The Virtual Grid Application Development Software (VGrADS) Project

Overview

Ken Kennedy

Center for High Performance Software

Rice University

http://vgrads.rice.edu/site_visit/april_2005/slides/kennedy-overview

VGrADS

Virtual Grid Application Development Software Project

Supplementary Materials

- Packets
  - Agenda
  - Participant bios and pictures
  - Slides
  - Poster abstracts
  - Annual report for 2005
- Web Site
  - Linked agenda
  - Full participant bios
  - Electronic (full-page) slides
  - Publications
  - Annual reports
  - PDFs of Posters

Graphical Agenda

Progressive Revelation

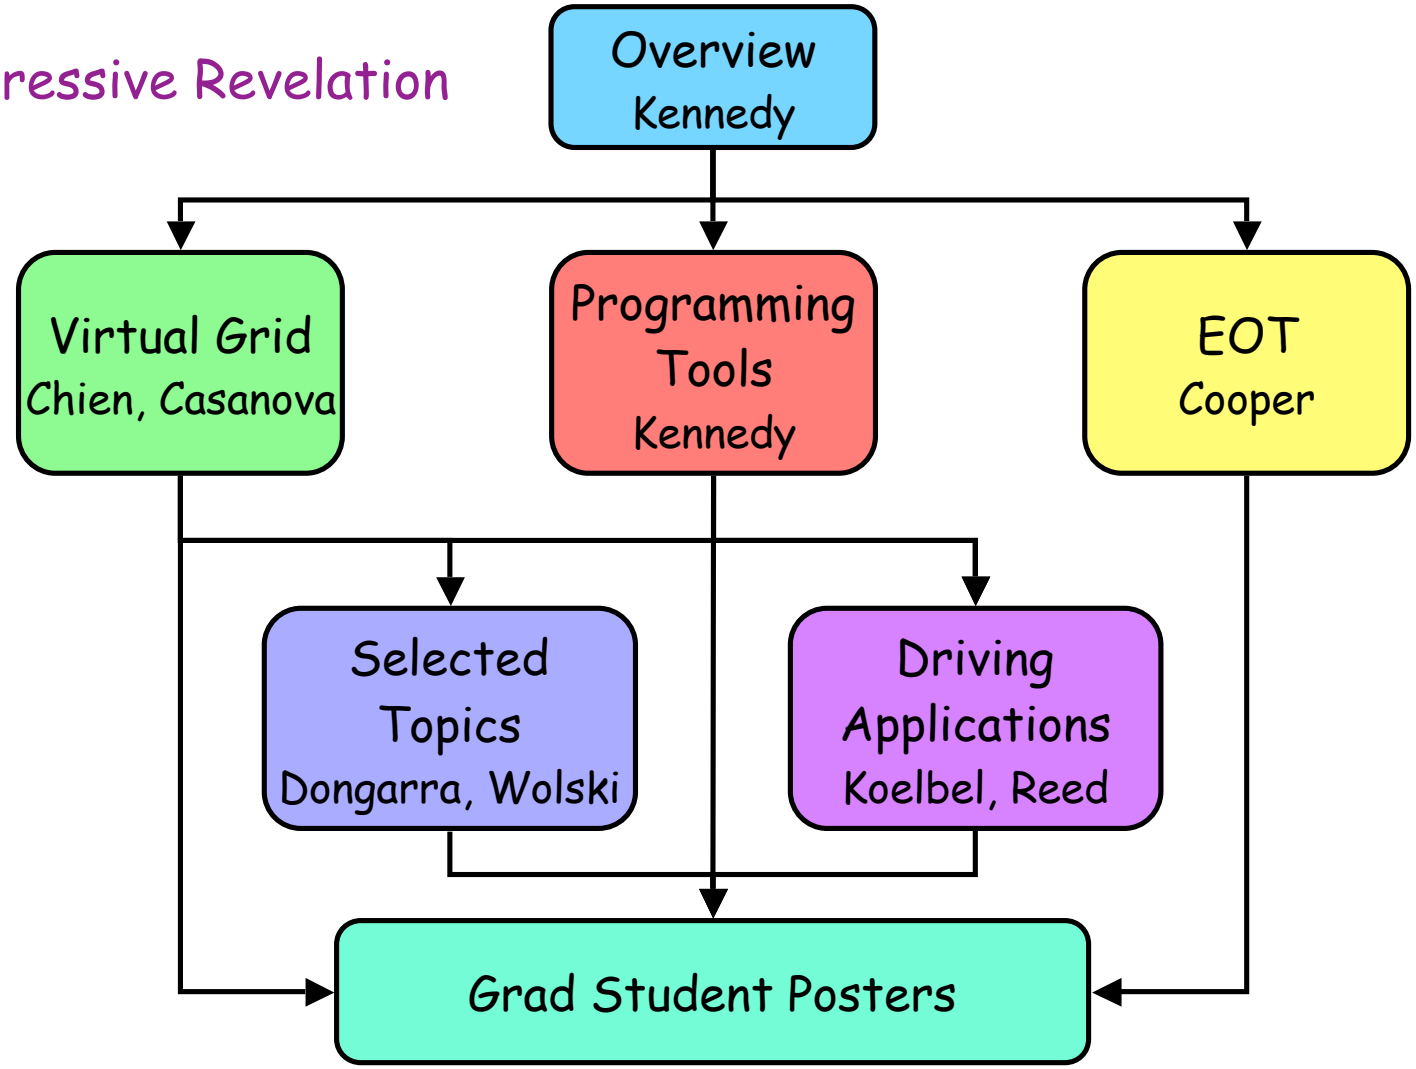

The VGrADS Team

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

Rich Wolski

Fran Berman

Andrew Chien

Henri Casanova

Keith Cooper

Ken Kennedy

Charles Koelbel

Richard Tapia

Linda Torczon

Jack Dongarra

Carl Kesselman

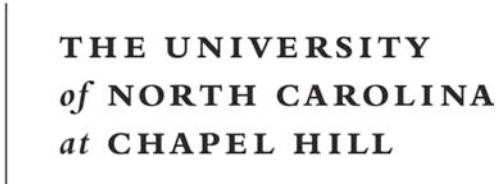

Dan Reed

Lennart Johnsson

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

VGrADS
Virtual Grid Application Development Software Project

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 development
- What Do We Need?
  - —A more abstract view of the Grid
    - Each developer sees a specialized "virtual grid"
  - —Simplified programming models built on the abstract view
    - Permit the application developer to focus on the problem

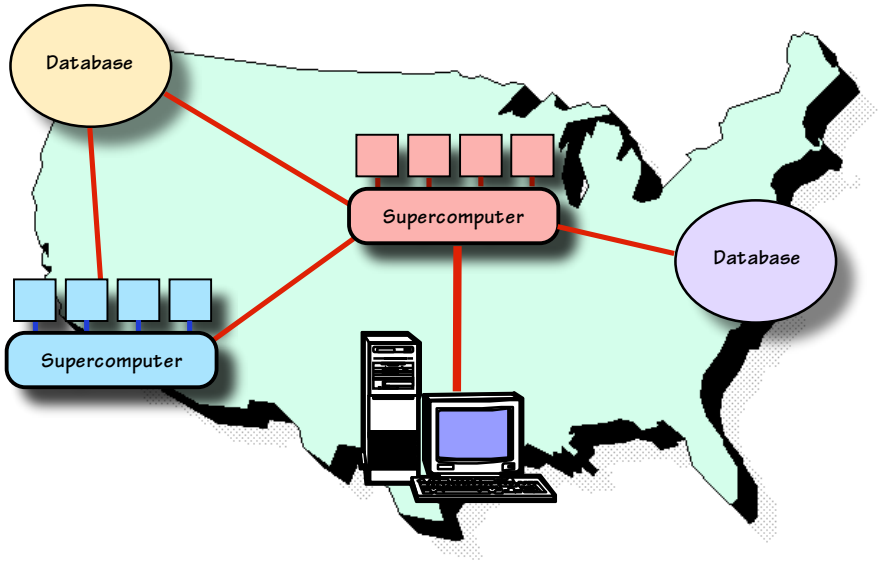

The Original GrADS Vision

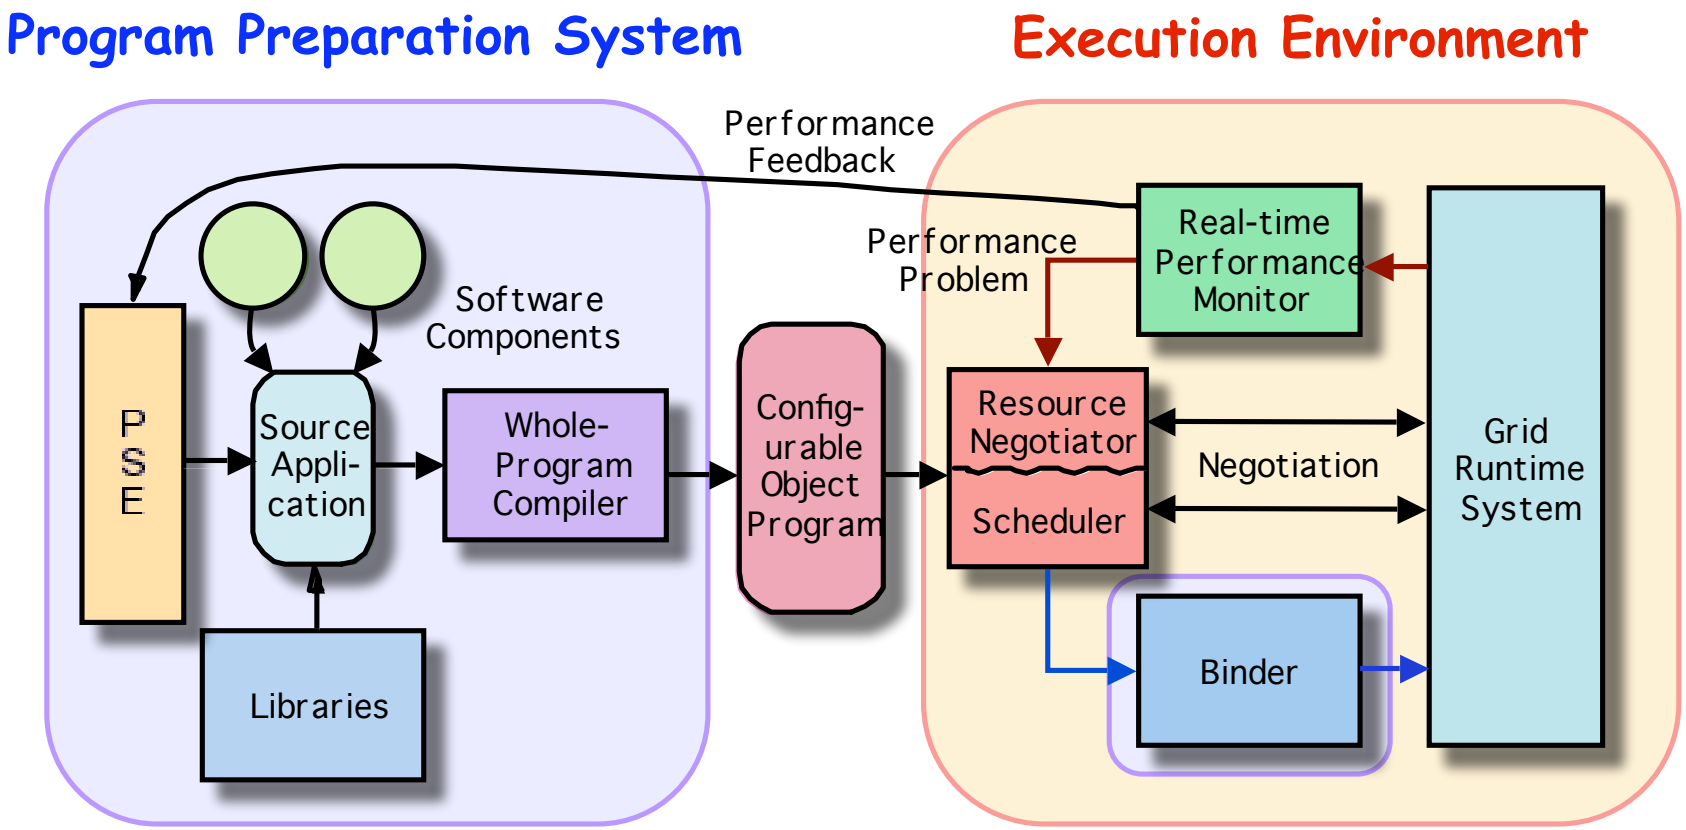

Lessons from GrADS

- Mapping and Scheduling for MPI Jobs is Hard
  - Although we were able to do some interesting experiments
- Performance Model Construction is Hard
  - Hybrid static/dynamic schemes are best
  - Difficult for application developers to do by hand
- Heterogeneity is Hard
  - We completely revised the launching mechanisms to support this
  - Good scheduling is critical
- Rescheduling/Migration is Hard
  - Requires application collaboration (generalized checkpointing)
  - Requires performance modeling to determine profitability
- Scaling to Large Grids is Hard
  - Scheduling becomes expensive

VGrADS Virtual Grid Hierarchy

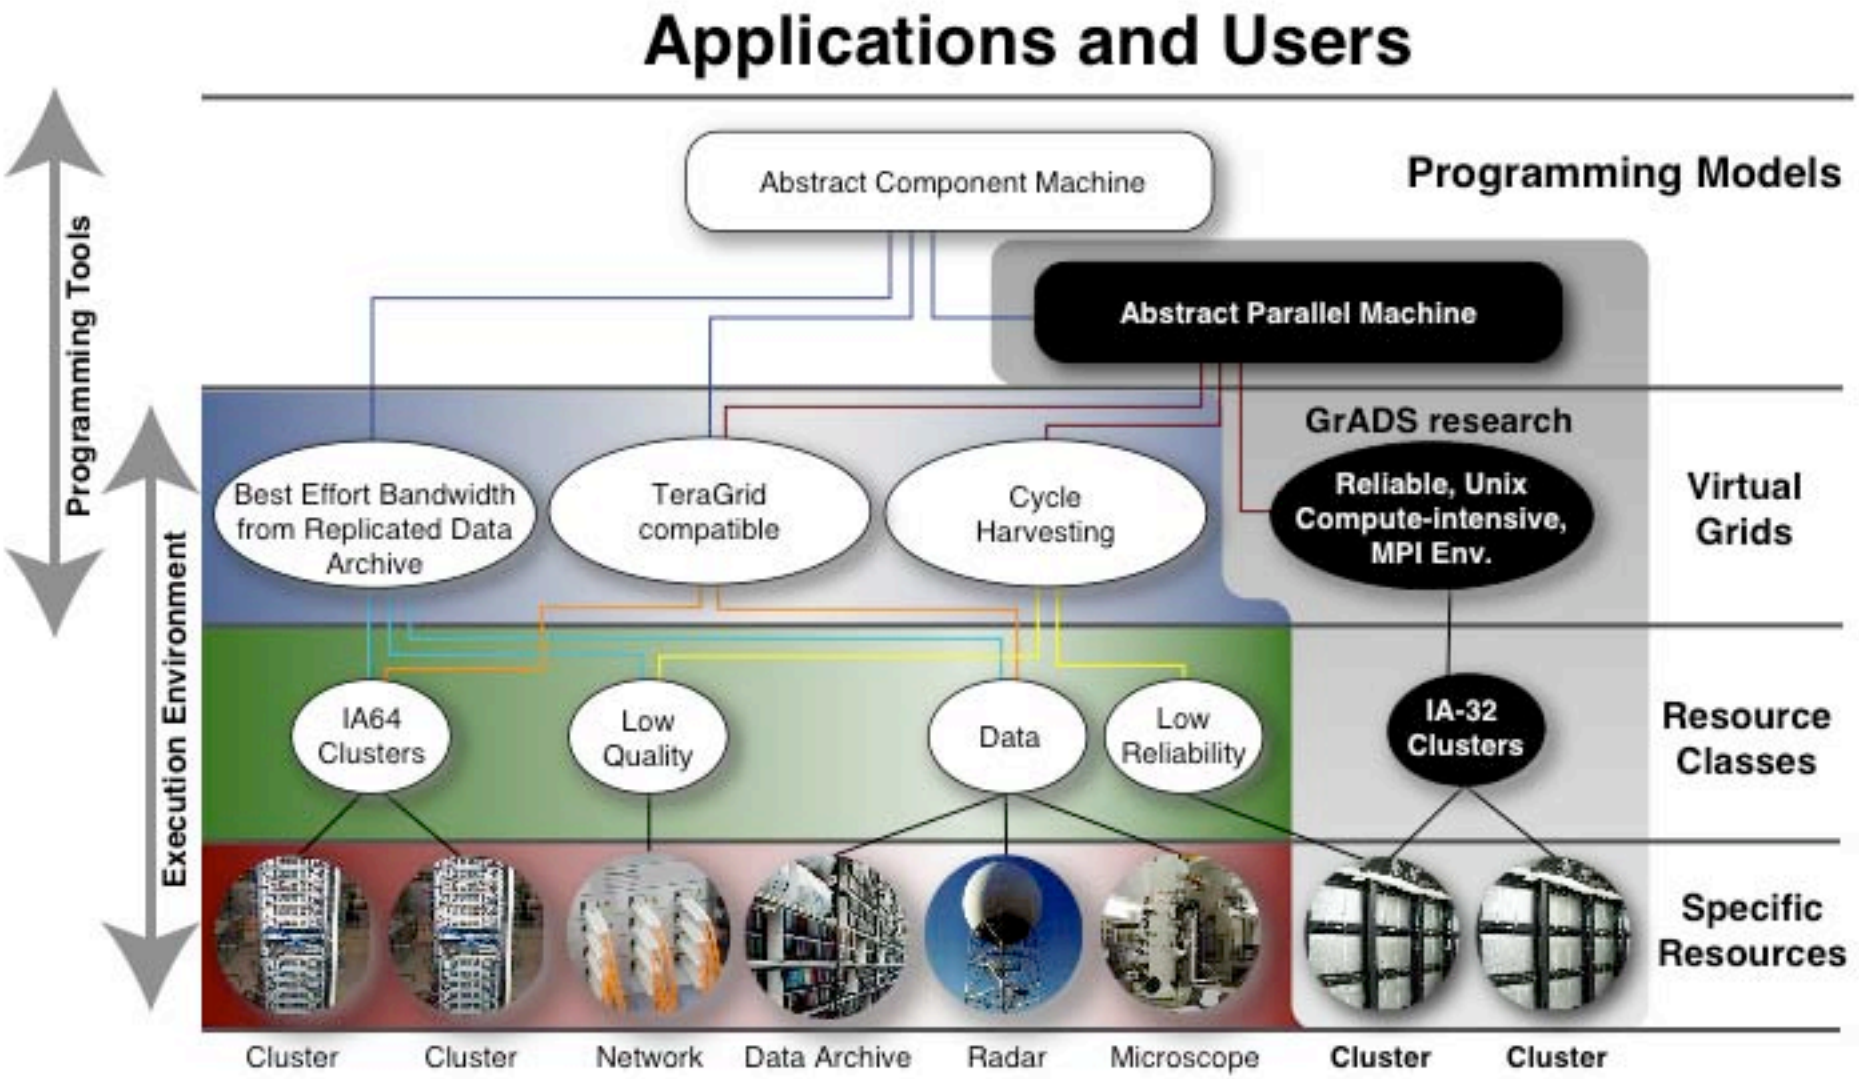

VGrADS
Virtual Grid Application Development Software Project

VGrADS Project Overview

- Virtual Grid Abstraction and Virtual Grid Execution System (vgES)
  - Application-level resource abstraction separates concerns
  - Permits true scalability and control of resources
  - Virtual Grid Execution System enables simple resource management
    - Supports fault tolerance, reasoning about app behavior, rescheduling
- Tools for Application Development
  - Abstract programming interfaces
  - Easy application scheduling, launch, and management
    - Workflow graphs and tightly-coupled computations
- Support for Fault Tolerance and Rescheduling/Migration
  - Collaboration between application and vgES
- Research Driven by Real Application Needs
  - EMAN, LEAD, GridSAT, EOL

Architecture: vgES implements Virtual Grid

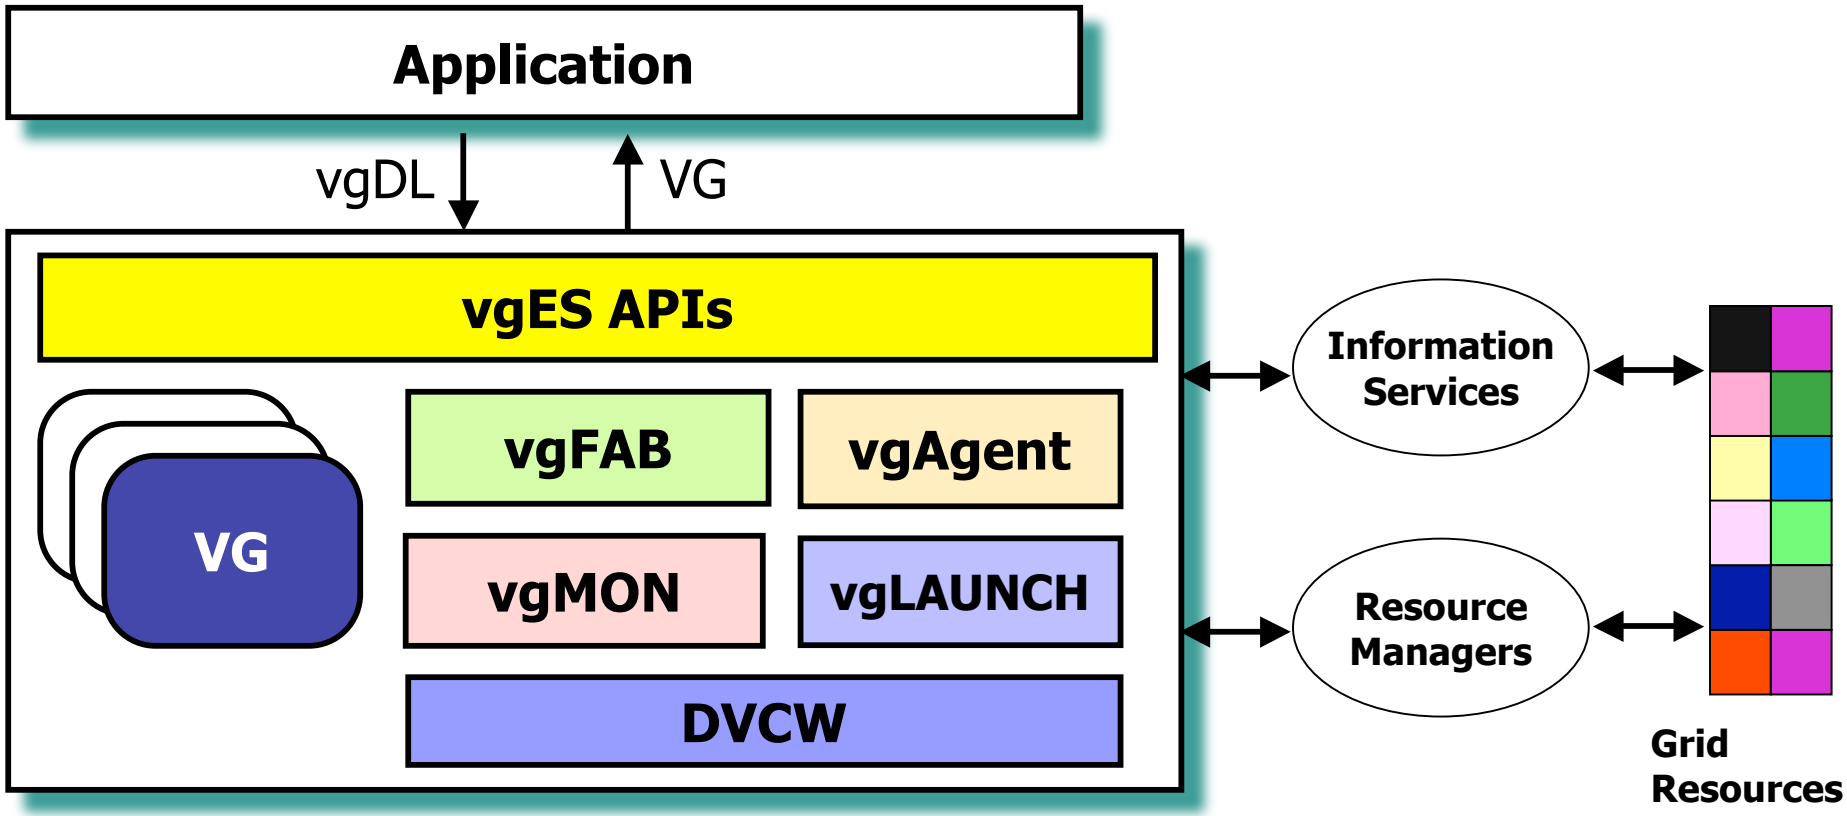

Virtual Grid Approach

- Separation of Concerns
  - Application Planning and Management
  - Complex Grid Resource Environment Mgmt
  - => vgDL and Virtual Grid
- Scalable Selection and Binding
  - Large Resource Pools
  - Competitive, Dynamic Environments
  - => Finding and Binding
- Application-Driven Resource Management
  - Application-Level Abstraction
  - Grid Information
  - => Virtual Grid Explicit Resource Abstraction

Programming Tools

Focus: Automating critical application-development steps:

- Initiating and managing application execution
  - Optimize and launch application on heterogeneous resources
  - Support for fault tolerance and rescheduling/migration
- Scheduling application workflows
  - Whole-workflow scheduling using performance models
- Constructing performance models
  - Automatically from application binaries
    - Cross-platform modeling
- Building workflow graphs from high-level scripts
  - Examples: Python (EMAN), OGRE (LEAD), Matlab

VGrADS Applications

Philosophy: Computer Science Research Driven by Applications

- EMAN
  - Contruction of 3-D models from 2-D electron micrographs
    - Talk by Charles Koelbel, Posters by Anirban Mandal and Bo Liu and by Gabriel Marin
- LEAD
  - Collection of applications for real-time weather prediction, sponsored by NSF ITR Large
    - Talk by Dan Reed, Poster by Emma Buneci
- GridSAT
  - Satisfiability on the Grid
- EOL (Encyclopedia of Life)
  - Reduced emphasis due to software issues

VGrADS Education

- Graduate Education
  - Graduate student support: 23 students
    - Rice: 7; UCSB: 1; UCSD: 5; UH: 2; UNC: 4; USC: 1; UTK: 3
  - Cross-institution graduate student exchange
  - Grid courses
    - UCSD (Chien) and UCSB (Wolski)
- Undergraduate Education
  - Internet and Grid Architectures Course
    - Under development
- Underrepresented Groups
  - Support for AGEP program
  - Activities at Tapia and Hopper Conferences
  - CS-CAMP

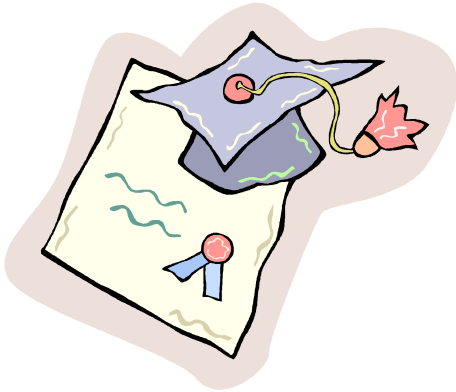

Outreach

- Web Site
  - http://vgrads.rice.edu/
- Conferences and Talks
  - CCGrid and HPDC papers
  - Invited addresses
- Application Collaborations
- Participation at SC Conferences
  - Demonstrations of two applications

VGrADS Demos at SC04

- **EMAN - Electron Microscopy Analysis [Rice, Houston]**
  - 3D reconstruction of particles from electron micrographs
  - Workflow scheduling and performance prediction to optimize mapping
- **GridSAT - Boolean Satisfiability [UCSB]**
  - Classic NP-complete problem useful in circuit design and verification
  - Performance-based dynamic resource allocation and scheduling

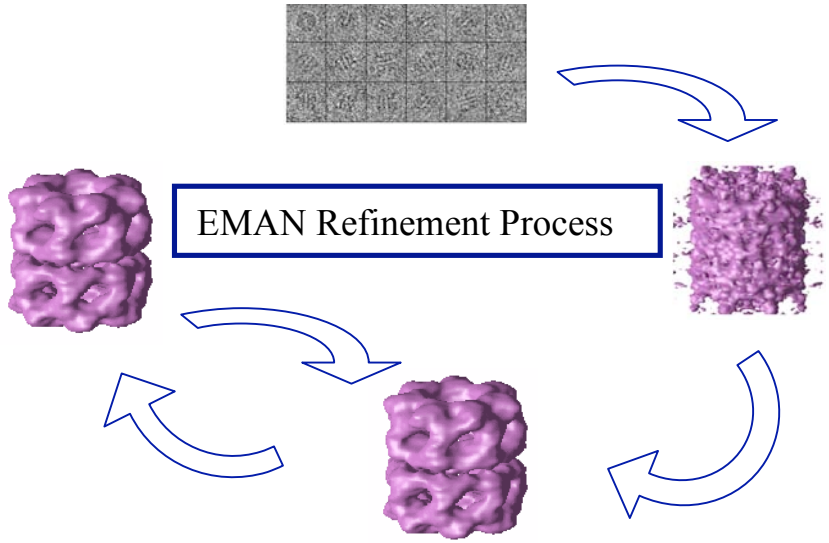

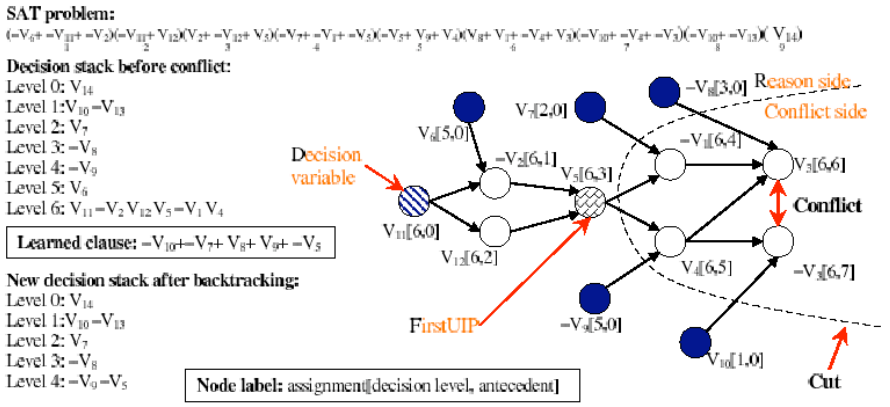

Figure 1: Example of conflict analysis with learning and non-chronological backtracking

VGrADS
Virtual Grid Application Development Software Project

Leverage

- NMI
  - Community infrastructure that can be counted on
  - Vehicle for deployment of successful research
  - Outlet for technology
- State of Texas Support
  - LEARN:
    - Support for statewide networking
    - 33 institutions $7.5 million over two years
  - TIGRE
    - Application driven software stack
    - $2.5 million for two years shared by 5 institutions
    - Rice and UH both participants
- Houston Research and Education Network
  - Rice, UH, Texas Med Center, NLR (lit in May)
    - Abovenet contribution $3.4 million

Management

- Value of Virtual Organization
  - —No one site could take on topic of this breadth
  - —Project vision requires an integrated approach
    - Thinking about broad range of issues from the beginning
- Management
  - —Executive Committee
  - —Technical Working Group
- Strategy for Collaboration
  - —Workshops, Telecons
    - Overcome project management challenges through interactions
  - —Grad student exchange

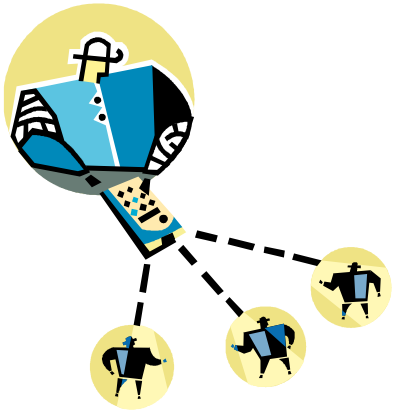

Research vs Development

- Original Proposal
  - Research
  - Significant investment in software infrastructure development
  - Application collaborations (funded staff members)
- Revised Statement of Work
  - Focus on research (guidance from NSF)
- Situation Today
  - To conduct research we are building prototype software
  - Leveraging NMI and Texas
  - Application collaborations still drive research

Annual Budget

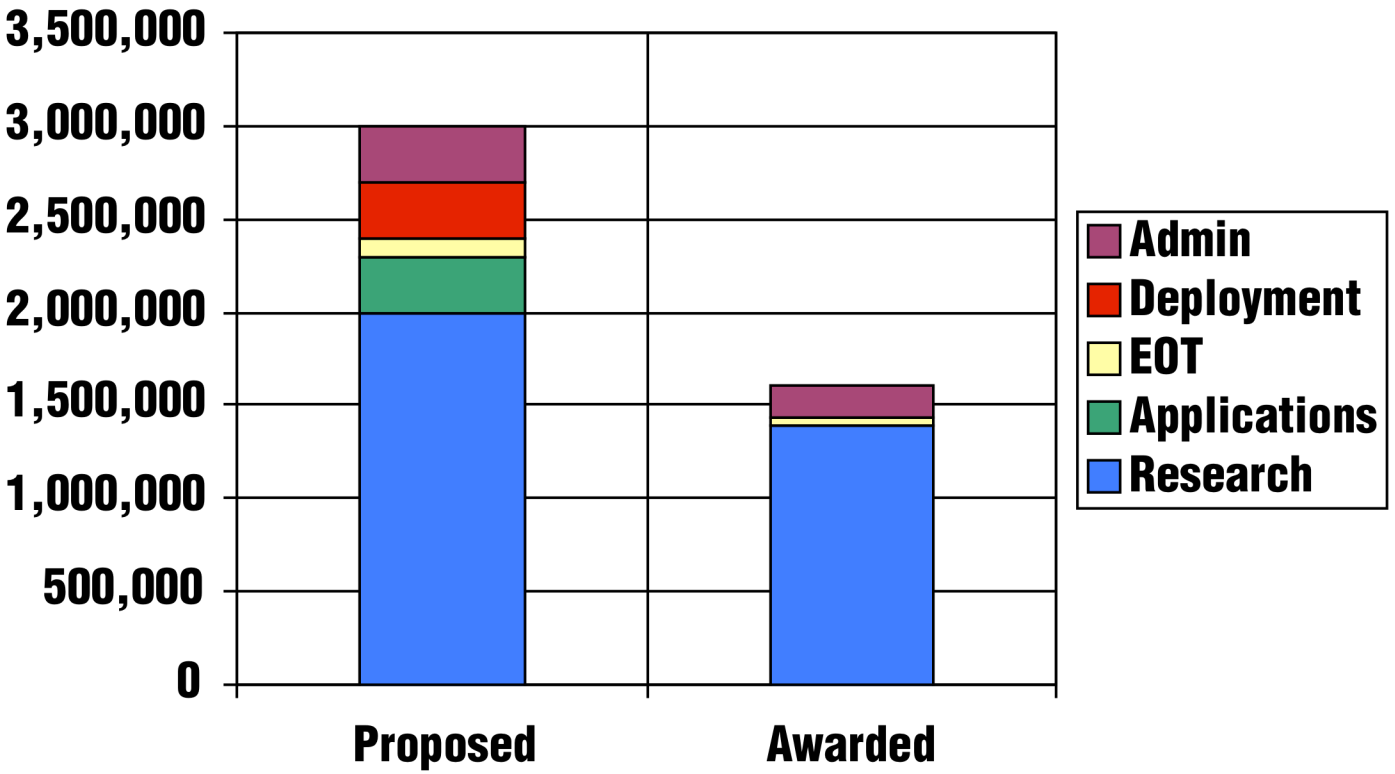

VGrADS Summary: A Holistic Approach

- What justifies a Large ITR?
  - Community: no one institution covers everything
  - Unified project vision
  - Shared infrastructure
  - Integration: would not happen without a unified project

- VGrADS
  - Built on GrADS insights and experiences
    - community of leading researchers who work together effectively
    - broad coverage of requisite topics
  - Vision for extremely simple application development interface
    - grid virtualization to hide complexity
  - Shared software stack and testbed
    - vgES toolkit, policies and application drivers
  - Many interrelated layers require and benefit from the integrated effort
    - program tools, provisioning, scheduling, measurement
    - prediction, fault tolerance, infrastructure, applications

Annual Budget by Institution