Answers to Questions

CGrADS Response Session

CGrADS PIs

http://hipersoft.rice.edu/stc_site_visit/talks/answers.pdf

CGrADS: The Critical Science

- CGrADS scientific research theme
  - Exploration of formalisms and strategies for adaptation in dynamic environments
- Ten year challenges
  - Characterizing emergent Grid dynamics
  - Developing effective compilation techniques for dynamic environments
  - Determining bounds on achievable efficiency and performance
  - Establishing stability and adaptation constraints
  - Creating high-level application development methodologies based on component composition
  - Validating formalisms via empirical investigation
- CGrADS would be uniquely positioned to help focus the community on the exploration of long-term Grid research foundations
  - Leverages Grid infrastructure developed by other projects
  - Near term implementation efforts enable investigation of long term scientific issues

Why an S&T Center?

- Comprehensive, long-term, integrated research effort
  - Many researchers from different institutions and different academic backgrounds needed to address the problem
  - Focus will be needed to drive toward a common goal
- Enables large-scale experimental Grid Research
  - Can't be done without a group effort
- Multidisciplinary nature of the research
  - We must be able to pursue new approaches as they emerge
- Critical mass and visibility
  - To foster education and outreach programs
  - Catalyst for the community

(1) ERT Timeline and Success Measure

- What is your measure of success for undergraduate research training?
  - Numbers of students from under-represented groups who get degree in the S&T discipline
  - Numbers of students from under-represented groups who go to graduate school in the S&T discipline
  - Numbers of students from under-represented groups who take jobs in industry in the S&T discipline
  - Numbers of students who have gone through the general education courses at all institutions
    - Course evaluations
    - Number of course adoptions at institutions within and outside of CGrADS

Goals For Success of Graduate Programs

- Goals for 4 years:
  - Undertake programs that will double the numbers of graduate students from underrepresented groups enrolled in graduate programs in the participating departments
    - Current: 54/1390
  - Increase numbers of women enrolled by 50%
    - Current 318/1390
  - Maintain US citizens and permanent residents percentage at 50% or better
  - Increase percentage of currently enrolled students who complete Ph.D.
    - As compared to the four years preceding CGrADS start
- Strategies:
  - Focus on the undergraduate minority students in current institutions with the goal of getting them to GS in other CGrADS institutions
  - Build a community along the lines of SaS/AGEP
  - Develop CGrADS EHR Leadership Committee to develop and manage graduate support groups at all institutions

Timeline: Graduate

- Year 1
  - Establish Grid projects course
  - Initiate cross-institutional visits
  - Continue Graduate support groups at Rice and instantiate at one other institution
  - Form CGrADS EHR Leadership Committee
- Year 2
  - Design Grid programming course
  - Export Grad/Undergrad community-building groups to two other institutions
- Year 3
  - Teach Grid programming course at one or more institutions
  - Continue export of graduate student support groups and cross-institutional visits

Timeline: Undergraduate

- Year 1
  - Initiate undergrad research projects at Rice with focus on Grid-related research (Internet, systems, compilers)
  - Teach version 1 of general education course
  - Restart SC-COSMIC
- Year 2
  - Expand undergrad research programs to one or two other sites
  - Teach version 2 of the general ed course with detailed notes prepared for export
  - Revitalize SC-COSMIC and plan curriculum export programs
- Year 3
  - Continue to expand undergraduate research efforts
  - Export general ed course to several other institutions, including at least one minority-serving institution (TSU)

Timeline: K-12

- Year 1
  - Design of expanded TeacherTECH program with materials from Information Technology Architectures
  - Initiate program for parents with TTOI & TeacherTECH
  - Plan collaboration with FIRST
- Year 2
  - First offering of revised TeacherTECH program
  - Continue TTOI collaboration on parent programs
  - Initiate FIRST collaboration
- Year 3
  - Expand revised TeacherTECH program
  - Continue and review other programs
    - Evaluation collaboration with HISD

(2) Dependence on Other Projects

- **Are there projects not in your center on which your success depends, such as Condor-G?**
  - No.
    - We do depend on some projects that we control, notably Globus and NWS
    - We do plan to leverage external projects when and where appropriate and will consider stability, longevity and willingness to collaborate of project personnel in project selection
  - External projects leveraged by CGrADS but on which our success is *not* dependent may include
    - Condor
    - GriPhyN
    - APST (AppLeS Parameter Sweep Template)
    - BIRN (Biomedical Imaging Research Network)
    - TeraGrid
    - Etc.

(3) Intended IP Policy

- Our goal is to ensure the broadest possible availability of the software developed by the project
- We recognize the importance of liberal licenses and centralized ownership to potential industrial partners and adopters
- To this end, we intend to adopt a common, liberal open source licensing policy for the core software developed within CGrADS
  - A FreeBSD-like license, as used e.g. for Globus and (Sca)LAPACK
- The establishment of this policy will require negotiation with the 8 institutions involved in CGrADS; this may not be easy, but we are committed to pursuing it
- We note that U.Chicago, NCSA, and USC/ISI all have approved open source policies of this form already
- We will consult with our industrial council as we move forward in this area

(4) What value does CGrADS add to NMI, DTF, and PACI?

- NMI, DTF, PACI are integration, deployment & support projects
  - NMI identifies best of breed and deploys and supports
  - DTF deploys current Grid software on TeraGrid facilities
  - PACI supports only hardening and deployment
- Only CGrADS will advance our understanding of the Grid
  - What are the basic methodologies for developing Grid applications?
    - How do these deal with dynamics in Grid environment?
    - What types of development environments are needed?
  - Prototype development of technology
  - Initiate technology transfer (perhaps to NMI, DTF or PACI)
- Value added to long-term future of NMI, DTF, and PACI:
  - Establish research foundations
  - Explore next generation Grid tools and technologies
  - Increase usability, efficiency, and performance

(4) What are the CGrADS Deliverables (Years 1-3)

- Planning Documents, Research Papers, Graduates
- Research Results/Software Prototypes
  - Prototype runtime binder (K. Cooper, L. Torczon)
  - Prototype domain specific language generation strategies (K. Kennedy, K. Cooper, J. Mellor-Crummey)
  - Automatic performance model generator (J. Mellor-Crummey)
  - Scalable network simulator, methodologies for performance extrapolation (A. Chien)
  - Economic models for resource allocation (R. Wolski)
  - Intelligent performance monitoring and contract specification (D. Reed)
  - Scheduling strategies for new application classes (F. Berman)
  - Configurable information services to support dynamic contract monitoring (C. Kesselman, I. Foster)
  - Tools for dynamic testbed config, monitoring (C. Kesselman, R. Wolski)
  - Grid-aware adaptive library framework (J. Dongarra, L. Johnsson)
  - Methods for automatic resource selection, prototype tool (I. Foster)

(6) Future Directions

- Some aspects seem quite mature (e.g., ScaLAPACK). What will the relevant researchers work on in the out years?
- In all CGrADS areas, we build on some existing technologies but also have an aggressive, long-term research agenda
- In the case of numerical libraries, in particular:
  - The design of smart libraries; libraries that can analyze the data and search the space of solution strategies to make optimal choices
  - The development of agent-based methods for solving large numerical problems on both local and national grids
  - The design of a telescoping language framework for expressing the "software component architecture" of grid applications to make development easier while resolving multi-language/multi-library interface issues
  - Development of a prototype framework based on standard components for building and executing composite applications

(7) Currently Targeted Application Milestones

- Three codes concurrently, each engaged for two-three years
- 2002
  - Cactus: traditional PDE solver, aggressive application scenarios
  - CAPS: dynamic data acquisition and real-time data ingest
  - ChemEng Workbench: application service scenarios prototypes
- 2003
  - Cactus: by now transitioned to operational use by application group
  - CAPS: adaptive execution for high-speed prediction
  - ChemEng Workbench: application service scenarios operational
  - CMS/GriPhyN: query estimation and dynamic scheduling
  - BIRN-like distributed bioscience: emergent behavior issues
- 2004
  - CAPS: by now transitioned to operational use by application group
  - CMS/GriPhyN: large-scale experimentation in production settings
  - NEES: application service and real-time data analysis scenarios

Application Milestones

- We began GrADS by leveraging internal application expertise
  - ScaLAPACK due to its relative simplicity & internal domain expertise
  - Collaboration with the Cactus group, an aggressive early adopter with strong commitment and vision in Grid area
- Within CGrADS, we expand the application base to include a broader range of application domains and usage scenarios

Application Milestones

- Chosen applications are exemplars of Grid behavioral domains
  - Distributed application services
  - Heterogeneous component composition
  - Computation and data management
  - Commercial and research codes

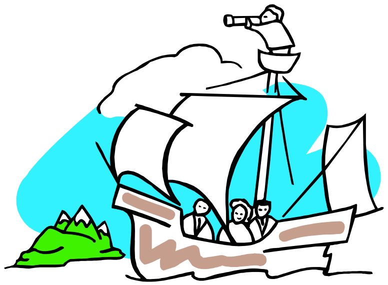

- We will manage applications via "research expeditions"
  - ScaLAPACK and Cactus were the two chosen for GrADS
  - Only a small number of codes can be managed concurrently
    - Bounded to maintain intellectual focus and coordinated activity
    - Chosen on basis of intellectual challenge <u>and</u> engagement
    - CGrADS will work with 2-3 codes concurrently

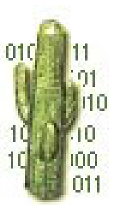

- Cactus, an example
  - Initially, component partner collaborations with Cactus
  - Later, integrated prototype coordination

Application Milestones: Cactus

- Automatic configuration of unigrid Cactus configurations on heterogeneous collections of uniprocessors and/or clusters
- Automatic configuration of AMR Cactus configurations on heterogeneous collections of uniprocessors and/or clusters
  - Using Globus/NWS resource characterization and scheduling
  - Demonstrate robust performance on range of system configurations
- Automated dynamic resource discovery, acquisition, and migration across Grid resources
  - Using resource selector, application manager, contract violation detection, rescheduling models
  - Demonstrate efficient and robust migration

Application Milestones: ChemEng Workbench

- Y1: Application service scenarios prototype
  - Negotiation of performance contracts with users
  - Dynamic scheduling and resource acquisition based on performance contracts, using contract monitoring, Globus real-time scheduling
- Y2: Application service scenarios operational
  - Production deployment of ChemEng workbench application server
  - Delivery of application service tools
  - Integration with commercial application service technologies

Application Milestones: CAPS

- Y1: Dynamic data acquisition and real-time data ingest
  - Establishment and monitoring of performance contracts to meet real-time data ingest requirements
  - Resource monitoring and prediction
- Y2: Adaptive execution for high-speed prediction
  - Dynamic application configuration and resource acquisition to optimize high-speed, data-intensive mesoscale prediction capabilities
  - Configurable object programs
  - Composable performance contracts for pipelines
  - Delivery of real-time data-driven application tools
- Y3: Transitioned to production deployment

Application Milestones: CMS, NEES

- CMS/GriPhyN
  - Y1: Query estimation and dynamic scheduling
    - Apply program preparation, performance modeling, contract monitoring techniques
    - Use to investigate application-specific dynamic scheduling
  - Y2: Large-scale experimentation in production settings
    - Application experiments on thousands of processors and dozens of sites
- NEES (Y3)
  - Application service scenarios
    - Demonstrate ability to schedule uni- and multi-processor NEES computations onto available Grid resources
  - Real-time instrument coupling scenarios
    - Demonstrate ability to perform robust real-time coupling of data analysis and simulation components with NEES instruments

(8) Benefit to Students of Long-Term Exchanges

- What is the benefit to the students of long-term exchange of graduate students?
  - By "long-term," we intend a visit of several months – say a quarter or a summer, or more if appropriate. During such visits, a student will
    - collaborate closely with researchers and students other than those at their home institutions
    - experience another research culture
    - build a research network of their own
  - The benefits of these visits are
    - Students will develop their own research collaborations and contacts
    - Students will experience a fresh perspective and alternative approaches to research
    - Students will develop greater professional maturity
- Note: Students on the GrADS project have developed a strong collaborative culture and professional network by routinely visiting other GrADS groups to discuss their work

(9) How Will EHR Programs be Exported?

- To CGrADS Institutions:
  - Leadership Committee with a representative from each site will meet on a regular basis to discuss export
  - Course export will be via packages of materials
  - Support programs and K-12 will require local coordinator
  - Graduate projects courses will be collaborations over the Internet
- To Other Institutions Nationally
  - Via existing PACI mechanisms and collaborations
    - PACI EOT Leadership Team

(10) Application Scientists in EAC

- Do you intend to include applications scientists in your external advisory committee?
  - Yes, we believe that applications scientists will provide valuable advice and perspective to the proposed research activities
  - We will include several distinguished computational leaders on the External Advisory Committee. Bill McCurdy, Bill Tang, Tom Jordan, Warren Washington, and Paul Woodward are exemplars of the type of application scientist we plan to include.

(11) Participation of Administrators

- Why are there no administrators from non-local institutions participating in the site visit?
  - Administrators from the Lead Institution (Rice) and one of the partner institution are participating in the site visit. The Rice Provost is serving as a proxy for the administrators of the other institutions and will consult with administrators from the other institutions as the need arises.
  - Letters of support have been received from the administrators of each institution, and were included as part of the proposal.
  - We requested guidance on whether participation of administrators from all partner institutions was necessary, and they indicated that it was our decision. Participation of the Rice administrators was required.
  - Because of the number of institutions involved (8), we decided that a teleconference with so many sites would provide little opportunity for interaction with each, and thus be counterproductive.

Risk Mitigation

- Definition of success
  - Insight into the deep questions defined earlier
  - From this insight, useful systems for programming Grids
- What are risks?
  - We fail to synthesize insights from the body of experience
    - Center structure guards against this by enabling long-term, coordinated investigation of deep issues
  - Failure to impact the community
    - Center structure guards against this by enlisting the community in building a shared vision
- General observations
  - We see no single point of failure for our research approach
  - This is a research project and failures themselves provide insights and lead to new research approaches
- Strategies
  - We have a portfolio of shorter- and longer-term research activities

How Do We Achieve User Buy-in?

- User communities are already committed to Grid computing concepts
  - Many funded to develop Grid applications
- Globus experience: focus on delivery of simple, modular components
- Leverage strong links with existing user communities
- Build tools that leverage and interoperate with standard infrastructure and tools
- Demonstrate significant added value: e.g., ease of development, performance, ease of modification, …
- Promote development of standards

CGrADS: The Critical Science

- CGrADS scientific research theme
  - Exploration of formalisms and strategies for adaptation in dynamic environments
- Ten year challenges
  - Characterizing emergent Grid dynamics
  - Developing effective compilation techniques for dynamic environments
  - Determining bounds on achievable efficiency and performance
  - Establishing stability and adaptation constraints
  - Creating high-level application development methodologies based on component composition
  - Validating formalisms via empirical investigation
- CGrADS would be uniquely positioned to help focus the community on the exploration of long-term Grid research foundations
  - Leverages Grid infrastructure developed by other projects
  - Near term implementation efforts enable investigation of long term scientific issues

Science!

- 3-5 year problems
  - Scalable simulation techniques
  - Adaptive runtime scheduling algorithms
  - Effective policies for federated resource control
  - Formulation of a stable resource economy supporting adaptation
  - Efficiency metrics for dynamic application execution
  - Compilation techniques for generic domain-specific languages
- 5-10 year problems
  - Compilation, algorithms, and application techniques for dynamic systems
  - Unifying intellectual framework for managing adaptation in dynamic environments
  - Rigorous understanding of overall Grid performance dynamics
  - Comprehensive measurement theory for Grid performance evaluation