B Project Summary

Advances in networking technologies make it possible to use the global information infrastructure in a qualitatively different way—as an *active, computational* resource, rather than as a passive information resource. As described in *The Grid: Blueprint for a New Computing Infrastructure*, and demonstrated in numerous compelling proof-of-concept applications, the "Grid" promises to connect the nation's computers, databases, instruments, and people in a fabric of computing and distributed intelligence that can be used on-demand as an integrated problem-solving resource in diverse fields—in particular, for science and engineering research.

However, the power and flexibility of the Grid come at the price of increased complexity in infrastructure and applications. Although ambitious research projects in Grid computing are targeting key challenges such as security, resource discovery, and resource management, one fundamental problem is not being adequately addressed—the programming of these highly complex and dynamic systems to create applications and services. New software tools are required to make Grid resources useful and accessible to scientists and engineers, while delivering high levels of performance and resource efficiency. Those tools must embody major advances in both the theory and practice of distributed computation and of software development.

For the past three years, in the *Grid Application Development Software (GrADS) Project*, we have conducted fundamental research on software frameworks that have the potential to simplify application development and improve performance for the Grid's distributed, on-demand, heterogeneous computational systems. This effort has produced an integrated system architecture and prototype software called *GrADSoft* that schedules and executes *configurable object programs (COPs)* on the Grid. A COP augments the application with routines for mapping and performance estimation, which programming tools can generate semi-automatically. GrADSoft uses these additional routines and innovative information services to intelligently select resources, producing an efficient execution. GrADS research has produced initial versions of COP construction tools, information services and the execution environment, and has applied this software to several applications.

We are now in a position to propose new research that exploits and significantly extends GrADS to both simplify and accelerate the development of Grid applications and services, while making the Grid accessible to a larger community of users and developers. The **intellectual merit** of the five-year *Virtual Grid Application Development Software (VGrADS) Project* is to explore, define, and implement a hierarchy of virtual resources and a set of programming and service models that provide a comprehensive framework for effectively programming heterogeneous, dynamic systems. We will conduct research in three key areas: (1) *virtual Grid (vgrid)* architectures that enable a separation of concerns between high-level programming tools and applications on the one hand and the complexity of dynamic Grid scheduling, fault-tolerance and resource management on the other; (2) programming models, languages, compilers, environments, component libraries, and tools to support creation of Grid applications, services and problem-solving environments; and (3) core software technologies, such as scheduling, online performance monitoring, and resource virtualization to support execution of performance-efficient, reliable, adaptive Grid applications.

We will pursue this research agenda using a strongly application-driven methodology, using a series of leading science and engineering applications to elicit key challenges and validate results. Initially, we have established partnerships with two application groups in biosciences and one in atmospheric science. We will experiment with these applications on two infrastructures developed under GrADS: the *MicroGrid,* a repeatable, observable Grid testbed built by combining simulation and direct execution, and the *MacroGrid,* a modest-sized experimental Grid based on resources at VGrADS sites. These testbeds augment the evolving national "cyberinfrastructure," which we will use for full-scale experiments and demonstrations.

To enhance its **broader impacts**, VGrADS will build on the past successes of our PIs in education and human resource development. We will work through existing programs at our sites, such as the PACI EOT program, to attract and retain women and minorities in computational science. At the graduate level, we will develop new courses incorporating Grid material and pursue long-term student exchanges among research groups. At the undergraduate level, we will create new courses that bring Grid-enabled computing into the classroom, bring students into our research program, and help build support communities for those students to increase retention. To further broaden the impact of Grid computing, we will aggressively pursue technology transfer, including a web presence, software distribution, and public tutorials. Taken together, these education, outreach and technology transfer programs will contribute to revolutionary changes in the ways scientists, engineers, and other professionals solve their problems on the Grid.