Response to Reviews of the CGRADS STC Proposal

While we are gratified by the generally positive nature of the reviewers' comments, we would like to address several issues raised in those reviews.

1. Software Development

The reviewers expressed concern about the quality of software that must be produced by the proposed STC if the effort is to succeed. This sentence from the panel summary is representative:

> "It will be difficult to produce 'industrial strength' software which will be useful for applications on the grid. The previous STC produced a lot of excellent research but not very many useful products that computational scientists and engineers could use."

It is clear to us that we cannot succeed unless the CGrADS research leads to software that makes the Grid easier to use. Over our careers, we have generally accomplished this goal through two strategies. First, we have produced software that was widely distributed and used. LAPACK is a good example of this. On the other end of the spectrum, we have worked with vendors to produce commercial software that embodies research ideas. High Performance Fortran (HPF) is an example. Although HPF has met with mixed success in the marketplace, more than 20 vendors have offered HPF compilers. All of the commercial HPF compilers, and most parallelizing compilers, are based in part on research from the Center for Research on Parallel Computation. A second example of this approach is MPI, whose development was initiated by the CRPC.

We anticipate producing three classes of software during the lifetime of the Center: (1) high-quality "research-grade" software that will provide the basic infrastructure for our experiments and will be distributed externally; (2) experimental software for internal use to support specific experiments and investigations; and (3) prototype or throwaway software for use by a small number of investigators.

The first category, high-quality research-grade software, is critical to the Center's success. Therefore, the key question is: how will we produce it? Our approach is based on two principles:

- *Use of professional programming staff:* We cannot expect students to produce or maintain high-quality software. While students may be more than capable, the support of software for even a moderate-size user community is inconsistent with our educational goals for students.

- *Use of best-practice software engineering practices:* Although we will not be producing "industrial-quality" software, we will follow practices used in many industrial software development activities. These include: coding standards, well-defined software specifications, creation of test suites, regular build and test cycles, use of commercial software testing tools (such as Purify and Insure), and use of revision control systems (such as CVS).

This approach has been used by the PIs to produce a number of software packages that have seen widespread distribution and use (PVM, HPVM, NWS, PCN, LAPACK, ScaLAPACK, CMSSL, CMRTS, Nexus, Pablo, ATLAS, and HPCview). The Globus system, which was also developed with this methodology, is currently in use at hundreds of installations, including installations at NASA and the DOE.

Given the interest in the Grid shown by the national and international scientific community, we have an opportunity to work with that group to deploy more robust implementations of the research software produced by CGrADS. One way to drive this activity is through our collaboration with the NSF PACIs, which have committed to Grid computing for the proposed Distributed Terascale Facility (DTF).

2. Time Commitments

Several of the reviews expressed concern about the perceived commitment level of the investigators. The following quote from the summary document is representative:

> "The PIs are already involved with several other large projects, and there is some skepticism that they will be able to invest much of their time and energy into this Center."

In response, we observe that the proposed CGrADS effort lies at the heart of the research agenda of each of the research directors, most of whose current and proposed projects relate to the Grid in some way. We believe that these projects cannot succeed in the long term unless the CGrADS effort succeeds. Indeed, CGrADS can be viewed as an integral component of an R&D pipeline—it will conduct basic research needed by the next generation of the investigators' deployment (and research) projects. Thus, CGrADS will represent a critical portion of the investigators' intellectual effort, in spite of the fact that it is only a part of their total research support.

As evidence of this, note that the investigators are already collaborating and meeting regularly, running experiments across sites, writing papers together, and jointly organizing workshops, symposia and conferences, despite their busy schedules. The CGrADS project will extend and enhance their collaboration, rather than add an unrelated project to their plates.

Several of the reviews commented on the success of management strategies employed in the CRPC, a previous STC led by the Project Director. The CRPC involved a similarly busy group of researchers at seven different institutions. These management strategies have been extended and refined in the preliminary GrADS research sponsored by the NSF Next Generation Software Program. These experiences demonstrate that the CGrADS leadership can be effective in bringing together a diverse group of busy researchers to achieve a shared vision.

3. Lack of Details

Several reviewers commented on the lack of details regarding how the investigators will address specific research problems and achieve their vision. We admit that the length restrictions of the proposal and the scope of the project made it difficult to include much detail. While the site visit

may be the best time to explore these issues, the preliminary GrADS effort has developed a number of design documents that the site team may wish to peruse. The URL for the GrADS project (funded by NSF Next Generation Software program) is:

http://hipersoft.cs.rice.edu/grads/

The "publications" link on this page leads to a number of project planning documents, although these were written for the three-year horizon of the NGS-sponsored effort.

At the site visit, we will provide more detail on plans. However, we insert a note of caution that CGrADS is a research effort rather than a development project, so plans are subject to change based on what we learn in the early years. One major advantage of an STC is that it can accommodate precisely this type of change.

4. Application Collaborations

One review commented on our policies for working with applications groups, in the following quote:

> "There is also a concern that 'the PIs will work closely with a small number of application groups,' which sounds like they will actually continue in their current mode of making progress one project at a time for those high-profile groups involved with the PACI sites."

We believe that this comment reflects a misunderstanding of the goal of our application collaborations, which is to enable CGrADS researchers to develop a deeper understanding of the problems that affect Grid performance and the challenges that confront users of Grid programming tools and systems. The PIs have already established several such collaborations.

As a part of CGrADS, we intend to seek out new collaborations representative of major new classes of Grid applications (on-demand applications, ubiquitous applications, robust portable applications, and integrated data analysis and simulation). While available resources will limit the number of concurrent collaborations, the PIs will rotate through a variety of applications to ensure that a diverse pool of applications is considered. Once again, we will collaborate with application teams on their codes to better understand the issues they face, to focus CGrADS research goals, to help set research priorities, and to critically evaluate and assess the Grid application development software resulting from CGrADS research.

That having been said, we believe that the work will benefit a far larger community of scientific application developers. CGrADS researchers have close ties to a variety of organizations, programs, and companies interested in high-performance computing applications. The PIs will work through these organizations to attract new users and to effect technology transfer through tutorials, workshops, talks, software, and publications. To reach a larger community of potential users, newsletters, publications, software, and other information about the CGrADS project will be available from the CGrADS web site.

5. Education

The reviewers were generally positive about our education plans. However, we noted one concern regarding the connection between research and education:

> "I would have liked to see more direct connections between research results and educational programs, perhaps in the form of computational science modules developed in conjunction with teachers and parents."

We are already running a two-week pilot program on K-12 computational science materials for teachers who plan to use the materials with their students. The curriculum is described at the following web site:

http://ceee.rice.edu/Books/CS/index.html

In CGrADS, we plan to tie teacher professional development closely to the Center's science. For example, a prototype of what we might do as the Grid develops is the Bugscope activity we initiated this year, in which teachers operated the electron microscope at UIUC's Beckman Institute via the Web for a science lesson. A similar type of activity could be done with the Grid once we are ready to bring it to K-12.

6. Ethics

The reviewers raised two concerns related to ethics and the proposed project:

Training of staff and students in research ethics: Each of the participating institutions has a set of policies that relate to research ethics and intellectual property. Although we view training in research ethics to be an institutional responsibility, rather than one to be addressed by a single project, the investigators will work within their institutions to ensure that all faculty, staff and students achieve an appropriate understanding of these issues.

As pointed out by one reviewer, the investigators have established a reasonable record of dealing fairly with graduate students, including time to graduation, appropriate sharing of credit, and recognition of their intellectual property.

Ethical issues related to the Grid: Widespread use of the Grid is likely to raise a number of ethical issues. Those easily foreseen issues include property rights associated with Grid resources, privacy issues for users and data, and accessibility. The investigators are aware of these issues, and have designed mechanisms into the first generation GrADS architecture that address property rights. For example, a system manager must take active steps to place a system in the Grid. The system itself is active in the resource discovery and allocation process, providing the system manager with direct control over its use.

Although we cannot predict all of the ethical issues that may arise in the future, we recognize the need to be vigilant and to build flexible mechanisms into the system to accommodate a variety of policies.