NC STATE UNIVERSITY

40

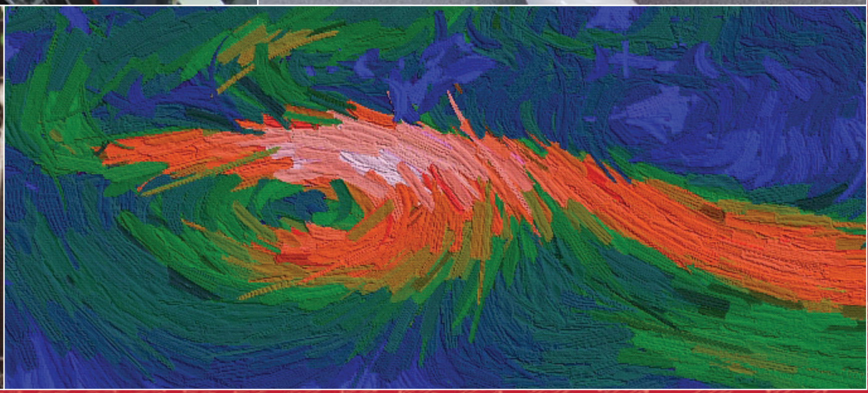

Department of Computer Science

1967–2007

40 years of research and education

Dr. Mladen Vouk, Head
Department of Computer Science

Welcome

Message From the Head

When the first computers were installed in Patterson Hall in 1957, there must have been great speculation about the impact of computing. Ten years later, in the fall of 1967, the NC State Department of Computer Science was "officially" established. Over the next 40 years, we have emerged as one of the largest and most productive computer science departments in the nation. We now have 44 outstanding faculty, 15 NSF CAREER award winners, and expertise that ranges from systems to theory, from artificial intelligence to networks, and from software engineering to bioinformatics.

As part of a land-grant institution adjacent to the world-famous Research Triangle Park, we have held true to our mission—serving citizens of North Carolina by educating a strong talent pool to meet the growing technology needs of our state and beyond. Award winning programs like our Senior Design Center have helped to ensure that our students graduate with real-world skills to accompany a solid technical foundation. With over 5,000 alumni, we rank among the leading producers of computer science talent nationwide, and are proud to consistently act as a highly valued source of new graduates for industry.

computer science
NC STATE UNIVERSITY
67 | 07

Since its launch in 1989, our graduate program has experienced steady growth and attained an outstanding reputation worldwide. Our researchers have been equally successful. New state-of-the-art facilities on NC State's Centennial Campus house over 30 research groups, laboratories and centers. Strong industry ties have helped drive curriculum evolution in areas such as Security and Services Sciences, and technology innovation in areas that include our award-winning Virtual Computing Laboratory.

Over the years I have had the privilege to know and work with many who have shaped the department. I am extremely proud of all our students, faculty and staff, past and present, and of our alumni and corporate friends. I would like to thank them, as well as the former department heads and the NC State University, UNC, and State of North Carolina administrations for the vision, support, and effort that allowed this department to excel.

The pages that follow provide some insights into the paths the department has followed, its current education and research activities, and its future directions. While it is hard to give appropriate credit to all who have contributed to our success, I hope that these snapshots will whet your interest. Please join us in celebrating our 40th anniversary.

Contents

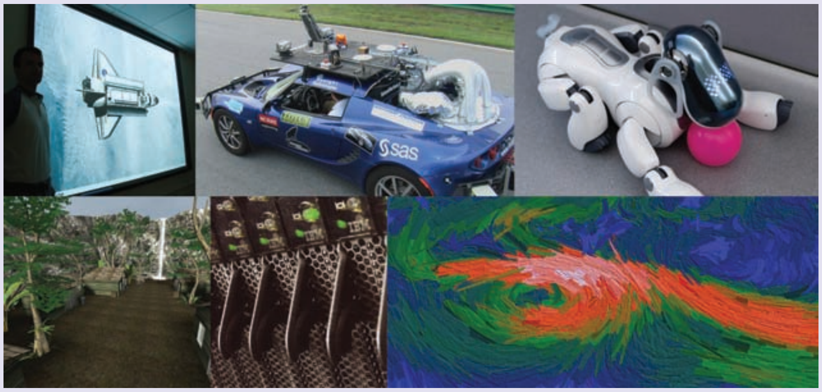

About the cover (clockwise from top-left): four-projector display in Dr. Watson's Design Graphics Laboratory; autonomous Lotus Elise entered by NCSU COE in the DARPA Urban Grand Challenge; Sony AIBO from Dr. St. Amant's Knowledge Discovery Laboratory; artistic visualization of a simulated supernova collapse; network servers in the Systems Research Laboratory; Crystal Island *online learning environment from Dr. Lester's IntelliMedia Center*

Department of Computer Science
North Carolina State University
890 Oval Drive
Raleigh, NC 27695-8206
919.515.2858

http://www.csc.ncsu.edu

Fall graduation ceremony

Facts and Figures

Achievements by the NC State Department of Computer Science now match or exceed those of other top departments throughout the country. Our reputation for excellence has also improved. The ASEE[1] consistently ranks us among the top 10–20 departments in enrollments, graduate rates, and research funding. We were 47th in the nation in the 2006 US News & World Report ranking of graduate Computer Science programs (27th among public universities), up from 49th in 2002.

Student Enrollment. In 2006–2007 the department included 628 undergraduate, 254 MS, and 144 PhD students. Our PhD enrollment per faculty matches the average for departments nationally ranked 13–24[2]. We also have one of the largest MS programs in the country.

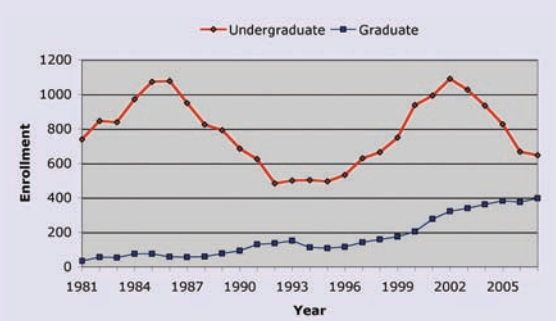

We ranked 7th in the nation in undergraduate enrollment for Computer Science departments in Engineering schools[1]. The rate of decrease in our undergraduate enrollment is below the national average, with only slightly fewer students (607) anticipated in 2007–2008.

Graduation Rates. The department graduated 105 undergraduate, 101 MS, and 18 PhD students in 2006–2007.

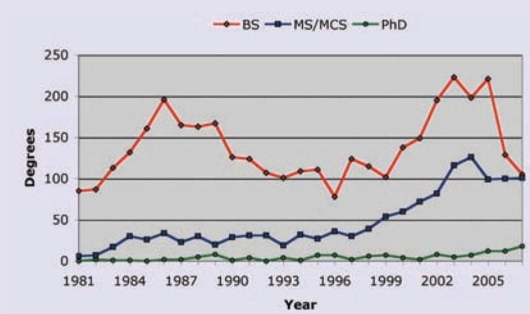

While our undergraduate totals follow national trends, overall our graduation numbers continue to strengthen. Our PhD graduation rate matches the average for departments ranked 13–24, and our MS rate exceeds the average for departments ranked 1–12[2].

Research Expenditures. Annual research expenditures continue to grow. In 2006–2007 we had $22 million in active grants, with total research expenditures estimated to exceed $7.5 million, an increase of 6.7% over the previous academic year.

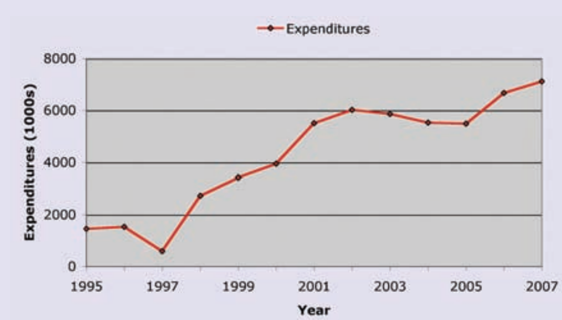

1 American Society for Engineering Education, http://www.asee.org

2 2007 CRA Taulbee Survey, http://www.cra.org/CRN/articles/may07/taulbee.html

Awards


2006 IEEE Computer Society International Design Competition winners

Selected Faculty Awards

- **Alcoa Foundation Engineering Research Achievement Award**, George Rouskas
- **American Association for Artificial Intelligence Fellow**, Jon Doyle
- **CRA Digital Government Fellow**, Annie I. Antón
- **DOE Early Career PI Award**, Xiaosong Ma
- **Distinguished University Research Professor**, Donald Bitzer
- **Emmy Award**, Donald Bitzer
- **IDA/DARPA Defense Science Study Group Member**, Annie I. Antón
- **IEEE Fellow (2)**, Donald Bitzer, Mladen Vouk
- **IFIP Outstanding Service Award**, Harry Perros
- **National Academy of Engineering Member**, Donald Bitzer
- **National Academy of Television Arts and Sciences Award**, Donald Bitzer
- **NSF CAREER Award (15)**, Annie I. Antón, Rada Y. Chirkova, Vincent W. Freeh, Khaled Harfoush, Christopher G. Healey, James C. Lester, Xiaosong Ma, Frank Mueller, Peng Ning, Injong Rhee, George N. Rouskas, Munindar P. Singh, Ben Watson, Laurie Williams, R. Michael Young
- **NCSU Academy of Outstanding Teacher Award (8)**, Edward Davis, Christopher G. Healey, James C. Lester, Carolyn Miller, George N. Rouskas, Alan Tharp, Laurie Williams, R. Michael Young
- **SAS Professor of Computer Science**, Jon Doyle
- **Woman of Influence in the Public Sector Award**, Annie I. Antón

Dr. Donald Bitzer receives an Emmy Award for co-inventing the plasma display

Selected Student Awards

- **Fulbright Fellowship**, Bushra Anjum
- **Google Anita Borg Scholarship Finalist**, Claris Castillo
- **GEM Fellowship (2)**, Beth M. Adams, Lenard Williams
- **IEEE CS International Design Competition Winners, 2005**, David Coblnetz, Dakota Hawkins, Jonathan Lewis, Ben Noffsinger
- **IEEE CS International Design Competition Winners, 2006**, Hunter Davis, Josiah Gore, Eric Helms, Blake Lucas
- **IBM PhD Student Fellowship (2)**, Sarah Heckman, Travis Breaux
- **NSF Graduate Research Fellowship**, Kristy Elizabeth Boyer
- **U.S. Department of Homeland Security Fellowship**, Steven McKinney

Department formed July 1967

Head Paul Lewis 1967-1973

Head Donald Martin 1974-1985

Master of Computer Studies 1977

Accreditation 1986

Head Robert Fund 1986-199

1967 — 1970 — 1980


IBM/360 1964


Martin Luther King 1929-1968


Apollo 11 1969

Altair 8800 1975

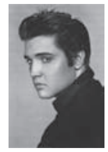
Elvis Presley 1935-1977


Columbia launch 1981


IBM PC 1981


Apple Macintosh 1984

Computer science began at North Carolina State University in 1957, when the Department of Experimental Statistics installed the first computers in the basement of Patterson Hall. In 1965, a National Science Foundation grant fostered collaboration among NCSU, Duke, and UNC Chapel Hill to form the Triangle Universities Computation Center. In 1966, a committee on computer science called for a Computer Science Department (CSC) offering a Bachelor of Science degree. In 1967, Chancellor Caldwell formally submitted a proposal to establish the Computer Science Department. In July of that year, the new department was approved.


Dr. Paul Lewis
Head, 1967–1973

Dr. Paul Lewis became the first Computer Science Department head. Student major numbers increased from 49 in Spring 1968 to 302 in 1973. After beginning in Harrelson Hall, the department expanded to occupy 10 offices in the newly opened Dabney Hall in 1970. The department also recruited its first two computer science PhDs in 1969—Dr. Alan Tharp (Northwestern) and Dr. Robert Fornaro (Penn State).

In 1973 Dr. Lewis stepped down as department head, and Dr. Norm Williamson acted as interim head for over a year. In 1974, Dean Menius asked Dr. Donald Martin to become the new head of the Computer Science Department. In the first four years of Dr. Martin's leadership, the department averaged 350 majors and enrolled more than 2,000 students in all CSC courses.


Dr. Donald Martin
Head, 1974–1985

In 1979, CSC made a major move from Dabney Hall to the newly renovated Daniels Hall. During the same period, Governor James Hunt instructed the North Carolina General Fund to invest $500,000 in the department. In Fall 1982, CSC dedicated a Data General MV 8000 to the support of its instructional program. In Fall 1983, CSC installed a Sage microcomputer system to serve 2,000 students in the basement of Leazer Hall.

Dr. Robert Funderlic took over as head of the Computer Science Department in 1986. In order to maintain the latest technology, the department partnered with ECE to form the Computer Systems Laboratory (CSL) in 1987. In its first year, 50 graduate students worked


Dr. Robert Funderlic
Head, 1986–1992

erlic
2

Move to COE 1989 | MS, PhD programs 1989 | Head Alan Tharp 1993-2004 | Senior Design Center 1994 | SAB 1994 | MSCN program 1999 | ePartners program 2000 | Head Mladen Vouk 2004- | Move to EB-II 2006

1990 | 2000 | 2007


Berlin wall
1989


Mosaic web browser
1993

Google inception
1998


iPod
2001


Mars rover
2004

Harry Potter final volume
2007

on projects at CSL. In 1989 Dr. Donald Bitzer, member of the National Academy of Engineering and co-inventor of the plasma display and the Plato educational system, moved from the University of Illinois to join the department.

One of the department's long-standing desires was the creation of a graduate program administered solely by Computer Science. Through the strong efforts of Dr. Funderlic, with support from the Dean of Engineering Dr. Larry Monteith and Electrical Engineering Head Dr. Nino Masnari, this goal was achieved in 1989 when the department moved from the College of Physical and Mathematical Sciences (PAMS) to the College of Engineering (COE), and instituted new MS and PhD programs.


Dr. Alan Tharp
Head, 1993–2004

Dr. Alan Tharp was named interim head in 1992, and department head in 1993. Dr. Tharp pursued a vigorous program of growth, resulting in numerous new faculty hires in areas such as e-commerce, network security, object technologies, visualization, optical networks, human-computer interaction, and virtual worlds.

The department's rapid expansion spread faculty members over nine separate buildings, including a move to the new Engineering Graduate Research Center (now MRC) on Centennial Campus. October 2003 saw the groundbreaking ceremony for the new Engineering Building II (EB-II), designed to combine the department's academic and research units into one space.


Dr. Mladen Vouk
Head, 2004–

Dr. Mladen Vouk was named interim head in 2004, and department head in 2006, This coincided with the department's move in 2006 into 210,000 square feet in the newly completed EB-II.

Although the department has changed dramatically over the last 40 years, its mission continues to be:

> *To create and disseminate knowledge by constituting a scholarly community focused on research and education in the area of Computer Science, better the state and nation through research in the science and technology of computing, and through our educational programs, equip our students to be competitive, to succeed in their profession, and to contribute to society.*

With continuing advances in technology and education, the Computer Science Department will remain at the forefront of the finest and most prestigious programs in the nation.

Undergraduate Programs

The Department of Computer Science at North Carolina State University offers a modern curriculum focusing on fundamental scientific and engineering principles and methods, exposure to cutting-edge technology, and the opportunity to work on exciting problems with real-world impact. From an initial undergraduate student body of 49 majors in the Spring of 1968, the department's undergraduate program expanded rapidly. We currently enroll 607 undergraduate majors pursuing degrees that include a Bachelor of Science in Computer Science, an Accelerated Bachelor's/Master's Degree, and an Undergraduate Minor in Computer Programming. Curriculum objectives and outcomes are in keeping with the mission of the university, the college, and the department. Our Bachelor of Science program is accredited by the Computing Accreditation Commission of the Accreditation Board for Engineering and Technology, Inc.

All undergraduate students take a capstone project course in their senior year. One of our most popular capstone courses is the Senior Design Project. The course is run by the Senior Design Center, which was established in 1994 to provide undergraduates with realistic project experiences. Businesses provide a project and a contact engineer, and contribute financially to support the Center's activities. Student teams from the Senior Design Center finished first in both the 2005 and 2006 IEEE Computer Society International Design Competition, a worldwide competition between student teams who design and implement computer-based solutions to real-world problems. These are the only times entries from the United States have won this prestigious award.

North Carolina State University

NEAT

Networks for Endangered Animal Tracking

NEAT combined GPS and wireless sensors to track endangered animals, winning the 2005 IEEE CS International Design Competition

The job market for Computer Science graduates has improved dramatically in recent years. Indications are that this trend will continue. Prior to graduation, our students have the opportunity to participate in the university's Cooperative Education Program, allowing them to alternate between semesters of full-time study and full-time paid work experience at companies like Cisco, Red Hat, and the SAS Institute. Informal exit interviews suggest that approximately 60% of our May 2006 graduates had secured positions in industry, with another 20% electing to continue on to graduate school.

Students in the EB-II atrium interacting ... ePartners Caree

Graduate Programs


Students attend a seminar on Service Sciences, Management and Engineering

Graduate degree programs in the Department of Computer Science offer excellent education and research opportunities to outstanding students from across the United States and throughout the world. This includes a nationally ranked PhD program, and Master of Science, Master of Computer Science, and Master of Science in Computer Networking programs.

We currently enroll 478 students (162 PhD, 158 Master of Science, and 158 Master of Computer Science). This includes 198 U.S. students and 103 female students (41% and 22% of our graduate students, respectively). Enrollment has been at a record high for several years, with our PhD numbers growing by about 300% over last four years. Dr. David Thuente, Director of Graduate Programs, attributes this growth to a recognition of the quality of our faculty, to the reputation of the program, to active recruiting of our best U.S. applicants, and to the success of our graduates in obtaining excellent positions in the world's top research laboratories and companies. Our students have consistently lead the COE and the university in terms of Dean's and Alumni Fellowships awarded and have received many prestigious IBM and NSF fellowships.

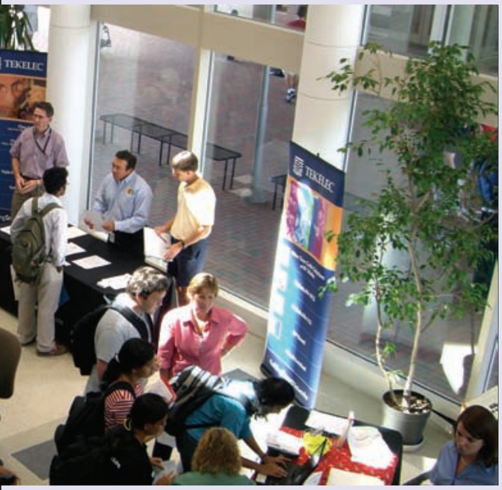

with corporate partners during the annual
Connection event

Numerous seminar series run within the department offer our graduate students an opportunity to meet and attend presentations by well-known researchers and luminaries. Recent visitors include Steve Wozniak, co-founder of Apple Inc., Alan Kay, winner of the Turing Award for work on object-oriented programming, and Douglas Comer, Vice President of Research at Cisco Systems.

Our graduate students also have an opportunity to participate in the university's Cooperative Education Program. During the summer of 2007, 84 Computer Science graduate students held co-op positions, the most of any NCSU department. Participating companies included IBM, Intel, Network Appliance, and Goldman Sachs. Informal exit interviews suggest nearly all of our students are employed following graduation at companies that include Google, Yahoo!, Microsoft Research, Cisco Systems, SAS Institute, and others. Our PhD graduates also occupy academic positions at various universities, including UNC Charlotte, UC Irvine, UT Arlington, and the University of Nebraska.

Characters help high school students explore the research station on Crystal Island, a game that teaches microbiology

Getting Serious About Games

The Digital Games Research Center is a multi-disciplinary center whose focus investigates the scientific, engineering, social and educational challenges of digital games technology. Housed in the Department of Computer Science, the center's faculty include seven computer scientists working on a wide range of research and educational initiatives that study new modes of entertainment and interaction in digital media. Colleagues from the colleges of Education, Design and Humanities and Social Sciences also join in the center's efforts, some of which are listed here.

"The DGRC investigates scientific, engineering, social, and educational challenges of digital games technology"

Intelligent game-based learning environments. To explore the role of intelligent games in education, associate professor Dr. James Lester has developed *Crystal Island*, a narrative-centered learning environment featuring a science mystery designed to teach students elements of microbiology. The mystery is set on a recently discovered volcanic island where a research station has been established to study the island's unique flora and fauna. The student plays the role of the daughter (or son) of a visiting scientist and must uncover the origins of an unidentified illness at the research station. Through the course of her adventure she must gather enough evidence to correctly choose among candidate diagnoses including botulism, cholera, salmonellosis, and tick paralysis as well as identify the source of the disease relying on her knowledge of microbiology to solve the mystery.

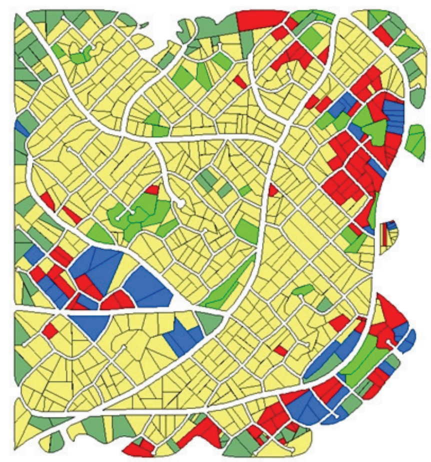

Images of virtual cities that match the complexity and structure of real-world urban environments

Procedural modeling of urban landscapes Cities are important settings in both computer games and film. Generating the digital representations of these rich environments needed for the latest special effects technology is extremely expensive and time consuming, even with the latest 3D modeling tools. The work of Associate professor Dr. Ben Watson seeks to automate the genera-

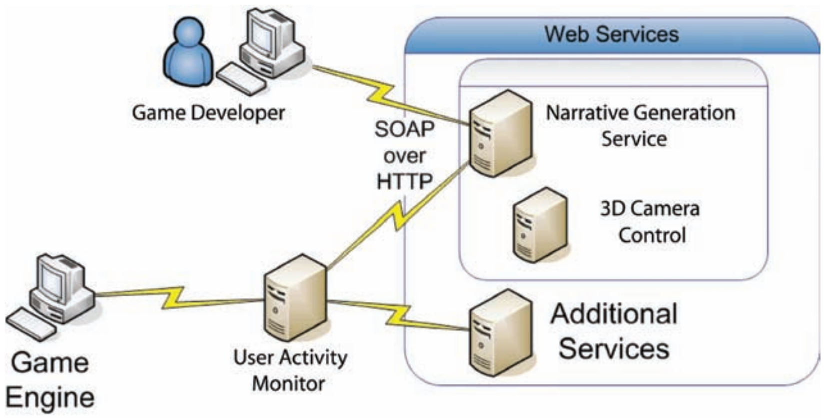

Zocalo *allows network-based services to connect directly to commercial game engines, opening game worlds to a range of research tools*

tion of these urban models using swarm technology, while also providing game and film artists the control they need to meet the demands of their narrative.

Service-oriented architectures for controlling 3D worlds. Commercial off-the-shelf virtual environments like computer games offer rich visual experiences for their users, but the task of integrating research tools (like those described above) into these systems can be complex and computationally expensive. Associate professor Dr. Michael Young's team has developed the *Zocalo* service-oriented architecture for intelligent control of virtual worlds. Zocalo defines a set of open protocols that allow developers to easily connect existing games with a range of computing services across the Internet. Dr. Young is using Zocalo to integrate sophisticated artificial intelligence tools into games like Half-Life 2™ and Unreal™, creating new narrative-based experiences for game users that go beyond current games' abilities to creating exciting storylines and respond to novel user interaction.

Exploiting the allure of games to engage students in science and math. Researchers at NCSU's College of Education are taking advantage of childrens' love for computer games by integrating science learning with game development. Dr. Young is developing a tool called *Virtuoso* that uses a simple, drag-and-drop interface to help novice programmers build 3D computer games. Virtuoso is already being used in the HI-FIVES project, a collaborative effort led by NCSU's College of Education, to engage North Carolina middle and high school students in their science and math curriculum. Using Virtuoso, students and their teachers work together to build games tied to their course curriculums, then share those games with other students in their class and throughout the state.

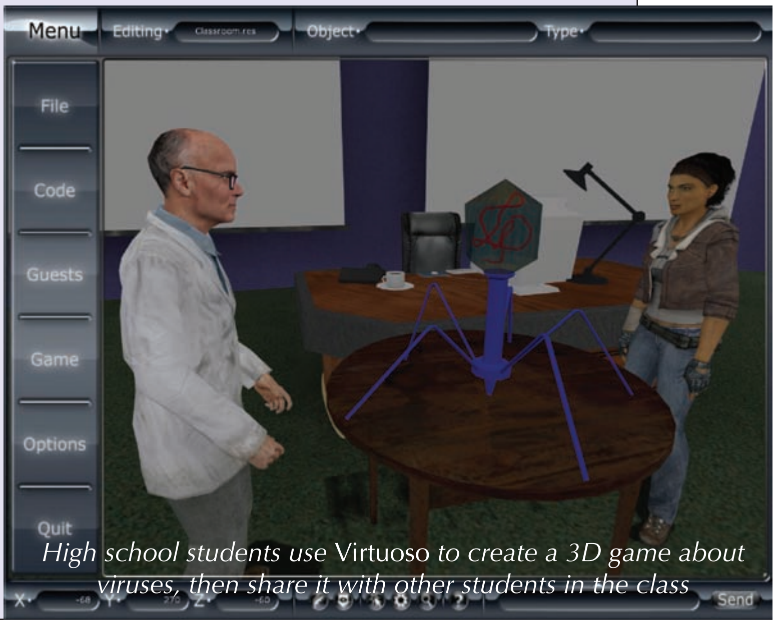

High school students use Virtuoso *to create a 3D game about viruses, then share it with other students in the class*

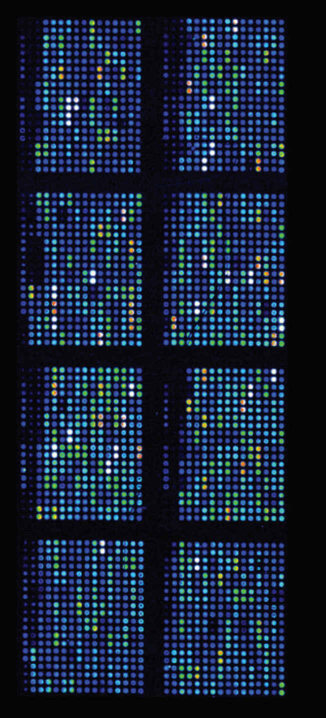

DNA microarray data

Data-Driven Modeling of Biological Complexity

What are design principles for a biological system that produces low-cost ethanol from plant biomass? Or how do microbes convert toxic waste to nontoxic substances? Answers to such questions are of groundbreaking significance, but are hard to obtain because biological systems are inherently complex. For example, in humans fifty or more proteins may influence the activity of a single gene. Computational modeling that reproduces and predicts such behavior forms the Holy Grail of systems biology.

Our understanding of biological systems can be measured by our ability to reconstruct their inner workings and to predict their dynamic behaviors in response to changes in environmental conditions. Unlike physics, where the four Maxwell's equations describe all the electromagnetic phenomena, predictive simulations of biological systems' dynamics require *data-driven model building*. This promising approach aims to interrelate emerging disparate and noisy "omics" observations by relying on mathematics, computer science, information technology, and computing.

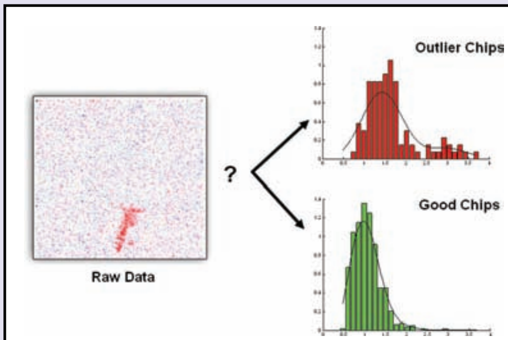

Automated quality assessment allows users to handle large data volumes

From experimental data to components of biological systems. High-throughput experimental technologies such as DNA sequencing and microarray technology have created a unique opportunity to screen thousands of genes in a matter of hours. Yet there remains the growing gap between the high-throughput biological data and the analytical tools capable of deriving all the working "components" from the data—the key step in building models of biological systems. Having a tsunami of data is becoming "a curse" rather than "a blessing," and today's computational limitations in mining these data are tomorrow's nightmares when attempting full systems-level understanding.

For example, analyzing Affymetrix GeneChip microarray data is a sophisticated, time consuming process with many potential sources of anomalous variation that could compromise the results if left uncontrolled. Quality assessment for microarray data is an important but challenging task due to enormous data volume and data complexity. Robust statistical methods often allow users to effectively utilize data that contains small artifacts. However, in some cases, microarrays are beyond correction, and removing the defective arrays from the data set is warranted.

Dr. Heber and his colleagues aim to apply elements of software design, artificial intelligence, and machine learning

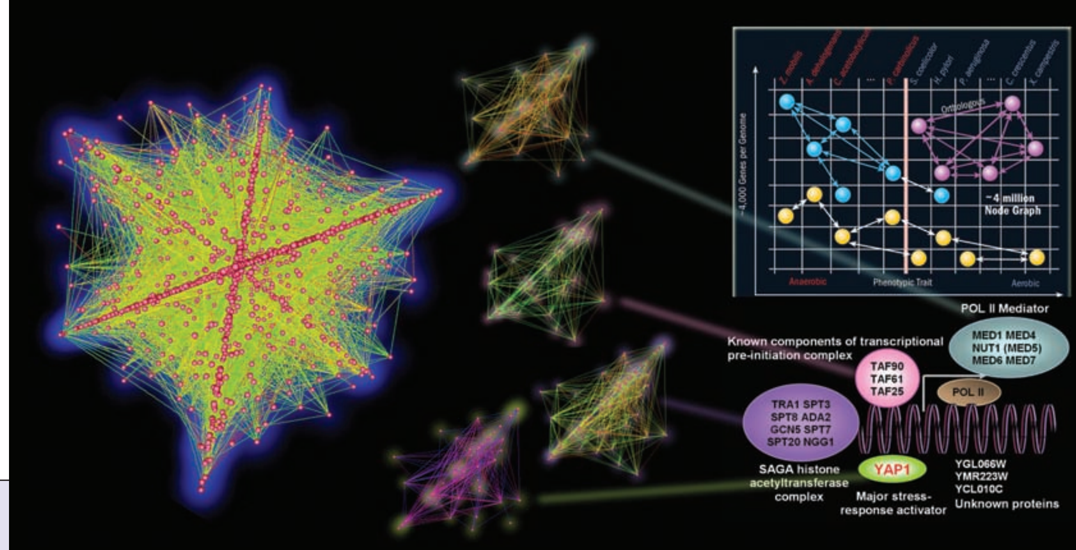

The cellular machinery involved in stress-induced transcriptional reprogramming of yeast cells discovered by a graph-theoretical computational framework developed by Dr. Samatova's research lab

to condense microarray analysis expertise into an intelligent, automated system. They developed and implemented a computational model of microarray quality that approximates expert opinion, and which allows users to perform automated quality judgment. Active research directions include computational analysis of alternative splicing and algorithms to investigate gene order permutations.

"Models that predict the behavior of complex biological systems form the Holy Grail of systems biology"

Connecting the dots. While analytical tools that derive the components from high-throughput experimental data significantly reduce the amount of data and increase the quality, the challenge still remains of how to "connect the dots," that is, to construct predictive *in silico* models of these biological systems. The combinatorial space of feasible solutions is enormous, and advanced methods for constraining such a space and for efficient search of optimal and biologically meaningful solutions are in great demand.

Dr. Samatova's research lab is developing advanced algorithms for modeling and comparative analysis of biological networks. Biological networks can be mathematically represented as graphs where nodes might be genes or metabolites and the edges represent some kind of relationship between them. The types of relationships might be regulation or physical interaction. Questions about these biological networks could then be translated to problems on graphs. For example, the question of identifying all protein complexes can be reduced to the problem of enumerating and merging maximal cliques in the network of pairwise protein interactions.

While the idea of leveraging graph theory and graph algorithms sounds attractive, in reality, most of these algorithms fail for large-scale biological networks. For example, the question of finding critical genes for aerobic microbial growth may require running clique enumeration algorithms on graphs with millions of nodes. The difficulty lies in the exponential time complexity of these algorithms with the graph size and the lack of scalable parallel implementations of these algorithms on high performance computing systems. By advancing the theory of combinatorial search space reduction and developing a library of scalable parallel graph algorithms, Dr. Samatova's research lab tackles many important large-scale biological problems related to bioenergy production and bioremediation.

Visualizing our Changing World

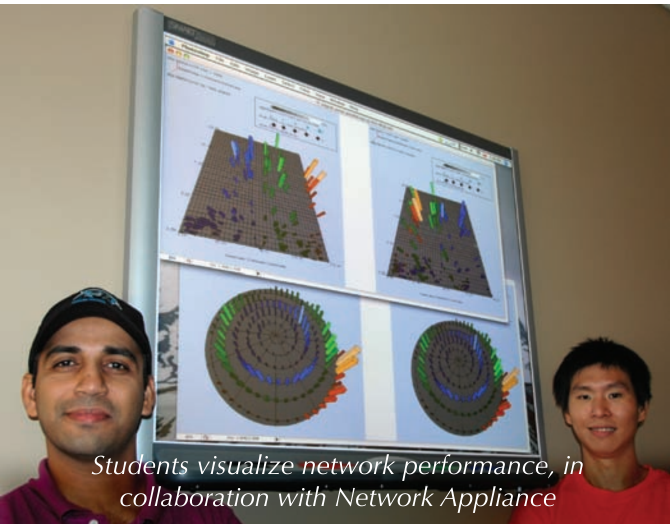

Students visualize network performance, in collaboration with Network Appliance

An important research area in computer graphics is visualization, which converts collections of strings and numbers (or datasets) into images that viewers can use to explore, discover, and analyze. Rapid growth in our ability to generate, capture, and archive vast amounts of information has increased the need for effective visualization techniques. Unfortunately, methods to display information in useful and meaningful ways have not always kept pace.

Displays themselves fall short, offering only a few million pixels, while today's data sets have grown to include billions of elements. Imagine instead that your walls are displays showing detail sharper than your eye can see. With billions of pixels, these displays can show every element of a massive data set, render text that is as crisp as the printed page, and in a video conference, reveal every flicker of emotion on a colleague's face.

> **"Imagine a world where your walls are displays showing detail sharper than your eye can see"**

Unless we make fundamental changes to our imaging technologies, we will not see these displays in our lifetimes. We are realizing these changes by exploiting our knowledge of the human visual system. Computer images are traditionally formed into frames: a rectangular array of color pixels, each representing the same moment. By discarding the notion of a frame, we can generate high quality images that match human sensitivities to both spatial and temporal change, at

Given the same rendering speed, frameless imagery (center) looks better than framed imagery generated at 60 Hz (left), and is more up to date than framed imagery generated at the same spatial resolution (right)

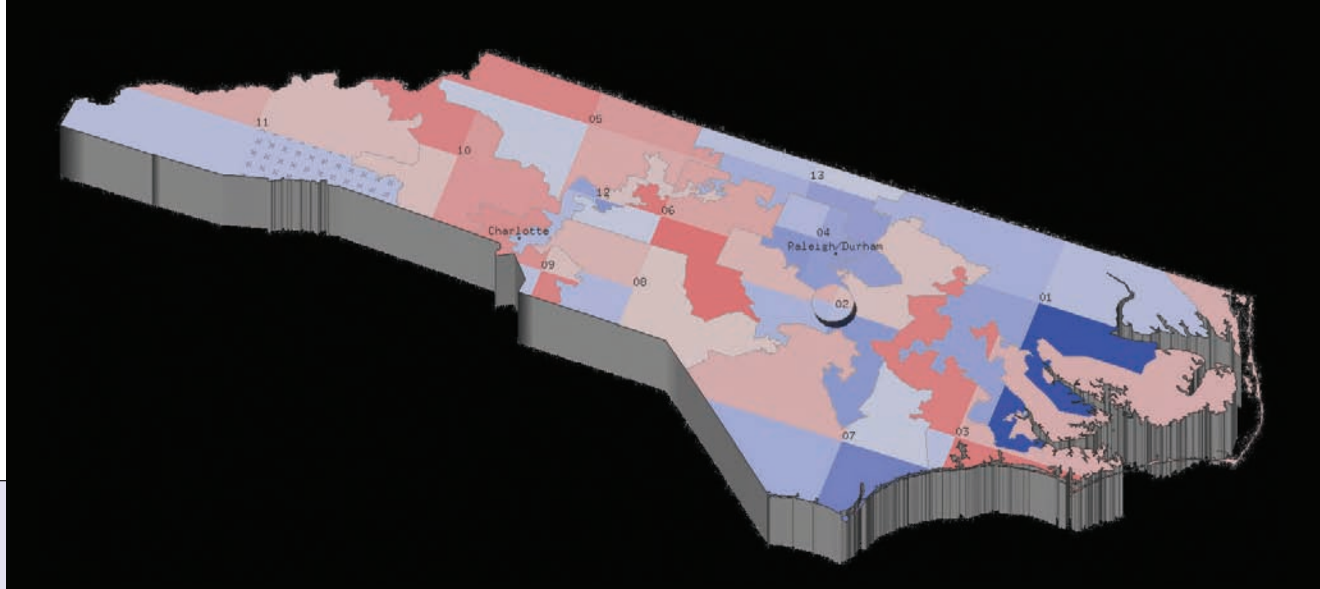

Visualizing President, U.S. Senate, U.S. House, and state Governor results by congressional district in North Carolina; color represents winning party, and saturation represents winning percentage

ten times traditional rates. For further speedups, we embed higher level visual elements such as edges and textures into our imagery.

> **"Although it is common to refer to a state as *red* or *blue*, this simple dichotomy can often be misleading"**

A second project visualizes U.S. election results. Understanding elections has always interested the voting public, even more-so given the recent close races and shifting demographic throughout the United States. Although it is now common to refer to a state as "red" or "blue" depending on whether a majority of its participants choose Republican or Democratic candidates, this simple dichotomy can often be misleading. Consider North Dakota, a "red" state that voted Republican in its Presidential and state Governor elections, but Democratic for the U.S. Senate and U.S. House, or Maine, a "blue" state that voted Democratic for President, U.S. House, and state Governor, but Republican for the U.S. Senate.

Variations in election patterns occur not only within a single election, but also across different elections. This raises the question: do common groups of individuals tend to vote in similar ways for different elected offices? That is, can groups of voters really be characterized as "red" or "blue", or do their choices vary based on the candidates running and the positions being filled?

Dr. Healey tabulated and visualized the 2004 Presidential, the most recent U.S. Senate, the 2006 U.S. House, and the most recent state Governor elections. Voters are divided by congressional district. For all 435 districts spread across the 50 United States, the candidate the district's voters selected in each election is identified. The district is subdivided into four quadrants representing the four elections of interest: President (upper-left), U.S. Senate (upper-right), U.S. House (lower-right), and Governor (lower-left). Within each quadrant, hue identifies the winning candidate's party: red for Republican, blue for Democrat, and green for Independent. Saturation visualizes winning percentage: more saturated hues are shown for higher percentages. A texture pattern of small X's is displayed to highlight incumbent party losses. Finally, state-wide totals are presented in a small disc floating over the state. The disc is subdivided into the same quadrants as the congressional districts.

SILO–Designing for Change in the Global Internet

How well will the Internet hold up in the next year, the next decade, or quarter century? *Will* it hold up?

The Internet, conceived in the era of mainframe computers and 56Kbps links, has evolved into a complex world-girdling system of importance equal to that of the power grid and the transportation infrastructure. The Internet has changed every aspect of our lives in the past few decades, and has itself changed nearly beyond recognition in the same time. It has gone from an academic curiosity, to a global commercial enabler, to a universal social nexus. Next, the incorporation of small sensors and actuators, and the constant access to such devices provided by networking, may change our world. Again.

Nevertheless, there is widespread perception in the networking community that key limitations of the Internet's design might be bringing it close to a breakdown point, and evolution is urgently needed. In 2005, the National Science Foundation created an initiative to fund clean-slate design research on the Internet. The SILO project is one of those funded by this initiative.

> **"The SILO philosophy: transport in layers, control across layers"**

The main focus of SILO is to create a framework in which evolving networking services can be seamlessly and flexibly introduced. The objective is not to design the next system or even the best next system, but a system that can sustain continuing change. The SILO framework consists of (1) reusable building blocks of fine-grain functionality, (2) explicit support for combining elemental blocks to accomplish highly configurable complex communication tasks, and (3) control elements to facilitate

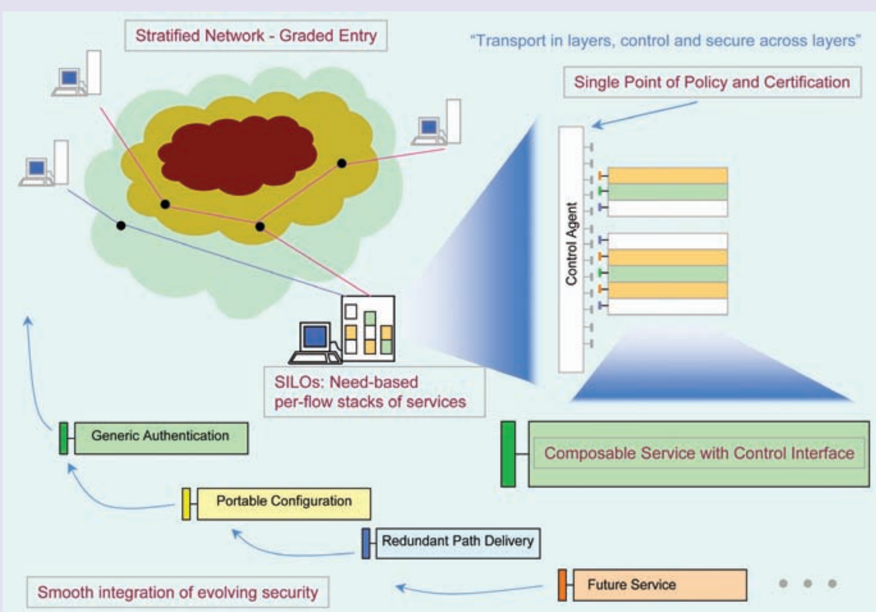

A diagram of the proposed SILO architecture what is currently referred to as cross-layer interactions.

"The objective is not to design the next system or even the best next system, but a system that can sustain continuing change"

Fundamentally, the SILO architecture generalizes the concept of layering. The building block is a *service*, which takes the place of a protocol layer. Like a protocol layer, it presents a data interface to a served (upper) and serving (lower) service (layer), but in addition, it provides (1) a control interface, which communicates with a unified control agent, and (2) a set of rules for composability, which states what other services this service may be composed with, and in what relation. Because the framework does not limit the services that may be presented to the *control agent* for composition, incorporating new services (e.g. those reflecting an evolving security policy) is seamlessly supported by the SILO architecture. The control agent also acts as repository for security policies and certificates, and can provide a first level of function verification services to support security. The SILO approach can be viewed as "transport in layers, control across layers."

In such an architecture, composing the services that make up the software layers for specific data flow requirements emerges as an essential part of the architectural system. We have provided a minimal set of precedence constraints to express service interactions, and an algorithm that obtains correct compositions under this set. Our ongoing research efforts address cross-service tuning algorithms, an ontology and underlying theory for expressing the effects of control knobs, and other related issues. We are also building a prototype to demonstrate basic service composition and tuning functions. Application Programming Interfaces, as well as the current ontology, are available from the project website.


The SILO team (from left to right):
Dr. Ilia Baldine (Renaissance Computing Center), Anjing Wang, Dr. George Rouskas, Manoj Vellal, Dan Stevenson (RTI International), Dr. Rudra Dutta

Compliance in Information Systems

Dr. Antón testifies to Congress

Privacy is important because it helps us maintain our individuality, autonomy and freedom of choice. Properly protecting the privacy of information is in all our best interests because once our privacy is lost is can seldom be fully recovered. To properly protect this information, systems must be designed holistically within the broader regulatory and legal compliance context. Researchers at The Privacy Place research group are developing methods and tools to help software engineers build systems that keep sensitive data secure and private, while ensuring regulatory and legal compliance.

Privacy laws require companies to enforce their policies and consumers are increasingly concerned about privacy violations. In addition, companies are increasingly being held accountable for their privacy practices. However, machine-readable and machine-enforceable polices are needed to consistently apply privacy practices and prevent breaches. To this end, Drs. Annie I. Antón and Ting Yu are collaborating with colleagues at Purdue University to develop a privacy policy lifecycle framework to tackle this problem.

> **"Properly protecting privacy is in our best interests; once lost, it can seldom be fully recovered"**

JetBlue Airways. Understanding data breaches in information systems is complex because each data breach involves numerous stakeholders adhering to different policies and governed by various laws. In 2003, JetBlue Airways gave five million customers' travel records, in violation of its privacy policy, to a US Department of Defense contractor. The JetBlue case is complex because of the large number of parties (actors) involved. Researchers at NCSU's The Privacy Place modeled this complexity by focusing on the actors, the actual information each actor obtained and used, and whether each actor published privacy policies on their Web sites. By

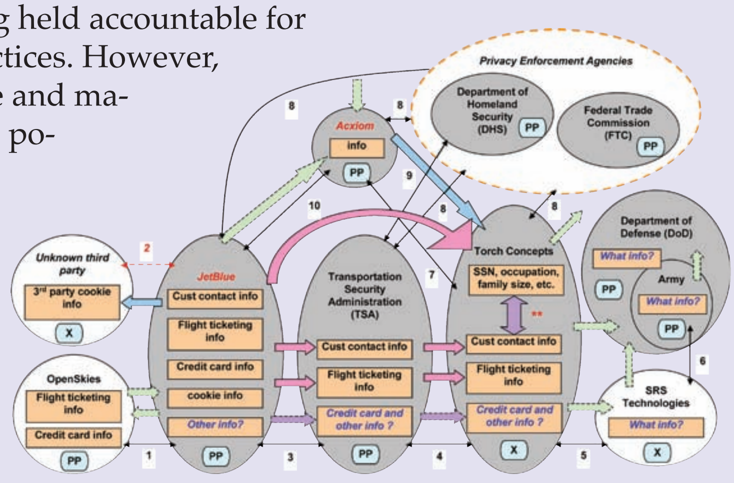

JetBlue information transfer and privacy policy violation model, actors' (shaded) relationship and information flow model depict the handling of sensitive customer information modeling the contractual relationships in this information transfer and privacy policy violation case, they revealed vulnerabilities that resulted in unfortunate privacy breaches and two lawsuits. JetBlue's privacy-policy violation case is quite complex, involving many actors and resulting in apparently unintended consequences. This study caught the attention of the Chief Privacy Officer of the U.S. Department of Homeland Security. Dr. Antón now serves on the U.S. DHS's Data Privacy and Integrity Advisory Committee.

ChoicePoint. In 2005, fraudulent parties posing as legitimate businesses accessed data-broker ChoicePoint's databases. By the end of 2005, ChoicePoint had notified roughly 163,000 victims that their personal information had been fraudulently accessed. Criminal investigators discovered more than 800 instances of identity theft where stolen data was used to access the personal information that ChoicePoint had stored. According to ChoicePoint, these security breaches eventually cost the company US $27.3 million in 2005 alone to cover legal fees, notify victims, and seek audits. The Privacy Place team of J.D./PhD student Paul Otto, Dr. Annie I. Antón, and Dr. David Baumer investigated this data breach and developed a set of recommendations for data brokers that have been presented at national meetings, resulting in an invitation to meet with the top executives at ChoicePoint in Alpharetta, Georgia. In addition, Dr. Antón recently testified before a congressional subcommittee on protecting the social security number from identity theft.

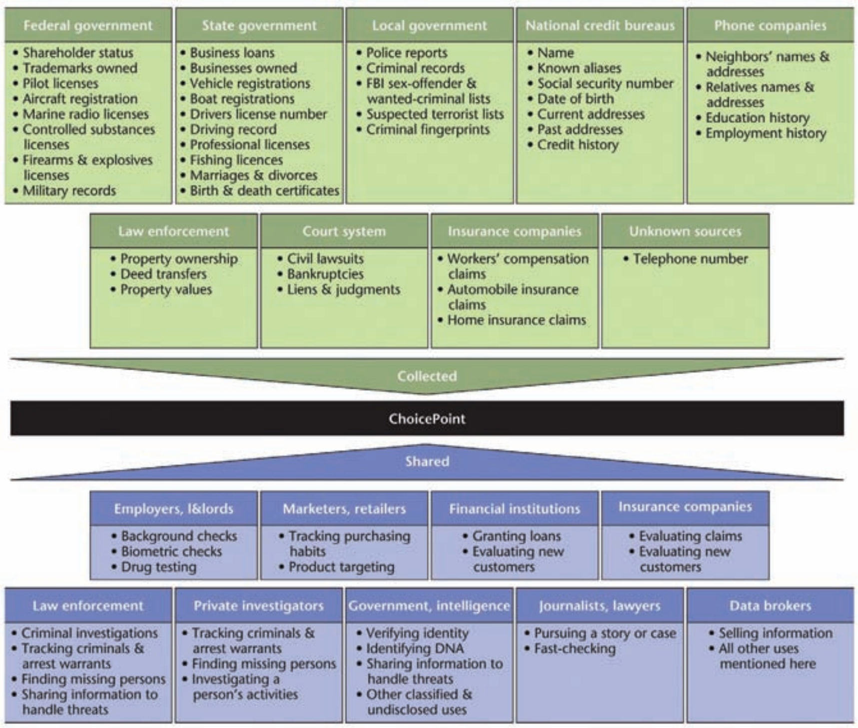

ChoicePoint's information flow: blue boxes represent data leaving the data broker, green boxes represent sources of data entering the company

Productively Building Trustworthy Software

Research interests of our software engineering group are diverse and range from software development lifecycle, to requirements engineering, to system and software testing and reliability, to software process and risk management, to open source issues. We distinguish ourselves through a focus on practical application and verification technologies. Our industrial and government partners and sponsors include ABB, ARO, Cisco, DARPA, DOE, EMC, Ericson, FDA, IBM, Intel, Microsoft, Nortel, NSA, NSF, ORNL, RedHat, SAS Institute, and Tekelec. We teach our students how to productively build software that can be *trusted by its direct and indirect users in all aspects important to the users*, and through our research we are advancing the frontiers of the techniques for doing so. The following illustrates our research interest areas. For more information we invite you to visit our web pages.

“We teach students to productively build software that can be trusted by both its direct and indirect users”

Process. Software organizations face intense competition to deliver software faster, cheaper, and of higher quality while being responsive to change. This intense competition is creating a crisis within software organizations and is driving them to reexamine their processes. For many organizations, the question has switched from, “Should we use an agile or a plan-driven methodology?” to “Which agile practices shall we integrate with our current process so we can be more responsive and have higher quality?” Dr. Williams’ research group has been successfully working with industrial teams to help them selectively transition to using a variety of agile software development practices. Empirical analyses indicate that small to medium-size teams improve their productivity, quality, and customer satisfaction through the use of these techniques.

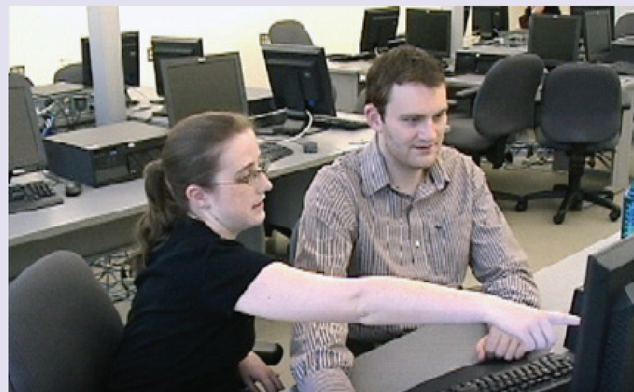

Sarah Heckman and Lucas Layman collaborate via pair programming

“The question has switched from ‘Should we’ to ‘How will we use agile practices?’”

Testing and automation. In addition to appropriate processes, automation and tools also play a very important role in software development. Both Dr. Williams's and Dr. Xie's research teams have been producing tools that address practical problems in software development, especially in the domain of software testing, software verification, and software reuse. Testing remains the most common and yet quite costly practical means of assuring the trustworthiness of software and a number of our tools focus on that problem. This includes test generation, test-oracle construction, regression testing tools, and tools that support security testing.

Reliability and security. Software reliability and security have been of interest to a number of researchers for over 20 years. Interests range from software reliability and availability, to fault-tolerance, to software security and safety, to the development of new static and dynamic metrics, to security implications of the software requirements, designs, and development practices. More recently Dr. Xie's group has been working on integrated tools for mining open source code to help developers discover application programming interface (API) usage patterns during software reuse and API properties for detecting bugs in and during software verification. Over the years, our researchers have delivered seminal work on generalized modeling of the relationships between software coverage testing and product reliability, and in the area of specification-based testing. For example, work by Dr. K.C. Tai has shown that automated predicate-based testing may be one of the most efficient and effective software testing approaches.

"NC State is poised to take a leadership position in open source software engineering"

Open source. In 2007, our department launched the Center for Open Software Engineering (COSE, pronounced "cozy") under direction of Dr. Williams. The mission of the center is to both practically and formally study open source issues. Open source is an area of information technology that already has radically altered the way we develop, market, and service software and the way the software industry collaborates in the 21st century. Open source software products play a critical role in the nation's economy and in supporting the nation's infrastructure. NC State is fortuitously situated in the midst of open source giants such as Red Hat and IBM, and is poised to take a leadership position in open source software engineering.

SGI Altix supercomputer at ORNL Center for Computational Sciences

End-to-End Storage for Scientists

Scientific computing is increasingly data-intensive. While there are currently many efforts on storage and I/O for supercomputers, not enough attention has been paid to supercomputing end users' storage problems in their local computing platforms, or the data movement and coordination between the two environments. As a result, scientists often spend more time managing and moving their data, than actually computing on remote or local machines. We, as members of the system research group in the Computer Science Department at North Carolina State University, have been exploring several solutions to this problem.

One of our major approaches is to aggregate idle storage spaces from participating workstations into shared storage with parallel I/O capability. More specifically, we designed and implemented a shared storage layer on top of donated idle disk spaces, which appears as a transparent, unified storage volume. Files stored in this scavenged space in a striped manner will appear as a single file stored on one disk. This will greatly improve the appeal of desktop parallel computing to both application developers and users, since idle storage resources can be pooled together into larger capacities. Compared to what is normally available—local desktop disk space and small quotas on shared file systems—this would provide a low-cost solution that improves the utilization of existing storage resources and answers the bursty but relatively short-spanned scientific data processing needs.

> **"Scientists often spend more time managing and moving their data, than actually computing"**

To this end, we have built FreeLoader, a research prototype for scavenging idle desktop storage resources in LAN settings and aggregating them into a shared cache and scratch space for scientific data. FreeLoader utilizes donated idle spaces attached to unreliable workstation nodes using techniques including soft-state registration and software striping. We have verified that (1) a single manager suffices to handle system state maintenance

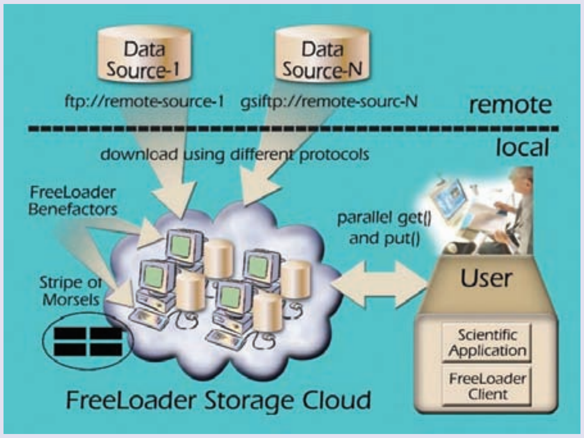

Envisioned FreeLoader distributed storage cache environment

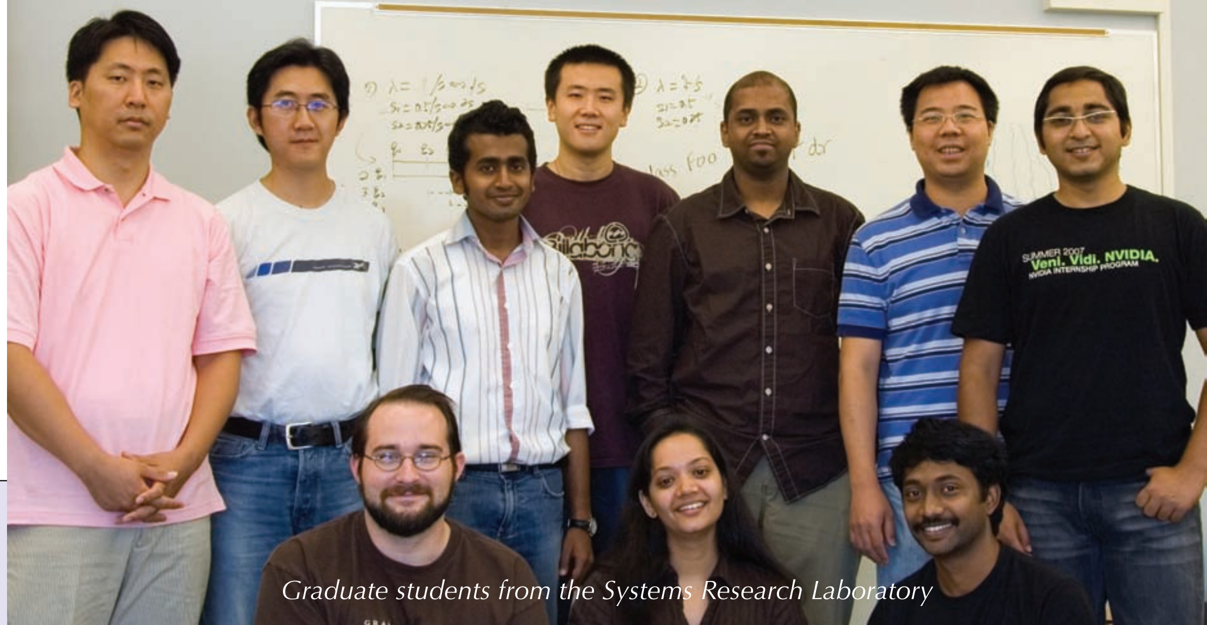

Graduate students from the Systems Research Laboratory

and metadata service, (2) file striping can deliver high data access rates in a desktop LAN environment with heterogeneous disk and network configurations, and (3) files distributed to donated workstations can be accessed transparently in scientific data processing applications through a set of wrapper interfaces exporting file system operations.

There are many interesting and innovative techniques that people can develop on top of platforms such as FreeLoader. For example, we experimented with combining the prefix caching technology used in multimedia systems, and the collective I/O technology used in high-performance parallel I/O, to partially cache data in FreeLoader. The uncached "tail" portion of each dataset can be retrieved on demand, with the transfer cost reduced by intelligent parallel downloading, and hidden from the user by the access time of the cached prefix.

In facilitating users' frequent data staging operations between their local computing environments and supercomputers, we built coordination between a supercomputing center's parallel file system and its batch job scheduler. Users may specify their data's staging-in and staging-out operations in a batch job script, and our enhanced job scheduler will automatically extract such operations when the job is submitted. These data operations are scheduled in a separate data queue, respecting the appropriate dependence relationship with the compute jobs to which they are associated.

"FreeLoader utilizes donated idle storage in LAN settings, aggregating it into a shared cache for scientific data"

This way, users' expensive data operations can be executed automatically, without tedious manual manipulation charges against precious computing allocations. Better, such a scheme allows for two additional benefits. First, automatically staged input files can bear additional file system metadata items that remember their remote source locations. Therefore, missing parts of those files can be transparently recovered from their sources when there are storage system failures. Second, automatic data staging and data job scheduling enables the entire scratch space, along with network resources at a computing center, to be optimized globally among many concurrent user jobs. Our prototype built in this research has enhanced the widely user Lustre file system and Moab scheduler, and is heading to production systems at Oak Ridge National Laboratory.

NP-Completeness in Practice: Theory and Experiments

Theoreticians study problems and algorithms from an abstract, structural point of view. This gives them a unique perspective, leading to the discovery of relationships among solutions in many diverse application areas. At NCSU an important aspect of this work centers on NP-completeness and its impact on practical problems.

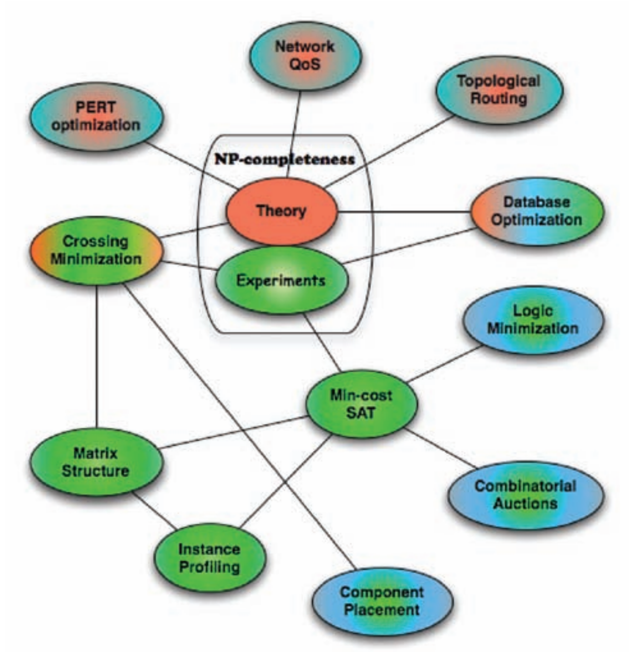

Connections among theoretical and experimental work related to NP-completeness

Since almost all the problems arising in application areas are NP-hard, the research focuses on finding easier special cases and on robust experimental evaluation of heuristics.

Collaboration with application-area specialists has led to the resolution of the complexity (polynomial-time versus NP-hard) of subtle variations of practical problems, most notably directed acyclic graph reduction (for PERT chart analysis), topological VLSI routing, and cost minimization for given quality of service guarantees (e.g., for network performance, reliability, and security).

Development of new heuristics is also essential, along with experimental work validating their performance, both runtime and solution quality. Application areas under experimental validation at NCSU include design automation (placement and routing, logic and state-machine minimization, technology mapping, and automated test pattern generation), database query optimization, data visualization, and combinatorial auctions.

"NP-hard problems require focus on easier special cases and robust experimental evaluation of heuristics"

Crossing minimization. Given a bipartite graph, arrange the nodes of each side along a horizontal line so that the number of edge crossings is minimized. Aside from the obvious application to drawing of diagrams, for example, PERT charts, UML diagrams, and relations in databases,

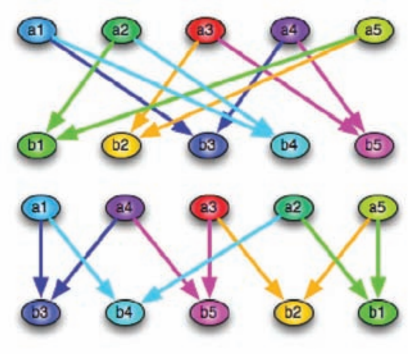

Minimizing graph crossings

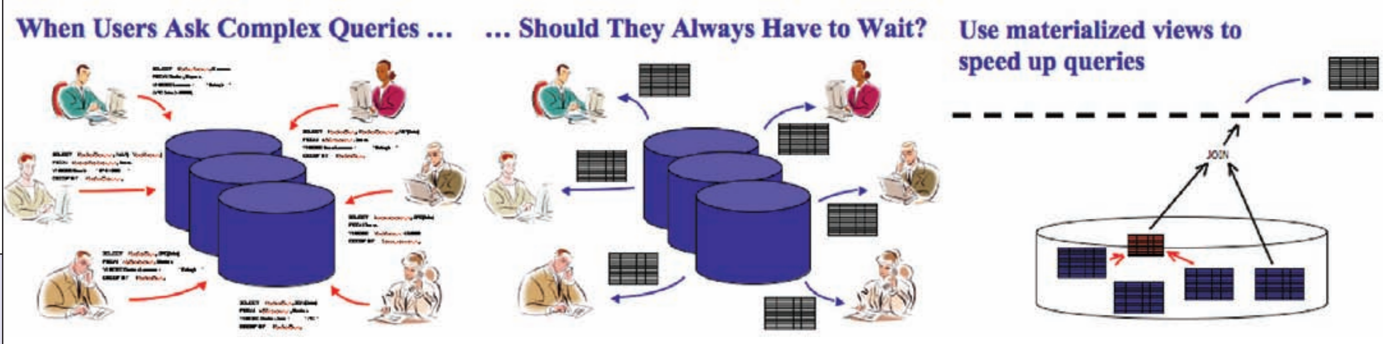

Materializing views to optimize database queries

there are also applications to VLSI layout, for example, placement of components in a circuit.

Visualization of large graphs. An important part of decision support is the ability to recognize patterns in large amounts of data. When that data takes the form of association diagrams that show relationships among products and services, the natural modeling tool is directed graphs. Traditional representations, even with crossing minimization, are hard to decipher. This work proposes a new perspective, in the form of quilts.

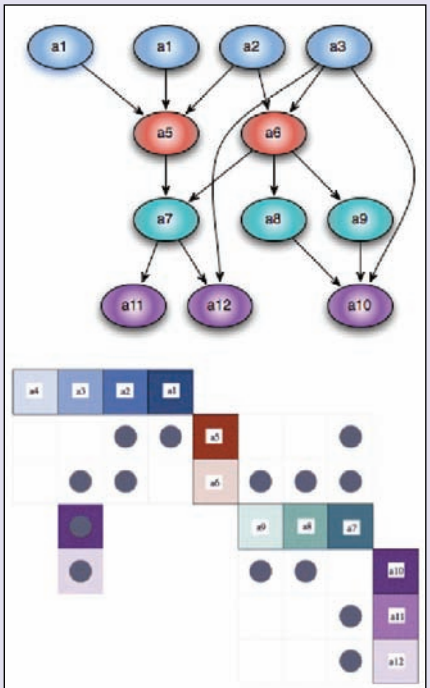

Visualizing large graph structure

View selection in databases. Database queries can be sped up significantly if relevant common parts are pre-computed as views. The goal of this work is to develop algorithms and heuristics for choosing optimal sets of views, those that satisfy a bound on the amount of storage available while improving the runtime of queries overall as much as possible.

Min-cost satisfiability (SAT). Many important problems in design automation center on minimization of basis logic elements in circuits, whether individual gates, states in finite-state machines, or off-the-shelf components (technology mapping). Min-cost SAT, minimization of the number of true variables to satisfy a formula, is an abstraction of these problems. Other applications include automatic test-pattern generation and combinatorial auctions.

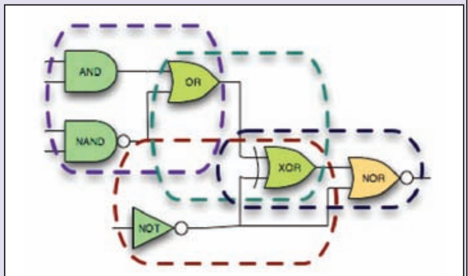

Technology mapping by min-cost SAT

Primary participants involved in these projects are: Dr. Franc Brglez, Dr. Rada Chirkova, Dr. Yahya Fathi (Operations Research), Dr. Jason Osborne (Statistics), Theresa-Marie Rhyne, Dr. Matthias Stallmann, Dr. Ben Watson, and the SAS Visualization Group.

Annie I. Antón, Associate Professor
PhD, Georgia Institute of Technology, 1997
Software engineering, requirements engineering, information privacy and security

Kemafor Anyanwu, Assistant Professor
PhD, University of Georgia, 2007
Semantic and services computing, bioinformatics, data and knowledge systems, data mining

Dennis R. Bahler, Associate Professor
PhD, University of Virginia, 1987
Artificial intelligence, machine learning, data mining, bioinformatics

Donald Bitzer
Distinguished University Research Professor
PhD, University of Illinois, 1960
High-speed networks, communications, bioinformatics, computer-based education

Franc Brglez, Research Professor
PhD, University of Colorado, 1970
Methods and infrastructure for reliable performance evaluation of combinatorial algorithms

Rada Y. Chirkova, Assistant Professor
PhD, Stanford University, 2002
Databases, computational logic

Edward W. Davis, Professor
PhD, University of Illinois, 1972
Computer architecture, parallel processing

Jon Doyle
SAS Professor of Computer Science
PhD, MIT, 1980
Artificial intelligence, mathematical foundations, rationality, knowledge discovery

Rudra Dutta, Associate Professor
PhD, North Carolina State University, 2001
Traffic grooming, fault tolerance, optical networks, ad-hoc wireless networking

Robert Fornaro, Professor
PhD, Pennsylvania State University, 1969
Concurrent programming in graphics, robotics, signal processing, operating systems

Vincent Freeh, Associate Professor
PhD, University of Arizona, 1996
Operating systems, compilers, programming languages, distributed and parallel computing

Robert E. Funderlic, Professor Emeritus
PhD, University of Tennessee, 1970
Scientific and parallel computing, numerical algebra, Google page rank, cluster analysis

Edward Gehringer, Associate Professor
PhD, Purdue University, 1979
Object-oriented software, parallel processing

Xiaohui Gu, Assistant Professor
PhD, University of Illinois, 2004
Distributed and operating systems, networks, autonomic computing, system mining

Khaled Harfoush, Assistant Professor
PhD, Boston University, 2002
End-to-end network diagnosis, topologies, routing, ad-hoc and peer-to-peer networks

Christopher G. Healey, Associate Professor
PhD, U. British Columbia, Canada, 1996
Visualization, computer graphics, perception

Steffen Heber, Assistant Professor
PhD, Universität Heidleberg, Germany, 2001
Computational biology, bioinformatics

Thomas L. Honeycutt, Associate Professor
PhD, North Carolina State University, 1969
MIS, modeling and simulation, computer literacy

S. Purushothaman Iyer, Professor
PhD, University of Utah, 1986
Formal methods for distributed and embedded systems, languages, concurrency

James C. Lester, Associate Professor
PhD, University of Texas, 1994
Artificial intelligence, intelligent tutoring systems, natural language processing

Xiaosong Ma, Assistant Professor
PhD, University of Illinois, 2003
High performance computing, parallel IO, storage systems, scientific data management

Richard Mayr, Assistant Professor
PhD, TU-Munich, Germany, 1998
Formal verification, logic, automata theory, term rewriting, tableau systems

David F. McAllister, Professor
PhD, UNC Chapel Hill, 1972
Computer graphics, speech processing

Faculty

Frank Mueller, Associate Professor
PhD, Florida State University, 1994
Parallel and distributed systems, compilers, real-time and embedded systems

Peng Ning, Associate Professor
PhD, George Mason University, 2001
Computer and network security, wireless security, intrusion detection, applied cryptography

Harry Perros
Distinguished Alumni Research Professor
PhD, Trinity College, Ireland, 1975
Networking, performance modelling, queueing theory, service science mgmt and engineering

Douglas S. Reeves, Professor
PhD, Pennsylvania State University, 1987
Network and software security, peer-to-peer computing

Injong Rhee, Associate Professor
PhD, UNC Chapel Hill, 1994
Computer networks, Internet protocol design, congestion control, multimedia networking

Robert D. Rodman, Professor
PhD, UCLA, 1973
Forensic speaker identification, automatic emotion detection, lip synchronization

George N. Rouskas, Professor
PhD, Georgia Institute of Technology, 1994
Network architectures and protocols, optical networks, performance evaluation

Nagiza Samatova, Associate Professor
PhD, CCAS, Russia, 1993
Computational biology, graph theory, scalable data analytics, data management

Carla D. Savage, Professor
PhD, University of Illinois, 1977
Enumeration, algorithms, combinatorics, discrete mathematics

Munindar P. Singh, Professor
PhD, University of Texas, 1993
Multiagent systems, trust, service-oriented computing, business protocols and processes

Robert St. Amant, Associate Professor
PhD, University of Massachusetts, 1996
HCI, artificial intelligence, cognitive systems, intelligent user interfaces

Matthias Stallmann, Professor
PhD, University of Colorado, 1982
Experimental algorithmics, combinatorial optimization, NP-hard problems, graph algorithms

William J. Stewart, Professor
PhD, Queen's U., Northern Ireland, 1974
Performance modeling, Markov chains, queueing theory and numerical linear algebra

Alan L. Tharp, Professor
PhD, Northwestern University, 1969
File structures, man-machine interfaces, databases, innovating in technology

David Thuente, Professor
PhD, University of Kansas, 1974
Communication system design, simulation, performance modeling, media access control

Mladen Vouk, Professor
PhD, King's College, United Kingdom, 1976
Software engineering, scientific and network-based computing, computer-assisted education

Benjamin Watson, Associate Professor
PhD, Georgia Institute of Technology, 1997
Computer graphics, design, interaction

Laurie Williams, Associate Professor
PhD, University of Utah, 2000
Pair programming, software security, agile software development, testing and reliability

Tao Xie, Assistant Professor
PhD, University of Washington, 2005
Software testing, verification, reuse, mining software engineering data

R. Michael Young, Associate Professor
PhD, University of Pittsburgh, 1997
Artificial intelligence, planning, natural language, interactive narrative, computer games

Ting Yu, Assistant Professor
PhD, University of Illinois, 2003
Security, trust management, privacy protection

Department Personnel and Alumni

Joyce Hatch, former Director of Advising, retired in 2006 after 30 years of service with the department

Administration

- **Dr. Barbara Adams**, Director of Advising
- **Dr. Dennis Bahler**, Director of Undergraduate Programs
- **Jason Corley**, Information Technology Manager
- **Ron Hartis**, Director of Operations
- **Linda Honeycutt**, Department Office Manager, HR and Executive Assistant
- **Ann Hunt**, Finance Director
- **Dennis H. Kekas, PE**, Director, Networking Technology Institute
- **Dana Lasher**, Lecturer and Scheduling Officer
- **Theresa-Marie Rhyne**, Director, Center for Visualization and Analytics
- **Missy Seate**, Contracts and Grants Manager
- **Ken Tate**, Director, Development and External Relations
- **Dr. David Thuente**, Director of Graduate Programs
- **Dr. Mladen Vouk**, Department Head

Staff

- **Ginny Adams**, Administrative Support Specialist
- **Carol Allen**, Administrative Support
- **John Bass**, Centennial Networking Lab Technical Director
- **Carlos Benavente**, System Programmer
- **Dare Cook**, Admissions Specialist and Counselor
- **Marhn Fullmer**, Networking Laboratories Manager
- **Michelle Healey**, CACC Administrator
- **Margaret Heil**, Associate Director, Senior Design Center
- **Jason Manners**, Lecturer
- **Carolyn Miller**, Lecturer
- **Trey Murdoch**, Operations and Systems Analyst
- **Tom Nelson**, Lecturer
- **Margery Page**, Graduate Program Manager
- **Susan Peaslee**, Accounting Clerk
- **Aaron Peeler**, Program Manager, IT Advanced Academic Computing Initiatives

Adjunct Faculty

- **Suzanne Balik**, Adj. Faculty
- **Dr. Charles Coleman, Jr.**, Adj. Asst. Professor
- **Dr. Kenneth Flurchick**, Adj. Asst. Professor
- **Dr. Brand Fortner**, Adj. Professor
- **Dr. Jaewoo Kang**, Adj. Asst. Professor
- **Dr. Scott Klasky**, Adj. Asst. Professor
- **Dr. Daniel Reed**, Adj. Professor
- **Dr. Andrew Rindos**, Adj. Asst. Professor
- **Dr. Xiaogang Wang**, Adj. Asst. Professor
- **Dr. Peter Wurman**, Adj. Assoc. Professor
- **Dr. Jun Xu**, Adj. Asst. Professor
- **Sammie Carter**, Adj. Lecturer
- **Dr. Eric Sills**, Adj. Asst. Professor

Associate Faculty

- **Dr. John W. Baugh**, Civil Engineering
- **Dr. Gregory T. Byrd**, ECE
- **Dr. Thomas M. Conte**, ECE
- **Dr. Alexander G. Dean**, ECE
- **Dr. Michael Devetsikiotis**, ECE
- **Dr. Erich Kaltofen**, Mathematics
- **Dr. Gianluca Lazzi**, ECE
- **Dr. Carl Meyer**, Mathematics
- **Dr. Thomas K. Miller**, DELTA
- **Dr. Michael Rappa**, Advanced Analytics Institute
- **Dr. Eric Rotenberg**, ECE
- **Dr. Jeffrey S. Scroggs**, Mathematics
- **Dr. Mihail Sichitiu**, ECE
- **Dr. Wesley E. Snyder**, ECE
- **Dr. Yan Solihin**, ECE
- **Dr. John MacKenzie**, Microbiology
- **Dr. Anna Stomp**, MRUFR
- **Dr. Yannis Viniotis**, ECE

Strategic Advisory Board

- **Nicholas Bowen**, IBM
- **Sidd Chopra**[1], Analytrix, LLC
- **Wayne Clark** (Chairperson), Cisco Systems
- **Keith Collins**[1] (Member Emeritus), SAS Institute
- **Christopher Crump**[1], Pacific West Financial
- **Jesse Fearrington**[1] (Vice-chairperson), retired, Wachovia Bank

1 Computer Science Department alumnus

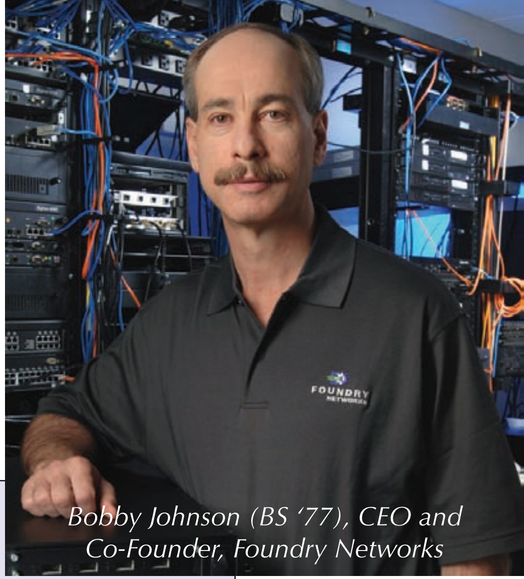

Bobby Johnson (BS '77), CEO and Co-Founder, Foundry Networks

- **Dennis Flynn**, The Sonar Group
- **Ken Hibbard**, Network Appliance
- **Vivien Joklik**, John Deere
- **Gopal Kakivaya**[1], Microsoft Corporation
- **Gail Kramer**[1], SAS Institute
- **Kathy Markham**[1], Kindred Healthcare, Inc.
- **Rudy Puryear**[1], Bain and Co.
- **William Riddick**[1], Computer Service Partners
- **Juan Vargas**, Google, Inc.
- **Eric Wagner**, EMC Corporation
- **William Weiss**[1], The Promar Group, LLC
- **Steve Worth**[1] (Member Emeritus), EMC Corporation
- **Mark Wyatt**[1], Duke Energy

Selected Distinguished Alumni

- **Marshall Brain** (MS '89), Founder of HowStuffWorks
- **Keith Collins** (BS '82), Chief Technology Officer, SAS Institute; NC State Distinguished Engineering Alumnus
- **Jim Davis** (BS '80), Sr. VP and Chief Marketing Office, SAS Institute
- **Robert Eason** (BS '72), President of QVS Software
- **Gary Funck** (BS '72), President and Co-founder, Intrepid, Inc.
- **Suzanne Gordon** (MS, '80), Chief Information Officer, SAS Institute; Member, NC State Board of Trustees
- **Wayne Harvey** (BS '94), VP, Chief Technology Office and Co-founder, Vicious Cycle Software
- **Gent Hito** (BS '94), President and CEO, n software
- **Dr. Larry Hodges** (MS '82), PhD '88), Professor and Chair, Department of Computer Science, UNC Charlotte
- **Richard Holcomb** (MS '89), CEO and Co-founder, StrikeIron, Inc.
- **Bobby Johnson** (BS '77), CEO and Co-founder, Foundry Networks
- **Richard Krueger** (MS '89), creator of xRes (now Macromedia Fireworks); President and CEO, Skinux, Inc.
- **Kathy Markham** (BS '80), VP of Information Systems, Kindred Healthcare, Inc.
- **Dr. Bill McKinnon** (PhD '97), Founder of ChanneLogics
- **David McPeters** (MS '82), VP of Information Technology, Deltacom, Inc.
- **Derek Meyer** (BS '79), FedEx Pilot; consultant, F/A-22 Raptor Air Combat Simulator
- **Dr. Elizabeth Mynatt** (BS '88), Associate Professor and Director, GVU Center at Georgia Tech; Internationally recognized expert in ubiquitous computing and assistive technologies
- **Dr. Raif Onveral** (MS '85, PhD '87), Co-founder of Orologic, Inc.; CEO of companies that include Litchfield Communications and ORD Bridge
- **Gerhard Pilcher** (BS '85), President of HB Rowe & Company
- **Rudy Puryear** (BS, '74), Director of IT Practice, Bain & Co.; Consulting Magazine's Top 25 Consultants in the World
- **Bill Riddick** (BS '74), President of Computer Service Partners
- **Ross Scroggs** (BS '71), Retired President and Founder, Telamon, Inc.
- **Matthew Squire** (PhD '95), Chief Technology Officer, Hatteras Networks
- **Dr. Dorothy C. Strickland** (PhD '95), Internationally recognized expert in virtual reality; President of Virtual Reality Aids, Inc
- **Troy Tolle** (MS '00), Chief Technology Officer and Co-founder, DigitalChalk, Inc
- **Erik Troan** (BS '95), Original Co-author of Red Hat Package Manager; Co-author of "Linux Application Development"; Chief Technology Officer, rPath
- **Kristopher Tyra** (BS '86), Co-founder, HiddenMind Technology, Inc.
- **William Weiss** (BS, '76), Chairman and CEO, The Promar Group, LLC
- **Ed Whitehorne** (BS '71), Successful entrepreneur; President, CI Partners, LLC
- **Josh Whiton** (BS '04), CEO and Co-founder, TransLoc, Inc.
- **Mark Wyatt** (BS '80), VP of Information Technology, Duke Energy

CISCO EMC² Fidelity INVESTMENTS Smart move. IBM JOHN DEERE

Corporate Partnerships

Since its launch in 2000, the ePartners Program has served as the department's cornerstone corporate relations program, providing a framework for developing and nurturing strong collaborative partnerships with the global business community.

Under the guidance of Ken Tate, Director of Development and External Relations, more than 60 companies are now actively engaged in partnership opportunities with the department. In addition to recognizing our corporate partners, the ePartners Program fosters ongoing communications and interaction with students and faculty, and allows our corporate partners to play a meaningful role in shaping the department's future direction.

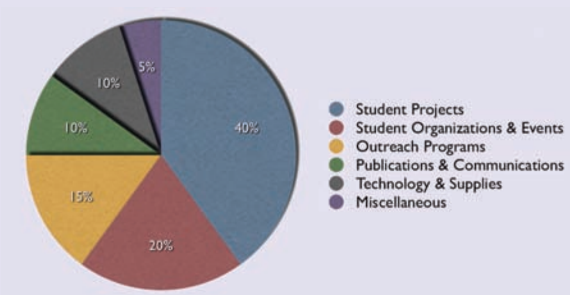

Typical ePartners program funding allocation

In recent years, unrestricted cash donations and awards provided by our corporate partners have exceeded $500,000 and overall support including software and equipment gifts-in-kind exceeds $4 million annually.

Funding and support provided by our partners are allocated across a wide range of strategic needs and initiatives within the department including student projects, minority and women outreach programs, student organizations and special events, department publications and communications, and learning and research technologies and supplies. In aggregate, our ePartners Program provides essential support which allows our department to continue to grow in emerging areas of computer science technology, while providing the highest quality education possible for our students.

One of the most meaningful and tangible ways corporations support our students is through the ongoing sponsorship of projects through our award-winning Senior Design Center (SDC). Since the center's launch in 1994, well over 1,000 students have participated on project teams, and corporate project sponsorships have topped $1 million. In recent years, SDC student teams have claimed international acclaim as the only U.S. school to win the CSIDC competition, and the only school in the world to win two years in a row. SDC project teams, in partnership with Insight Racing and with the support of corporate sponsors such as Tekelec, and involvement of the Advanced Vehicle Research Center (AVRC) and Lotus Engineering, have played a key role in the successful entry of an autonomous vehicle in the DARPA Urban Challenge.

Corporations and professional associations have led the way in providing much needed scholarships over

Microsoft NetApp Progress Energy sas TEKELEC

Including super ePartners Cisco Systems, Inc., EMC Corporation, Fidelity Investments, IBM, John Deere, Microsoft Corporation, Network Appliance, Progress Energy, SAS Institute, and Tekelec, the department enjoys strong and active partnerships with more than 60 companies, organizations, and agencies

the years. Our students benefit from substantial scholarship awards from Progress Energy, Nortel Networks, SAS Institute, Northrop Grumman, NC-SIM, and the Raleigh Chapter of ISSA, just to name a few. Some of the scholarship awards are truly transforming in nature. For instance, NC State was selected by Cisco Systems as one of just five schools nationwide to participate in the prestigious National Cisco Internet Generation Scholarship (NCIGS) awards program, providing $5,000 annual awards and internships to targeted incoming freshmen in an effort to increase the attraction and retention rates of underrepresented students (minority and female students) in the field of computer science. The inaugural awards were made to two freshmen in the fall of 2006 and both indicated that the awards played a significant role in their decision to study computer science at NC State.

As corporate partnerships mature, they often reach well beyond scholarships and student projects, as evidenced by Fidelity Investments' sponsorship of a premiere executive speakers series, our SAS Institute's support for a named professorship, IBM's involvement and support of our Service Sciences curriculum, and EMC's and NetApp's support of a new data storage curriculum. Corporate involvement and support is also an integral part of the launch of new centers. The creation of the new Digital Games Research Center (DGRC), with its multi-disciplinary focus on the sci-

The Raleigh Chapter of the Information Systems Security Association creates an endowment for student scholarships

entific, design, social, and educational challenges of design and construction of games and game technologies, was created in large part because of strong relationships with gaming companies such as Destineer Studios, Virtual Heroes, Emergent Game Technologies, Vicious Cycle, Electronic Arts, Epic Games, and Red Storm Entertainment.

Often, these partnerships result in truly outstanding results. For instance, IBM and NC State have partnered in the launch of the Virtual Computing Initiative (VCI), which will make tools and resources available to students at all educational levels to build 21st century skills across the state of North Carolina. The Computerworld Honors Program recognized NC State University as a 2007 Laureate in the category of education for its Virtual Computing Laboratory (VCL) project.

To all our corporate partners who have supported our department over the years through your gifts, your time, your insight, and ongoing engagement, we extend our deepest appreciation.

Department of Computer Science
North Carolina State University
890 Oval Drive
Raleigh, NC 27695-8206
Phone: 919.515.2858   Fax: 919.515.7896

http://www.csc.ncsu.edu