

Facilities, Equipment, and Other Resources

The Department of Computer Science maintains computing facilities to satisfy a variety of research and educational needs. The Computer Science Laboratory staff of 4 people installs, tests, and maintains computing facilities for the Department. Faculty, graduate student, and staff offices are equipped with workstations and/or personal computers. The department also maintains all needed peripherals including scanners, printers, video capture, and projection systems.

The Department also maintains research systems to meet the specific needs of individual research projects. This infrastructure, described below, is funded by donations from industry partners and equipment grants from the National Science Foundation CISE Directorate, including a 1999 Academic Research Instrumentation grant.

The Department, through the University, is linked to the 2.5 gigabit-per-second network ring established by the North Carolina Networking Initiative (NCNI). NCNI joins universities and industry in the Research Triangle Park area, with connections to Internet providers, the Internet 2, National LambdaRail, and the North Carolina Research and Education Network (NCREN).

The following sections summarize the Department's core facilities and the supplemental computing facilities available at Duke University and Research Triangle Park.

Department Infrastructure

The Department of Computer Science occupies quarters in the Leon Levine Science Research Center (LSRC), an \$80 million facility constructed to encourage innovative collaborations among the sciences, engineering, environmental studies, and medicine. The Department's space is designed for flexibility and to allow easy expansion of its state-of-the-art communications infrastructure. There is ample space for laboratories and workrooms to support research and group interactions.

The Computer Science Laboratory provides the basic computing infrastructure for the Department. The lab staff is responsible for installing and maintaining all Department computers and providing all required network services within the Department. In addition, the staff works with faculty and students to accommodate special research and academic needs not covered by the general computing infrastructure.

Workstations

The Department maintains a variety of desktop and server computing resources. All faculty and graduate student desktops are equipped with Dell Linux workstations with advanced graphics displays and at least 16 GB of RAM. In addition public computer kiosks are located strategically around the Department to provide student and guest access to Department resources.

File Servers

The collective disk space on the servers totals well over 70 TB; this space is available to all machines within the Department and does not include hundreds of gigabytes of local scratch space on each machine. Primary user file space is provided by a pair of mirrored Network Appliance FAS 3240 servers with 62 TB of disk space, connected to a 10GB fiber link for fast access. This space, along with many other project spaces, is mirrored and backed up on a nightly basis. Backups are copied and stored in fireproof tape cabinets in two monitored, climate-controlled machine rooms.

Cluster

The department also maintains a cluster of machines for computing with both CPU and GPU resources. The machines comprise approximately 1200 cores with memory ranging from 24-256GB of RAM. Jobs may be submitted to the cluster either individually or through the use of a batch submission system.

Networking

There are a variety of supporting communications facilities in the Department including 10GB and 1GB connections for servers, and high-speed ethernet connections for all computers and workstations, as well as 802.11/a/b/g/n wireless base stations.

Networking within the Department is supported with two Cisco 3560-E 10GB Switch/Routers which connect the department to the university over redundant 10GB fibers. These switches act as concentrators to 20 other Cisco switches that provide 1GB and 100MB connections to servers and desktop machines.

A campus-wide fiber-optic network, combined with redundant 10Gb connections, provides high-speed data paths to computing facilities at several other sites in the state of North Carolina, including National LambdaRail, and Internet2.

Network Services The Department provides many of its own services, independent of the University, in order to provide a flexible infrastructure that can respond rapidly to the research requirements of the faculty and students. These services include Domain Name Service (DNS) for the 14 subnets the Department maintains, inter-department routing, as well as independent email, printing, and web services. Dynamic Host Configuration Protocol (DHCP) service, to distribute IP addresses, is provided to allow faculty and students to operate their own personal computers on the Department network.

Classroom Resources

The department maintains 8 classrooms for the use of teaching and collaborating. All rooms are equipped with current generation projectors and sound systems. In addition, we house facilities to record and distribute presentations on the internet as well as iTunes-U.

In addition to the many general servers that the Department maintains, the Department currently has the following notable resources:

- * Shared cluster containing approximately 1200 cores
- * 10 Dell PowerEdge 730 40 cores and 256 GB RAM each
- * 18 Dell PowerEdge 610 with 48 GB RAM each
- * 10 Dell PowerEdge 620 48 cores and 256 GB RAM each
- * 10 Quantum Tensor TXR231-1000R 40 cores and 512GB RAM with 1x Tesla P100s
- * 3 Quantum TXR430-0512R 32 cores and 256GB RAM with 4x Tesla K80s
- * 4 Quantum TXR113-1000R 48 cores and 256GB RAM with 4x Tesla P100s
- * 2 Dell PowerEdge 610 with 96 GB RAM each
- * 2 Dell PowerEdge 610 with 128 GB RAM each
- * 2 Network Appliance FAS 8020s with over 62TB storage
- * 2 Network Appliance FAS 3240 with 62TB
- * 2 Cisco 3560-E 10GB Switches/Routers
- * 8 Cisco 3750-G 1GB Switches
- * 8 Cisco 3750 10/100MB Switches

Research Facilities

In addition to the common infrastructure, the Department maintains resources for specific research applications. The resources include the Systems and Architecture Laboratories, the Applied Geometry Collaboratory, AI, Robotics, Computer Vision Collaboratory, and the Donald Laboratory.

Systems and Architecture Laboratories

The systems and architecture group maintains server clusters and storage systems for experimental research, including interdisciplinary projects involving researchers in massive-data algorithms and computational sciences. This facility includes three machine rooms in two different buildings linked by a fiber network. The compute cluster as of 2011 consists of over 98 cores of server class machines including the latest processors from Intel. The group also has a 105 node cluster for computer architecture simulations. Sun Grid Engine schedules jobs to this cluster along with the CS department shared cluster. Data is stored on either department file servers or on a dedicated ZFS storage server.

The system group collaborates closely with the Duke's campus cluster (DSCR: <http://www.oit.duke.edu/comp-print/research/>), sharing equipment and machine room space. While the Duke campus cluster is primarily production oriented, those resources are available to a limited degree for systems research projects in managing large computing resources.

The systems group is a major partner in the region-wide BEN project (<https://ben.renci.org/>) based at RENCI (<http://www.renci.org/>). That relationship includes access to the BEN fiber ring connecting Duke with UNC-Chapel Hill, NC State, and RENCI. This network is fully experimental down to the physical layer. Fiber optic cables connect machine rooms to the RENCI's point of presence on the Duke campus.

This infrastructure is used for research projects in computer architecture, distributed systems, operating systems, networking and network storage, utility computing, large-scale Internet services, and data-intensive systems, and databases. It is funded in part by donations from industry partners including IBM, Intel, Network Appliance, Hewlett-Packard, and Cisco.

Applied Geometry Collaboratory

This laboratory applies computational geometry algorithms to problems in structural biology, Geographic Information Systems (GIS), and large-scale spatial and temporal databases. It supports high-end visualization and solid model hardware. In particular, this lab houses a front-projected H-1 stereographic visual display with four freefield speakers. The room supports real-time 3D human-computer interaction via a long-range Polhemus Fasttrack system that tracks the position and orientation of a person's head and hand. The graphics and audio display can be controlled by an SGI Onyx2, SGI Octane, or a PC. The Onyx2 has 3 Infinite Reality2 graphics pipelines, two audio rendering boards, each with 8-channel ADAT digital output, 6GB of memory, 10 250MHz R10000 CPUs, and 128GB of disk. In addition, the lab houses a 3D Thermojet solid object printer, four PC workstations, and a large DELL disk array.

AI, Robotics, and Computer Vision Collaboratory

In addition to a variety of Sun UltraSparc and Dell PC workstations, this facility is home to advanced robotics and computer vision research equipment. The lab includes a high-quality Riegl LMS Z390i laser range finder for determining shapes and positions. A six-camera Motion Analysis Hawk motion capture system allows recording detailed motions of people or

vehicles. This 200 frames-per-second system is synchronized through a SMPTE genlock circuit and time-code generator to three high-definition, 30 frames-per-second Canon XL-H2 video cameras with microphones. A full suite of digital still and video cameras, three stereo vision systems, as well as a photography-grade lighting system and background curtains, allow recording standard A/V data simultaneously with the motion and shape measurements. An iRobot ATRV-Junior robot vehicle, donated by Science Applications International Corporation (SAIC), is being used in both research and education programs. It features a laser range-finder, sonar sensing, stereo vision, GPS, voice recognition, and a series of other attachments that make it a versatile research vehicle.

Donald lab

Dr. Donald's laboratory contains multiple high-end Wintel and Linux workstations plus computational servers with large disks, RAID, and backup systems. Servers and clusters are maintained in a central machine room in LSRC. This includes access to experimental facilities for structural and molecular biology, including NMR, X-ray diffraction, CD spectroscopy, HPLC, column chromatography, gel electrophoresis, mass spectrometry of biomolecules, and PCR. The lab's computational resources include a 35-node, 107-processor, 216-core Dell Cluster with the following specifications: 16 PowerEdge 1950 nodes (32 Dual-Core Xeon 5150 processors) each with 8GB RAM; plus 19 PowerEdge 1950 nodes (38 Quad-Core Xeon E5420 processors) each with 16GB RAM. In addition to the facilities in the LSRC, Dr. Donald has 950 sq ft. of wet-lab space in the French Family Science Center, immediately adjacent to LSRC. Dr. Donald's laboratory is fully equipped to overexpress and purify proteins and protein complexes, and to perform binding and activity assays.

University Facilities

The backbone and other university-level infrastructure needs of the University are maintained by a central IT organization, the Office of Information Technology (OIT). OIT is responsible for the operation, testing, support, and engineering of the campus-wide data, voice, and video communications infrastructure. This includes the design and subsequent implementation of structured wiring and switching systems, enterprise-level servers, including Domain Name Server (DNS) and Dynamic Host Configuration Protocol (DHCP) servers, routing systems, and wireless systems.

Duke University's high-speed backbone, DukeNet, provides researchers, staff, faculty and students with a robust, redundant conduit for data. The backbone consists of Cisco routers with redundant 10 gigabit ethernet links. Most buildings on campus are wired with Category 5 cabling and have 10M/100M Ethernet ports supplied to each desktop. Servers and high speed research workstations can be provided with gigabit or ten gigabit ethernet ports as needed. Building networks connect to the backbone via dual gigabit or 10 gigabit ethernet uplinks.

The Duke Shared Cluster Resource (DSCR) facility maintains a shared computational cluster facility of over 600 machines (1152 processors) to which we have access. The processors range from 2.8GHz to 3.6GHz. While the cluster must be shared by the entire community at Duke, it provides a useful resource for computational science.

The Duke University High Resolution Nuclear Magnetic Resonance (NMR) Spectroscopy and X-ray Crystallography Shared Resource provides state-of-the-art instrumentation and methods to our laboratory. The Shared Resource operates and supports advanced instrumentation and technology and has a professional staff of three experienced Ph.D. scientists and an instrument specialist/engineer. Major equipment currently installed in this Facility includes two fully-equipped X-ray crystallography systems

with R-Axis IV and R-Axis II area detectors, focusing mirrors, and liquid nitrogen cooling systems. NMR equipment includes Varian Inova spectrometers operating at 800 MHz, 600 MHz, 500 MHz (2), 400 MHz and 300 MHz (2) with heteronuclear multi-channel capabilities. Both the 600 and 800 MHz NMR spectrometers are fully configured with four channels and ¹H, ¹³C and ¹⁵N triple-resonance cold probes with Z-axis gradients. All spectrometers have the capability of deuterium decoupling. A full complement of homo- and heteronuclear experiments have been implemented on these spectrometers and are routinely used to study biological macromolecules. Users of this facility also have access to the Southeast Regional Collaborative Access Team beamlines at the APS synchrotron X-ray source at Argonne National Laboratory and the 900 MHz NMR spectrometer at the University of Georgia. The Duke High Resolution Nuclear Magnetic Resonance Spectroscopy and X-ray Crystallography Shared Resource is located in a custom-designed wing of the LSRC, which is in close proximity to the Computer Science Department and the Biochemistry Department.

The University also maintains a campus-wide AFS file system infrastructure with terabytes of storage; a campus-wide electronic mail infrastructure supporting over 35,000 mailboxes and handling in excess of a million messages a day; a server-based file service, authentication services; directory services; web service; and name service and other network services.

External Computing Facilities

The campus connects to the commodity Internet, as well as the Internet 2, and the National LambdaRail through the North Carolina Networking Initiative. NCNI is a joint venture between Duke University, North Carolina State University, the University of North Carolina at Chapel Hill and the Microelectronics Center of North Carolina. The NCNI ring is comprised of 3 pairs of single-mode fiber that cover more than 100 miles. At the present time the ring uses DWDM (dense wave division multiplexing) technology to provide OC-48 data rates (2.4 gigabits/second) to users.

NCNI members share three commodity Internet gateways and one Internet 2 gateway. Current service levels and providers are as follows:

- * Sprint: OC-12 to commodity Internet (622 Mbps)
- * UUNET: OC-12 to commodity Internet (622 Mbps)
- * Qwest: OC-12 to commodity Internet (622 Mbps)
- * Abilene: OC-48 to Internet 2 (2.4 Gbps)

Routing decisions are made at NCNI, eliminating the need for researchers to determine routes. All traffic destined for an Internet2 member institution will automatically be routed out the Abilene link.

Additional computing power is provided to the Department through the North Carolina Supercomputing Center (NCSC), a division of MCNC. NCSC supports and promotes the use of computational science for education, research, and commercial applications. High Performance computing (HPC) resources are provided by NCSC for use by higher education, commercial, and government organizations. HPC resources available at NCSC for workstation users in the Computer Science Department include high performance computing platforms, a diverse application software suite, visualization resources, training support, and collaboration. The computational resources have been regularly updated, moving from the original four-processor CRAY Y-MP, to the most recent addition of a more than 700-processor IBM RS/6000 SP. These regular technology upgrades insure that the Center's user community retains access to the state-of-the-art computational resources their work requires.