Areas of Education, Research, HPC Resources, Technology Transfer and Outreach

Education

HPC educational efforts will focus on the following goals:

- Physical science and engineering students --- both graduate and undergraduate --- should be educated in computational science techniques, thereby positioning them for careers in cutting edge science and technology.
- Scientists whose research can be improved via effective use of HPC technologies should be educated about HPC issues, hardware architectures, software technologies and problem solving environments.
- HPC researchers, not only faculty and staff but especially students, should be educated in the use of existing and emerging HPC technologies.
- Researchers who employ HPC in their scientific software should be educated in software design methodologies that allow maximal exploitation of HPC resources.
- Physical science and engineering students --- both graduate and undergraduate --- should become aware of, and incorporate into their research codes, techniques of sound software design, coding, debugging, testing, maintenance, porting and performance evaluation.
- Computer science students with an interest in HPC --- both graduate and undergraduate --- should be taught to act as consultants to application teams that use HPC (or that want to begin to do so), providing a pathway for disseminating computer science techniques that are likely to prove useful in many research projects, as well as teaching the next generation of computer scientists to engage in interdisciplinary computational research.
- Student researchers should have opportunities to present ongoing projects to an audience of HPC-educated colleagues, in order to obtain feedback on improving their investigations.

To achieve these goals, OSCER will provide a variety of HPC-oriented educational services, such as the following:

- OSCER will provide web-based programmer education materials, targeted to student researchers in physical science and engineering fields, in order to help them to develop sound habits of software design, coding, debugging, testing, maintenance, porting and performance evaluation. (Some of these materials have already been developed for CS 1313, the programming course for non-majors, which is targeted primarily to physical science and engineering students, especially geoscience students.)
- OSCER will assist interested instructors in coordinating coursework that has a direct relationship to HPC (e.g., Prof. S. Lakshmivarahan's Scientific Computing course sequence in the School of Computer Science), both to provide reliable access to HPC resources for use in pedagogical activities, and to assist in developing materials that will expose students to HPC issues and that will educate them in making effective use of HPC resources. Included in this activity will be the facilitation of the development of new HPC-based courses.

- OSCER will participate in externally funded student support initiatives such as the National Science Foundation's Research Experience for Undergraduates program. OSCER will be especially well-positioned to participate in such programs, because of its strong links to research teams in a variety of scientific disciplines.
- OSCER will conduct frequent workshops on HPC issues, open to the entire OU community, but targeted not only toward active HPC-based researchers, but also toward scientists with little or no HPC expertise but whose investigations could be enhanced by suitable application of HPC methodologies. These workshops will cover HPC issues such as current hardware architectures, parallel programming paradigms and strategies, scalar optimization, existing and emerging HPC software technologies, high performance communication, and integrated grid infrastructures ("the Grid"). In some cases, these activities will serve to introduce HPC to new users; other presentations will focus on the needs of experienced HPC consumers.
- OSCER will provide one-on-one instruction, both of HPC consultants and of research software developers. These mentoring activities will be conducted by researchers with strong HPC background and a well-defined educational mission with an appropriate reward structure. This reward structure will consist of improved financial remuneration for successful mentoring, opportunities for career advancement, and public recognition of achievement.
- OSCER will conduct frequent seminars at which researchers, especially student researchers, will present their projects and obtain feedback from other scientists. In some cases, these seminars may include guest speakers from outside the OU community.

The expected outcomes of this comprehensive educational effort will include:

- The University's ability to conduct computational physical science and engineering research will substantially improve, leading to increased external funding, publications and prominence for University scientists, as well as attracting an improved pool of student, faculty and staff candidates.
- Students in a variety of disciplines will conduct cutting edge, publishable research in computer science and in computational physical sciences and engineering.
- Students participating in HPC-based research will have opportunities to obtain feedback from experienced HPC users, and will become accustomed to presenting scientific methodologies and results.
- OU will prepare these students to take positions in the HPC marketplace, or to pursue (or continue) graduate studies in HPC, not only by gaining knowledge of HPC technologies and methodologies, but also through the apprenticeship opportunities associated with participating in cutting-edge scientific and technological research.
- Faculty who are developing funding proposals will have resources to refer to when discussing educational aspects of their projects, an issue that is becoming increasingly crucial to some funding agencies.
- OU graduates whose academic research has focused on HPC-based science will be well-qualified --- uniquely so within the State of Oklahoma --- to pursue Oklahoma-based entepreneurial opportunities that require HPC expertise, and these opportunities will be enhanced by OSCER's technology transfer support role, described below.

- OSCER will provide a base of HPC expertise from which to develop interdisciplinary computational science degree programs.

Research

OSCER will facilitate two kinds of research: research of applications that employ HPC technologies, and research about HPC technologies. It is likely that the majority of OSCER-based research projects will be of the former category, because of the ratio of the numbers of members of the former category to members of the latter among the target population.

OSCER will assist in HPC-based research in several ways:

- OSCER will attract researchers from disparate disciplines who have common interests, and will foster synergistic interdisciplinary collaboration, which will enhance research teams' ability to obtain external funding.
- OSCER will assist principal investigators in developing funding proposals for HPC-based projects, providing not only a set of resources to list in proposals but also expertise on HPC issues that are of interest to funding agencies.
- OSCER will provide not only HPC facilities but also HPC education for researchers with insufficient technical background.
- OSCER will work with participating researchers to develop appropriate paradigms for effective apportionment of resources, and will balance the needs of multiple disparate groups in a fair, unbiased manner. Unconventional usage paradigms will be accommodated to the extent possible, as long as sufficient notice is given to allow other researchers to manage their resource allotments appropriately.
- OSCER will provide HPC resources to members who are OU students, faculty and staff without assessing usage fees.

HPC Resources

OSCER will provide several categories of HPC resources, including hardware, software and networking infrastructure, as well as support for system administrator interaction.

Hardware

OSCER's HPC hardware will include:

- **HPC platforms**
  OSCER's initial HPC systems will be
  - a large cluster of inexpensive commodity processors with good memory bandwidth and a fast interconnect network, to provide distributed parallelism (e.g., MPI);
  - a small Symmetric Multiprocessor (SMP), to provide a modest degree of shared memory parallelism (e.g., OpenMP) and hybrid parallelism.
- **Condor**
  OSCER will also examine the usefulness and practicality of providing Condor services,

both to increase the available set of computing resources and to maximize OU's usage of its installed hardware. Condor, developed by a team at the University of Wisconsin, is a system for "High Throughput Computing." Specifically, Condor monitors a pool of computers (e.g., desktop machines in offices and labs); those machines that are discovered to be idle can be assigned compute jobs. A Condor job's I/O is not performed on the remote machine that is performing the computation; instead, I/O data are sent over the network to the machine from which the job was submitted, or to a designated I/O machine. When a remote machine that is running a Condor job changes its state from idle to active (i.e., when a key is pressed; when a telnet session is initiated), the Condor system checkpoints the run, migrates it to another idle computer, and restarts it appropriately. Thus, the only remote resources that are consumed are idle processors, idle memory (physical and virtual), and network bandwidth. The above mode of operation is called the "standard universe;" Condor can also operate in a stripped down "vanilla" mode, as well as in MPI, PVM and Globus modes. OSCER will provide not only Condor software and management support, but also a checkpointing server, which will allow Condor jobs to achieve maximal throughput.

- **Data Archive**
  To the extent possible, OSCER will provide hardware, software and services for OU researchers to archive their scientific data. This archive will be used for medium term storage; additional space will be provided for temporary storage while experiments are being conducted, and long term storage will be off site (e.g., at the National Center for Supercomputing Applications, the San Diego Supercomputer Center and/or the Pittsburgh Supercomputing Center).
- **Visualization**
  To the extent possible, OSCER will provide expertise, software and (where possible) hardware for OU researchers to visualize their scientific data.

Software

OSCER's HPC software will include such tools as compilers, parallelizers, debuggers, profilers, scientific computing libraries, I/O libraries, mathematical software systems, code management software, performance evaluation systems, visualization systems and integrated grid infrastructure software services.

OSCER will be responsible for maintaining a software pool that is sufficiently broad, deep and up-to-date to meet the needs of its members. In addition, OSCER will be responsible for having personnel sufficiently well-educated in the use of these software systems to be able to provide educational services for members who wish to use them.

Networking Infrastructure

OU is a member of Internet2, a consortium of over 185 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies, accelerating the creation of a new generation of high bandwidth communication infrastructure.[2] Internet2 is recreating the partnership among academia, industry and government that gave rise to the original Internet. The primary goals of Internet2 include:

- to create a leading edge network capability for the national research community;
- to enable revolutionary Internet applications;
- to ensure the rapid transfer of new network services and applications to the broader Internet community.

The two high performance connections from the university are a 1 Gbps connection within OU and an OC3 (155 Mbps) connection from OU to OneNet (our statewide network provider). The connection from OneNet to the Abilene (the Internet2 backbone) is OC12 (622 Mbps), which is 45,000 times faster than a standard 56k telephone modem.

OSCER will employ networking facilities of the highest possible bandwidth, not only internally within OU, but also externally to remote sites.

System Adminstration Support

Currently, OU research groups that employ HPC have local hardware systems that are managed by a variety of system administration teams. As is generally the case, these staff members are swamped by high workloads, and they have no reward structure for cooperating between independent groups.

Each group has a unique set of knowledge and experience that might very well prove useful to the others. But with no time and no resources earmarked for interacting with other groups, they suffer from a lack of shared information. The result is that the organizations that they serve have substantially less computational capability to draw on, because few if any of the administration groups have sufficient background to provide all of the services that their clients require.

OSCER will address this problem by devoting a meaningful amount of resources to interactions among system administration groups, particularly in helping the various groups develop the skills necessary to manage systems intended for HPC use.

Technology Transfer

HPC research is in a strong position to engage in technology transfer, because, in many cases, a fundamental product of such research is scalable, portable, extensible, maintainable software that has immediate utility for industrial and/or governmental users. OSCER will support technology transfer in several ways.

- OSCER will assist researchers in developing their codes in a manner consistent with the needs of potential corporate and/or governmental clients, and will engage in ongoing pursuit of information about the needs of such organizations, in order to provide the best possible information to HPC-based scientific teams.
- OSCER will establish linkages with the OU Office of Technology Development, in order to provide a conduit through which researchers can move their software into external use.

Outreach

OSCER will engage in outreach activities, to assist not only in disseminating the results of HPC-based research, but also to attract potential partners --- both scientific and industrial --- into such investigations.

- OSCER will maintain a website that includes links to descriptions of, and results from, research projects that employ OSCER resources. To achieve this goal, OSCER will provide web space to research groups that do not have their own websites, as well as links to existing websites.
- OSCER will maintain a list of publications and other scholarly works produced via the use of OSCER resources.
- OSCER will sponsor workshops and conferences that showcase the activities of its member researchers, and will foster the participation of its members in regional, national and international HPC conferences.
- OSCER will participate in student recruiting activities such as Engineering Open House and Geosciences Week, and will exhibit the work of participating students, depicting OSCER's mentoring role and its impact.
- OSCER will publish and disseminate a newsletter that calls attention to the ongoing activities and successes of the membership.