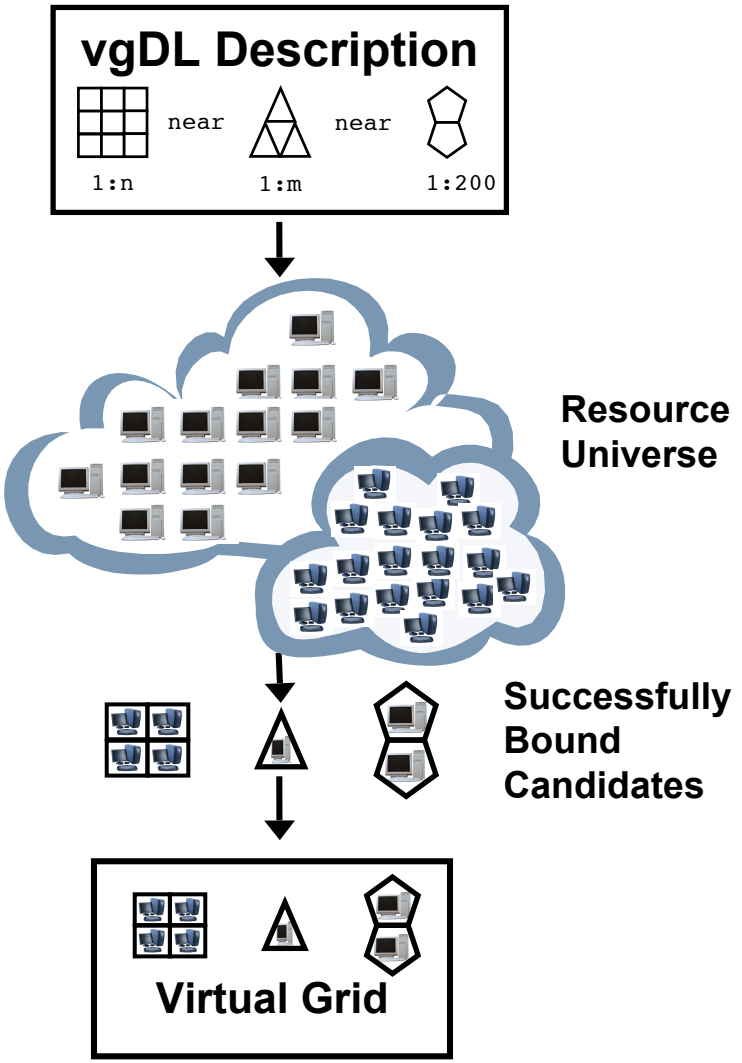

Virtual Grids (VGs)

- A Virtual Grid (VG) takes
  - Shared heterogeneous resources
  - Scalable information service
- and provides
  - A hierarchy of application-defined aggregations with constraints and rankings
  - E.g. Cluster of Opteron nodes
- Virtual Grid Execution System (vgES) implements VG
  - Brings order out of chaos
- Can be implemented on multiple platforms
  - Grids
  - Compute clouds
  - HPC
  - Manycore chips

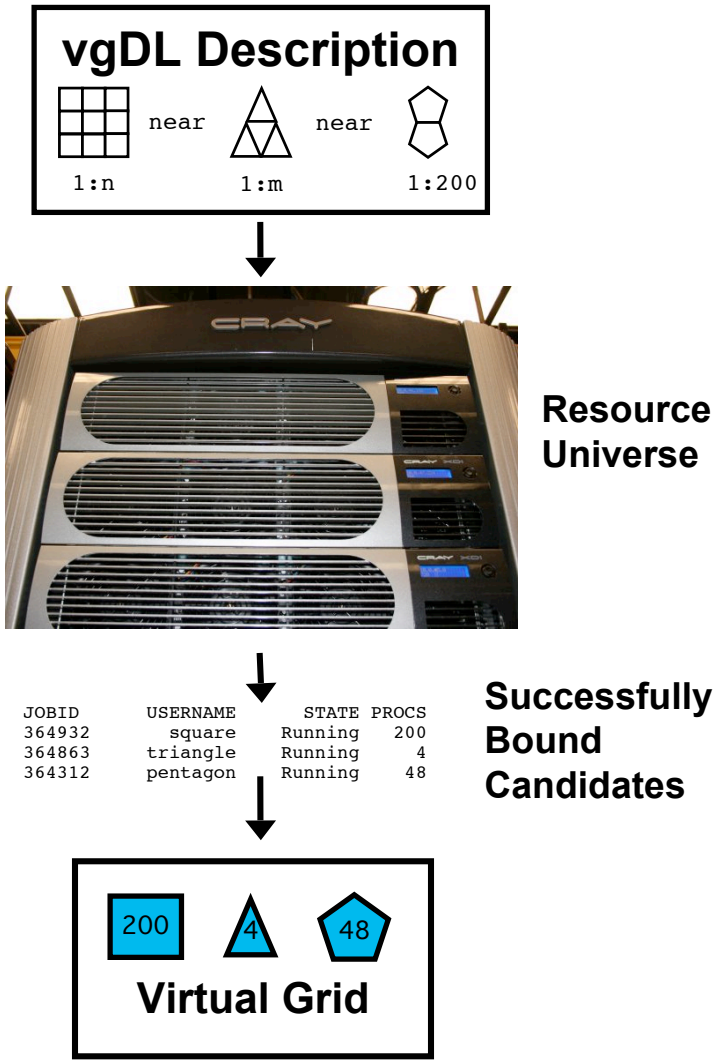

VGrADS
Virtual Grid Application Development Software Project

VGrADS Tool Research

- **Scheduling of workflow computations**
  - Off-line look-ahead scheduling dramatically improves total time
  - Accurate performance models significantly affect quality of scheduling
  - Batch queue behavior can be predicted accurately enough for scheduling decisions
- **Fault tolerance**
  - Diskless checkpointing for linear algebra computations (application-specific)
  - Qualitative performance analysis for fault detection
  - Optimal checkpoint frequency for iterative applications
  - Availability prediction

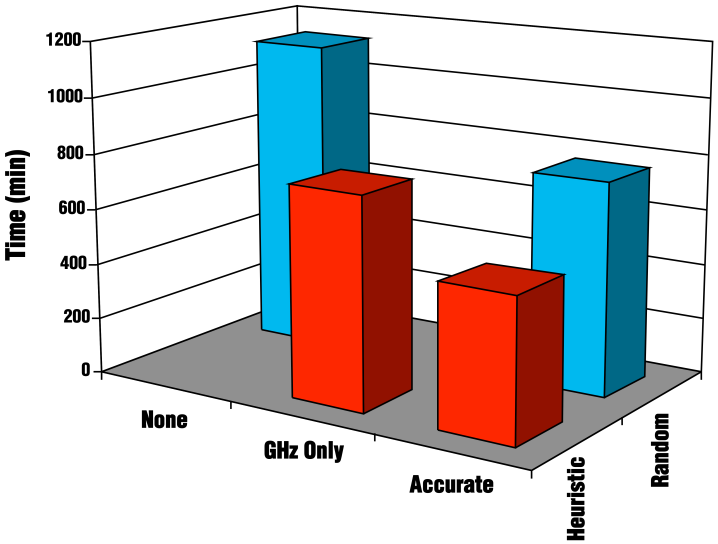

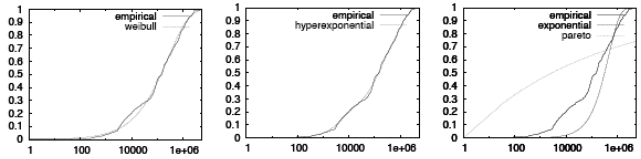

Fig. 1. CSIL data with Weibull fit

Fig. 2. CSIL data with hyperexponential fit

Fig. 3. CSIL data with exponential and Pareto fits

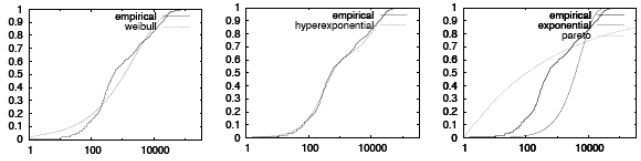

Fig. 4. Condor data with Wiebull fit

Fig. 5. Condor data with hyperexponential fit

Fig. 6. Condor data with exponential and Pareto fits

VGrADS
Virtual Grid Application Development Software Project

VGrADS: Application-driven Research

- SC04
  - Executing the EMAN application
    - Incorporated performance models
- SC05
  - Executing the EMAN application
    - Used batch queue predictions (and performance models)
- SC06
  - Executing the LEAD application
    - Scheduled to vgES virtualized resource reservation slots (leveraging queue predictions and performance models)
- SC07: *Fault tolerance through virtual grid abstractions*
  - Executing the LEAD application
    - Incorporates reliability information (and above features)
    - Combine replication and migration of workflow nodes

Overview of SC07 Demo

- Demonstrate fault-tolerance *on top of* vgES
  - Scheduler
  - Fault Tolerance & Recovery (FTR)
  - Workflow Planner (WP)
- Fault tolerance mechanisms
  - Scheduler, FTR, WP communicate during planning to decide *over-provisioning*
  - Scheduler, FTR, EM, vgES communicate during execution for *rescheduling*
  - FTR uses availability prediction service for reliability estimates
- vgES
  - Management of multiple copies
  - Passive and application agnostic
- LEAD – the driving application

The LEAD Vision: A Paradigm Shift

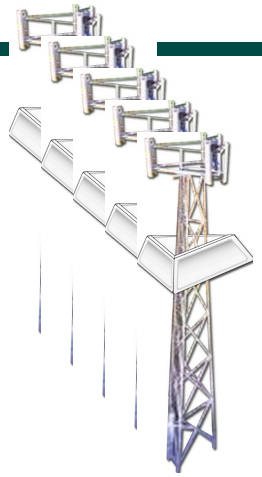

Dynamic Observations

Analysis/Assimilation

Quality Control
Retrieval of Unobserved Quantities
Creation of Gridded Fields

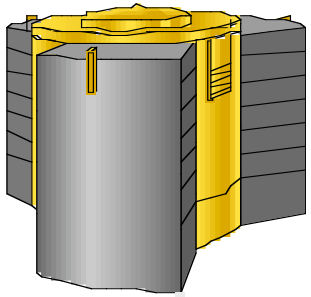

Prediction/Detection

PCs to Teraflop Systems

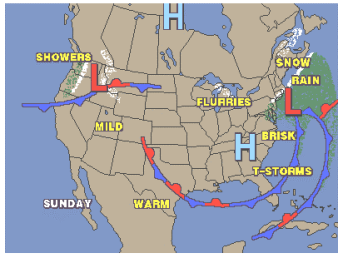

Product Generation, Display, Dissemination

Models and Algorithms Driving Sensors

End Users

NWS
Private Companies
Students

The CS challenge: Build cyberinfrastructure services that provide adaptability, scalability, availability, useability, and real-time response.

VGrADS
Virtual Grid Application Development Software Project

VGrADS at SC07

- Booth Talks and Demos
  - When and where
    - Tuesday, 4:00-5:00pm, *GCAS* booth (789)
    - Wednesday, 10:30-11:30am, RENCI booth (3215)
    - Wednesday, 3:00-4:00pm, SDSC booth (561)
- Other talks and posters
  - Poster: "Performability Modeling for Scheduling and Fault Tolerance Strategies for Grid Workflows" by Lavanya Ramakrishnan & Dan Reed
    - All Week (Reception, Tuesday 5:15-7:00pm), Ballroom Lobby
  - Paper: "Automatic Resource Specification Generation for Resource Selection" by Richard Y. Huang, Henri Casanova and Andrew A. Chien
    - Tuesday, 1:30-2:00pm, room A1/A6
  - Doctoral Showcase Presentation: "Qualitative Performance Analysis for Large-Scale Scientific Workflows" by Emma Buneci
    - Wednesday, 3:45-4:00pm, room A10/A11
  - Presentation: "A Framework for Qualitative Performance Analysis of Large-Scale Scientific Applications" by Emma Buneci and Dan Reed
    - Wednesday, 10:00-10:30, SC07 Education Booth (L1)