CS@CU

NEWSLETTER OF THE DEPARTMENT OF COMPUTER SCIENCE AT COLUMBIA UNIVERSITY VOL.4 NO.2 SPRING 2008

Shree Nayar

Elected to National Academy of Engineering

Shree Nayar, T. C. Chang Professor of Computer Science, has been elected to the National Academy of Engineering, one of the highest professional distinctions awarded an engineer. Academy membership honors those who have made outstanding contributions to "engineering research, practice, or education" and to the "pioneering of new and developing fields of technology, making major advancements in traditional fields of engineering, or developing/implementing innovative approaches to engineering education."

In announcing his election, NAE President Charles M. Vest noted that Professor Nayar was being recognized "for the development of computational cameras and physics-based models for computer vision and computer graphics."

As co-director of the Columbia Vision and Graphics Center and head of the Computer Vision Laboratory (CAVE), Professor Nayar's research efforts are focused on developing advanced computer vision systems. His laboratory has developed a new class of imaging systems called computational cameras. While the traditional world of cameras has its roots in the pinhole camera (camera obscura), Nayar's computational cameras use light rays from a myriad of sources, not just those that pass through the pinhole, or lens, of the modern camera. His computational cameras overcome several fundamental limitations of today's cameras, including field of view, dynamic range and spectrum. Nayar's inventions have been widely adopted by industry to address applications in digital photography, computer vision and computer graphics.

Professor Nayar, a native of Bangalore, India, received a B.S. degree in electrical and computer engineering from the Birla Institute of Technology and his Ph.D. degree in electrical and computer engineering from the Robotics Institute at Carnegie Mellon University. He has been on the faculty of The Fu Foundation School of Engineering and Applied Science since 1991 and was named to the T.C. Chang Professorship in 2002. His other prizes and awards include the David Marr prize twice, in 1990 and in 1995, the David and Lucille Packard Fellowship, National Young Investigator Award, NTT Distinguished Scientific Achievement Award, Keck Foundation Award for Excellence in Teaching and the Columbia Great Teacher Award. He has published more than 100 scientific papers and has more than 30 patents in imaging, vision and robotics.

An Interview with New Faculty Members

The Department of Computer Science is delighted to welcome two new junior faculty members, **Simha Sethumadhavan** and **Junfeng Yang**.

CUCS Newsletter roving reporter **Sean White** spoke with these new faculty members about their paths to computer science, their research motivations and challenges, and their visions of the future.

Junfeng Yang

How did you first become interested in computer science?

I became interested in Computer Science during my junior high school years. Back then, my school created a computer science course, which gave students a brief introduction to computers and taught them how to program in BASIC. I took the course and was thrilled by what I could do with computers. I liked that you could really build things. I became interested in Computer Science from then on and decided to pursue a career in Computer Science.

I found my interest in systems research grew gradually over the years. I'm more of an experimentalist. I enjoy learning the mechanisms of how fundamental systems work. I also enjoy building systems and applying them to real-world applications, so systems research is a good match for me.

How do you like to describe your field?

My broad area [of interest] is systems. I publish mostly in systems conferences like OSDI and SOSP. My past research has focused on practical software bug finding. Systems research used to center around performance–basically, how to make things fast. For the past few years, systems research has shifted to a variety of areas; one main area is software bug finding. My PhD started roughly at the time the area of bug finding became active in the systems community. I've been fortunate enough to work on the edge of this field.

As system researchers, we tend to focus more on the effectiveness of the checking tools we build. We care about how many bugs our tools find, how few the false positives are, and whether our tools can scale to large software, as large as the Linux kernel with several million lines of code. In order to get good results, we draw techniques from the field of program analysis as needed and adapt them whenever we see fit. We sometimes sacrifice the soundness or completeness of the program analysis techniques we use.

What motivated you to focus on your current research?

Software error is a real problem. Large, complex software systems inevitably have errors, and will continue to do so at least in the foreseeable future. These errors make computer systems unreliable and insecure. According to a study by the National Institute of Standards and Technology (NIST) in 2002, software bugs are estimated to cost the U.S. economy $59.5 billion annually. We have all seen enough headlines on software security holes.

One way to alleviate this problem is to provide programmers effective bug-finding tools, so that they can proactively detect bugs early in the development stages (e.g. coding and testing), before they ship their software. That's what motivated me to work on the field of building bug-finding tools.

How do you pick your research topics and what do you believe makes a good research question?

I want my research to be relevant to the world. I want to solve real problems that cause people to suffer. Of course, the research problems I work on must also interest me.

Do you have a method or approach that you like to use in your research?

My approach to research is experimental–I'm driven by practice, out of which the problems that matter most to me emerge. In my research on bug-finding, I focus on errors that can cause serious consequences (e.g., data-loss or security exploits). To find these errors, I build tools. I explicitly design tools to be effective and easy to apply so that system builders will use them. I focus on results. To achieve good results, I am willing to draw techniques from different fields, and adapt them whenever necessary.

What do you consider the most interesting and exciting challenges of your current research?

The EXPLODE approach has been successful in checking storage systems. I'm currently working on applying it to distributed systems. There are two main challenges. First, how do we deal with the effectively infinite state space defined by a real-world distributed system? I'm developing heuristics to search the space more effectively, to find more bugs and find bugs more quickly. Second, how do we make model checking even easier to use? Model checking is a heavyweight formal verification technique. My previous work has made it easier to use. I'd like to make it even more lightweight, ideally, push-button.

Do you have a favorite project that you've worked on?

My favorite project is EXPLODE, a general, lightweight system for finding storage system errors using model checking. This formal verification technique finds corner-case errors by comprehensively exploring the state spaces defined by a system. Storage systems have two dynamics that make them attractive for such an approach. First, their errors are some of the more serious ones, since they can destroy persistent data and lead to unrecoverable corruption. Second, traditional testing requires an impractical exponential number of test cases to determine if a system will recover if it crashes at any point during execution. Naive application of model checking to large storage systems is, however, prohibitive because it requires huge manual overhead. EXPLODE uses a novel in-situ checking architecture that drastically reduces the manual overhead needed for checking. In the best case, clients write a few hundred lines of C++ code to check a new storage system.

I applied EXPLODE to a broad range of 17 real-world storage systems (without requiring source code): three version control systems, an open source database library Berkeley DB, the Linux NFS V3 implementation, ten Linux file systems, a RAID system, and the popular VMware GSX virtual machine. EXPLODE found serious bugs in every system checked, 69 bugs in total, typically with little effort. One interesting result it found was that Berkeley DB could lose all committed transaction if it crashes.

How do you see your field evolving over the course of your career?

The software reliability and security problem is far from solved. I plan to continue to work in this direction. EXPLODE has shown that adapted model checking works well in storage systems.

In the near future, I am interested in trying out EXPLODE on other problems. I'm currently working on checking distributed systems at Microsoft Research. In the long term, I would also like to explore other opportunities where program analysis techniques can help build better systems.

Simha Sethumadhavan

How did you first become interested in computer science? Do you remember the moment or was it a gradually growing interest?

When I entered college I had no idea what I was going to major in. My parents suggested CS because they had heard that CS students get highly paid jobs. I was reluctant because I had never owned or used a computer or even played video games. Also I knew that several students choosing CS were doing so because they had been programming since they were young kids and they wanted to do more of the same. Anyway, upon the insistence of my parents, I decided to discuss this with the CS department chair at that time and he mentioned that he had never owned a computer and said that it was basically all about problem solving and I should not be worried about not having programmed. He encouraged me to choose CS.

So, I started taking CS classes and my interest gradually increased. I was really fortunate to have a totally awesome set of lab partners. These guys had very interesting views and opinions on everything and they had an awesome knack for getting anything to work; really anything from ICs and modems to Slackware Linux. I learned a lot just from hanging out with these guys. At the end of the second year, I participated in a programming contest and I came second, and it gave me some confidence. After that, I went back and thanked the CS department chair and he advised me to consider research. I did, liked it and then decided to go grad school.

How do you like to describe your field?

My research is in computer architecture. Basically, it's all about how to make applications go as fast as you want at the lowest possible cost. Several factors influence cost such as area, power consumption, programmability, and chip design complexity.

What motivated you to focus on your current research?

My current research is on cooperative hardware and software techniques for improving concurrent execution. Performance comes from two things: how fast you can do an operation, (frequency), and how many operations you can do in parallel. For a good 15 years starting from 1990 most computer performance came from ever increasing frequencies, but then it became difficult to cool the chips at very high frequencies. So processor designers, instead of trying to make things go fast, decided to create highly parallel hardware.

The predominant question for the architecture community, and arguably all of the CS community, is how do we exploit the highly parallel hardware? Some researchers argue that parallel software must be created to match the parallel hardware. Others argue that parallel software is too hard, has been tried before and had only limited success, and the solution must come from adding more hardware, which can assist in dynamic parallelization. My take on this problem is that we need to take another look at the abstractions that we have been working with and make adjustments so that the hardware and software can work together to solve this very difficult problem. This has to be done very carefully without increasing the cost and complexity of hardware and software.

How do you pick your research topics and what do you believe makes a good research question?

Computer architecture is influenced by applications and technology. Whenever there are changes in one of these two there is scope for innovation and research. I constantly look for these changes to decide what I want to work on next. I should clarify here that this does not mean that computer architecture research is purely reactive. Sometimes both technology and applications have to be manipulated in a particular way so as to be useful, and computer architects, with their understanding of applications and technology, are in the best position to bring these changes to the fore quickly; recent examples of such change include advances in nanomolecular memories and on chip nano coolers.

In my opinion, a good research question is one that you think you can solve (preferably, quicker and better than others), and when solved will make you and people who use your technology happy. To be more left-brained, it should have some value sometime in the future. A good research question should also inspire more good research questions.

Do you have a favorite project that you've worked on?

Yes, I'll mention two projects.

The first one was a multi-person, multi-million dollar, many year DARPA project in which we prototyped a full chip–the dual-core TRIPS processor. It was very exciting for me to see a project go from initial brainstorming and research all the way to working silicon. The chip building process was really gruesome, lots of low-level details and lots of abuse from the CAD tools but it was all worth when we had real SPEC programs running on the processor. I learned first hand where complexity arises in system design (and how to avoid it), how to build complex distributed systems and I also enjoyed the team camaraderie. It was definitely one of my best experiences.

The second project was a direct contrast to the first one; just one week's worth of work, no deadlines, just two people, and it was unfunded. It was a proposal for a cat warmer that worked on heat harvested from a microprocessor. We presented it at Wild and Crazy Ideas session at ASPLOS. The basic idea was to take the waste heat from a microprocessor and convert it back to energy to power the microprocessor. We argued that one day the use of futuristic nano-thermoelectric materials could improve the energy efficiency to make this worth considering. I had a lot of fun "brainstorming" with my advisor on the write up and presentation.

How do you see your field evolving over the course of your career?

Silicon based semiconductors are already on their last legs. I'm pretty sure that during my career we will see the decline of CMOS based computers. Hopefully, we will find replacement substrates and architectures to use these substrates effectively.

For individual profiles of Simha and Junfeng, please see page 4.

Columbia Computer Science Welcomes New Faculty Members **Simha Sethumadhavan and Junfeng Yang**

Assistant Professor
of Computer Science
Simha Sethumadhavan

Assistant Professor
of Computer Science
Junfeng Yang

Professor Simha Sethumadhavan obtained his PhD from The University of Texas at Austin in December 2007 and immediately thereafter joined Columbia University. During his tenure as a graduate student at UT, he was involved in two microprocessor designs. The first one was a two-core microprocessor–the TRIPS prototype processor–that could simultaneously execute up to 2K instructions. The second–the TFlex microprocessor–was a 64-core microprocessor that could execute up to 4K instructions from a single thread or up to 128 instructions from 64 threads or several combinations in between, so that the hardware can very nearly match the parallelism available in the application. Professor Sethumadhavan researched, designed, implemented and verified the memory system for the former and was involved in research and design of almost all aspects of the latter. Papers on both these topics have appeared as top picks from computer architecture conferences.

In the coming years, Professor Sethumadhavan and his students will be researching co-operative hardware/software techniques for improving concurrency on emerging many-core processors. Professor Sethumadhavan is also interested in system security with emphasis on hardware and on application-specific computers that will scale to Exaflops and software for such computers. Finally, Professor Sethumadhavan also wants to be a good teacher who is loved by all students, and is working at achieving this goal!

Professor Junfeng Yang's research interests span the area of operating systems, security, software engineering and program analysis. His past research has focused on practical software error detection. With his collaborators, he has built three error detection tools including MECA, which uses static analysis (i.e. without having to run the checked code) to detect security holes in Linux kernel; FiSC, which adapts model checking, a heavyweight formal verification technique, to detect serious file system errors; and EXPLODE, which improves upon FiSC to check general storage systems and drastically reduce the manual overhead needed for checking. He has also worked on using symbolic execution to automatically generate disk images of death that, when mounted, will cause a file system to crash.

Professor Yang's thesis research designs, implements, and evaluates EXPLODE, a general, lightweight approach to finding storage system errors using model checking. This formal verification technique finds corner-case errors by comprehensively exploring the state spaces defined by a system. Storage systems have two dynamics that make them attractive for such an approach. First, their errors are some of the more serious ones, since they can destroy persistent data and lead to unrecoverable corruption. Second, traditional testing requires an impractical exponential number of test cases to determine if a system will recover if it crashes at any point during execution. Naive application of model checking to large storage systems is, however, prohibitive because it requires huge manual overhead. EXPLODE uses a novel in-situ checking architecture that drastically reduces the manual overhead needed for checking: in the best case, clients write a few hundred lines of C++ code to check a new storage system. EXPLODE was applied to a broad range of real-world storage systems and found data-loss errors in every system checked, over sixty errors in total, typically with little effort.

Professor Yang received his PhD in Computer Science from Stanford University in January 2008, and his MS in 2002. He received his BS from Tsinghua University, Beijing, China in 2000. While at Stanford he won the Best Paper Award of OSDI 2004 for his work on model checking file systems. Professor Yang is currently spending a year as a post-doctoral researcher at Microsoft Research Silicon Valley; he will arrive on campus this summer. At Columbia University, Professor Yang plans to continue his research on reliability and security, to help build better computer systems.

New Computer Science Courses

Professor
Jason Nieh

Together with Scott Devine (co-founder of VMware), Professor **Jason Nieh** offered COMS E6998-2, **Virtual Machines**, in the Spring 2008 term. The current wave of virtualization technology from VMware, Xen, and Microsoft have revived interest in virtualization and virtual machine monitors. This course covers the design and implementation of virtual machine monitors as well other recent trends in virtualization. The first half of the course covers traditional virtualization techniques such as trap-and-emulate, binary translation, shadow page tables and device emulation. New hardware features introduced by Intel and AMD to assist virtualization are also covered. The second half of the course surveys the classic papers and other recent developments in virtualization.

In addition to the course instructors, the course features guest lecturers who in many cases will cover seminal virtualization papers which they have authored. The course provides a unique opportunity to learn from industry leaders in the field.

Visiting Professor
Dragomir Radev

Visiting Professor **Dragomir Radev** (University of Michigan) offered COMS 6998-6, **Network Theory**, in the Spring 2008 semester. The course is about naturally occurring networks such as the Web, social networks, citation networks, protein interaction networks, lexical networks, movie actor networks, etc. The course covered the mathematical and computational models that explain the behavior of these networks. Specific topics include random graphs, small worlds, scale-free networks, random walks and harmonic functions, spectral methods, descriptive analysis of networks, information diffusion and learning on graphs, etc.

Professor
Michael Chiang

Professor **Michael Chiang** (Columbia University College of Physicians & Surgeons) offered COMS W4560, **Introduction to Computer Applications in Health Care and Biomedicine (Biomedical Informatics)**, in the Fall 2007 semester. The course gives a broad overview of the field of biomedical informatics, combining perspectives from medicine, computer science and social science. It deals with the use of computers and information in health care and the biomedical sciences, covering specific applications and general methods, current issues, capabilities and limitations of biomedical informatics. Biomedical Informatics studies the organization of medical information, the effective management of information using computer technology, and the impact of such technology on medical research, education, and patient care. The field explores techniques for assessing current information practices, determining the information needs of health care providers and patients, developing interventions using computer technology, and evaluating the impact of those interventions.

Professor
Kathleen McKeown

Professor **Kathleen McKeown** offered COMS 6998-3, **Natural Language Processing for the Web**, in the Spring 2008 semester. Given the large amount of unstructured information on the web, whether text or spoken, natural language processing has the potential to have a large impact on accessing and harvesting information available on the web. The course focuses on applications using natural language processing that either have already been developed or are currently topics of research. Some of these applications aim to make it easier for end users to navigate the web (e.g., summarization and question answering) while others aim to make it easier to more accurately process information on the web (e.g., paraphrasing and entailment). The class covers the following topics: Text summarization; Question answering systems; Identifying meaning: paraphrasing, entailment; Extracting information from the web; Using language on the web as a machine learning corpus; and Multi Lingual Processing. The class has 4 invited speakers: John Prager (IBM) speaking on IBM's question answering system, Bill MacCartney (Stanford) speaking on Stanford's approach to textual entailment, Regina Barzilay (MIT) speaking on sentiment analysis for Web applications, and Ronald Kaplan (PowerSet) speaking about the use of NLP in web search. The class is seminar-style and focuses on reading research papers related the class topics. Students design and carry out a semester long project.

The **Computer Systems** Laboratory


An exterior view of the new Computer Systems Laboratory

The southeast corner of the fourth floor of the Computer Science Building has recently seen a major internal renovation with the creation of the new Computer Systems Laboratory (CSL) that brings together the research groups of Professors Luca Carloni, Stephen Edwards, Steven Nowick, and Simha Sethumadhavan.


Professor **Luca Carloni**


Professor **Stephen Edwards**


Professor **Steven Nowick**


Professor **Simha Sethumadhavan**

The renovated space is organized in two main areas: a large shared office space, with desks for the (currently thirteen) graduate students affiliated with CSL, and a separate laboratory for equipment including embedded computing devices, FPGA boards, microcontrollers, and set-top boxes. The new arrangement recovered about 700sf of previously unused space. The renovation was made possible in part by industrial support, including a four-year grant from Cablevision Inc., as well as support from the Fu Foundation School of Engineering and Applied Science and the Department of Computer Science.

CSL research activities encompass a broad range of topics in Computer Engineering, including digital integrated circuit design, computer architecture, computer-aided design (CAD), and embedded systems design with equal emphasis on theoretical foundations and practical applications. CSL researchers maintain active industrial collaborations with various leading semiconductor and information technology companies including Altera, Esterel Technologies, IBM, Intel, Philips, Boeing, and STMicroelectronics as well as government agencies such as NASA Goddard Space Flight Center.

System-Level Design

Prof. Luca Carloni

System-level design encompasses both the design of large integrated circuits produced by the semiconductor industry, such as multi-core processors and systems-on-a-chip, and the design of complex embedded systems which are a critical part in the products of many other industries including aerospace, automotive, health, building automation and public infrastructure. These systems are hard to design because their operations depend on the interaction of many devices that process data concurrently while communicating with non-negligible delays. Hundreds of millions of transistors switch concurrently on a modern chip at a speed that is faster than the time necessary for a signal to cross the chip. Many heterogeneous embedded devices such as sensors, actuators, and controllers work together exchanging messages over long distances to control the distributed subsystems in a modern car or airplane. Mapping concurrent behaviors on a distributed platform, assembling heterogeneous components, interfacing hardware and software, achieving a design implementation that is optimal under multiple criteria (e.g., performance and power), and guaranteeing robustness with respect to faults are some of the challenges of system-level design.

The System-Level Design Group led by Prof. Carloni develops modeling techniques, communication protocols, interface modules, and optimization algorithms that together yield new correct-by-construction design methodologies for complex systems. This research approach has its roots in Prof. Carloni's work on latency-insensitive design, a methodology that addresses the increasing impact of interconnect latency in nanometer integrated circuits and facilitates the reuse of pre-designed cores for building complex systems-on-chip. Ongoing research projects include a novel communication-based design flow for systems-on-chip, a public-domain software infrastructure for the automatic synthesis of interconnection networks (in collaboration with UC Berkeley), the design and fabrication of low-power network-on-chip prototypes (in collaboration with Prof. Shepard of the Electrical Engineering Department), the design of photonic networks for multi-core processing systems (in collaboration with Prof. Bergman of EE), and the development of programming frameworks for heterogeneous clusters of embedded devices and multi-core processors.

Embedded Systems

Prof. Stephen Edwards

Professor Edwards and his group explore automating the creation of software for embedded systems: application-specific computers hiding in a growing number of industrial and consumer systems. They have developed numerous compilation techniques for the Esterel synchronous language for real-time control and are also developing domain-specific languages.

Over the last few years, Edwards and his group developed the SHIM language (Software-Hardware Integration Medium). Designed for hardware/software codesign, SHIM is a concurrent imperative language whose semantics allow it to be implemented in both hardware and software. The current SHIM compiler can produce code that runs on a single processor, on multicore machines using shared memory and the Pthreads library, and work is ongoing to generate code for IBM's multicore Cell processor (found, e.g., in Sony's PS3).

Scheduling-independent concurrency is the central tenet of SHIM. Unlike most software concurrency models, SHIM's model does not allow data races: nondeterministic choices made by, say, the operating system schedule cannot affect the functional behavior of a SHIM program even though it may affect the execution speed. SHIM achieves this by prohibiting shared resources such as global variables, allowing tasks to only communicate through explicit messages. This makes SHIM programs deterministic, much easier to debug, and more easily ported to different hardware.

Edwards's group recently began work on the PRET (Precision Timing) project, which aims to rethink computer architecture for embedded systems. Today's high-performance processors all focus on improving average-case performance at the expense of predictability. For batch jobs, this is the right metric, but embedded systems usually have real-time constraints that demand predictable timing as much as predictable function. The end goal is a processor (and software development environment) able to guarantee timing. The architecture will have a hierarchy of scratchpad memories managed by software, thread-interleaved pipelines with no hazards, and time-triggered communication networks.

Asynchronous and Mixed-Timing Circuits and Systems

Prof. Steven Nowick

Traditionally, most digital systems are **synchronous**: operating under a global clock. All components operate in lock-step, and the fixed-rate clock must be evenly distributed across the entire chip. There are now fundamental challenges in constructing systems with centralized clock control, as chips approach one billion transistors, clocks achieve multi-Gigahertz rates, and device variability becomes severe. In contrast, an **asynchronous** digital system has no clock, and is organized as a distributed system, where components synchronize on local channels using handshaking protocols. Asynchronous design promises several benefits for complex systems: *low power; robustness to timing variability;* and *modularity and composability*. However, these systems are challenging to design, requiring new tools, circuit styles and methodologies.

Computer-Aided Design (CAD) Tools. With current and former students (Agyekum, Dearing, Fuhrer, Jeong, McGee, Theobald), Prof. Nowick has released the **CaSCADE tool suite** *(http://www.cs.columbia.edu/~nowick/asynctools)*. It includes MINIMALIST (for controllers), ATN_OPT (for highly-robust threshold networks) and DES Analyzer (for performance analysis). MINIMALIST was recently used at NASA Goddard Space Flight Center to design prototype chips for laser space measurement. These exhibited significant power and performance advantages over synchronous designs, and are being evaluated for use in future space missions.

High-Speed Asynchronous Pipelines. Pipelining is critical in designing high-performance systems. With a former student (Singh), Prof. Nowick developed three pipeline styles (all patented) that obtain comparable performance to synchronous, but with much greater adaptability to variable-speed environments. IBM used one style for an experimental FIR filter chip for disk drive reads, with lower latency than IBM's best commercial synchronous design.

Mixed-Timing Interfaces. While asynchronous circuits have large potential, in practice they will often be used in mixed-timing systems that incorporate clocked components. With a former student (Chelcea), he developed a family of robust, low-latency interface circuits to mediate between the different timing domains.

Flexible, High-Throughput Interconnect Mesh for Parallel Processors. Prof. Nowick is also developing a high-speed flexible asynchronous interconnect network for massively-parallel supercomputers-on-chip. Synchronous meshes for shared memory processors, which connect processors and memories, typically suffer from power, bandwidth and scalability problems. In contrast, the asynchronous mesh promises to flexibly interface with synchronous processors operating at varied clock rates, and avoid the power overheads of high-speed clock distribution. This work is in collaboration with the University of Maryland, and with seedling funds from the executive VP of research at Columbia.

Computer Architecture

Prof. Simha Sethumadhavan

Please see interview with Prof. Sethumadhavan on page 3 of the present issue.


A peek at the interior of the new Computer Systems Laboratory

DejaView:
A Personal Virtual Computer Recorder

Ricardo Baratto

Oren Laadan

Dan Phung

Shaya Potter

Professor
Jason Nieh

As users spend more time interacting with the world and their peers through their computers, it is becoming important to archive and later search the knowledge, ideas and information that they have seen through their computers.

As users spend more time interacting with the world and their peers through their computers, it is becoming important to archive and later search the knowledge, ideas and information that they have seen through their computers. However, finding the information one has seen among the ever-increasing and chaotic sea of data available from a computer remains a challenge. Exponential improvements in processing, networking, and storage technologies are not making this problem easier. Computers are getting faster at generating, distributing, and storing vast amounts of data, yet humans are not getting any faster at processing it.

Some tools address aspects of this problem. Web search engines focus on static information available on the web, but do not help with a user's personal repository of data, dynamically generated and changing content created at the moment a user has viewed a web page, or hidden databases a user may have seen but are not available through web search engines. Similarly, desktop file search tools return current files that may be of interest, but do not return results from files that are no longer available, or from information seen by the user but never actually saved to files. More importantly, none of these tools provides a way to go beyond individual static documents to capture aggregate information across all resources visible on a user's desktop.

Vannevar Bush's Memex vision was to build a device that could store all of a user's documents and general information so that it could be quickly referenced. Building on that vision, a team of researchers in the Network Computing Laboratory (NCL), Oren Laadan, Ricardo Baratto, Dan Phung,Shaya Potter, and Professor Jason Nieh, have created DejaView. DejaView is a personal virtual computer recorder that provides a complete WYSIWYS (What You Search Is What You've Seen) record of a desktop computing experience. The system enables users to playback, browse, search, and revive records, making it easier to retrieve information they have seen before.

Leveraging continued exponential improvements in storage capacity, DejaView records what a user has seen as it was originally displayed with the same personal context and layout. All viewed information is recorded, be it an email, web page, document, program debugger output, or instant messaging session. DejaView enables a user to playback and browse records for information using functions similar to personal video recorders (PVR) such as pause, rewind, fast forward, and play. DejaView enables a user to search records for specific information to generate a set of matching screenshots, which act as portals for the user to gain full access to recorded information. DejaView enables a user to select a given point in time in the record from which to revive a live computing session that corresponds to the desktop state at that time. The user can time travel back and forth through what she has seen, and manipulate the information in the record using the original applications and computing environment.

DejaView transparently provides these features by introducing lightweight virtualization mechanisms and utilizing available accessibility interfaces. DejaView virtualizes the display to capture and log low-level display commands, enabling them to be replayed at full fidelity at a later time. It utilizes accessibility interfaces to simultaneously capture displayed text and contextual information to automatically index the display record so it can be searched. It combines display and operating system (OS) virtualization to decouple window system and application state from the underlying system, allowing them to be continuously checkpointed and later revived, while only saving user desktop state, not the entire OS instance. Checkpointing at this finer granularity, shifting expensive I/O operations out of the critical path, and using various optimizations such as fast incremental and copy-on-write techniques are crucial to minimize any impact on interactive desktop application performance. DejaView combines logging and unioning file system mechanisms to capture the file system state at each checkpoint. This ensures that applications revived from a checkpoint are given a consistent file system view corresponding to the time at which the checkpoint was taken.

DejaView's ability to browse and search display content and revive live execution provides a unique blend of functionality and performance. By browsing and searching the display record, the user is able to access content as it was originally seen, and quickly find information at much faster rates than if the information had to be generated by replaying execution. By reviving the execution environment, the user can go beyond a static display of content to fully manipulating and processing information using the same application tools available when the information was first displayed.

The NCL research group has implemented DejaView as a set of loadable modules for Linux and the X Window System. It provides transparent operation without modifying, recompiling, or relinking applications, window systems, or OS kernels. To demonstrate its effectiveness, the group has evaluated its performance on a wide-range of real-world desktop applications. The results show that DejaView can provide continuous low-overhead recording without any user noticeable performance degradation of the system. Downtime due to checkpoints when running desktop application benchmarks is less than 10 ms, a time delay much shorter than what humans can readily detect. Storage requirements of DejaView records at highest quality are comparable to PVRs in recording HDTV resolution media programming. As terabyte storage capacities become commonplace, DejaView enables high quality WYSIWYS recording to be used for everyday use. The results also show that DejaView can provide much faster than real-time playback of records and supports browsing and searching of records fast enough for interactive use.

DejaView introduces a new computer recorder model to the desktop that opens up new directions for future research. Some areas of ongoing work include (1) conducting user studies to explore usage patterns to better understand how DejaView will be exploited by users over extended periods of time and how the user interface can be enhanced to better fit daily usage needs, (2) quantifying and improving the relevance and presentation of search results by exploring the use of desktop contextual information such as time, persistence, or the relationships among desktop objects, and (3) addressing the privacy and security ramifications of this emerging computing model.

For more information about the Network Computing Laboratory and the DejaView project, see http://ncl.cs.columbia.edu.

Columbia Robotics:

Building a Better Imaging Device for Minimally Invasive Surgery

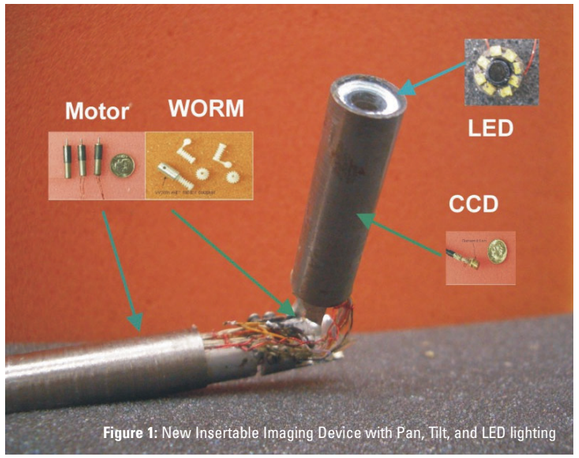

Figure 1: New Insertable Imaging Device with Pan, Tilt, and LED lighting

Anyone who has faced surgery in the last few years is familiar with the idea of minimally invasive surgery. There are many benefits to this kind of surgery including small incisions, less pain and trauma to the patient, faster recovery time, and lower health care cost. The most common type of this surgery involves using an endoscope, which is essentially a telescope inserted into a small opening in the body. The surgeon first cuts several small incisions and inserts the endoscope and instruments into the abdomen. Most endoscopic surgery is still dominated by the paradigm of pushing long sticks into small openings. This approach has a number of limitations for minimal access surgery, such as narrow angle imaging, limited workspace, counter-intuitive motions (e.g. moving the endoscope to the left inside the body requires moving it to the right outside the body as the device pivots about the insertion point), and making additional incisions for the endoscopic instruments used during the procedure.

A more compelling vision of minimally invasive surgery is being developed at Columbia by robotics professor Peter Allen working with his colleague Dr. Dennis Fowler from the Department of Surgery. The new idea being pursued is called "single port surgery", in which a single small incision is made in the body and multiple tools and sensors are placed inside the body which can then be used to perform the surgical procedure.

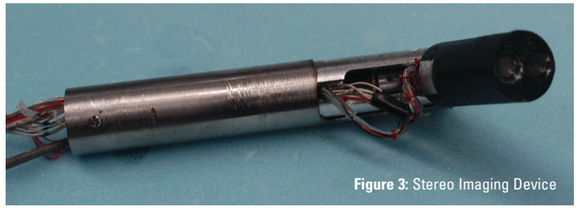

Figure 3: Stereo Imaging Device

As a first step, a new insertable, in-vivo camera was developed with assistance from post-docs Dr. Andrew Miller (Columbia CS Ph.D. '01) and Dr. Tie Hu. This device is notable for a number of innovations that have resulted in a U.S. patent on the technology. First, the device *(see Figure 1)* has multiple axes, providing pan, tilt and zoom functions while inside the body. Second, it has an integrated LED light source that provides direct illumination inside the body. Third, it is controlled by a joystick to allow an untrained operator to move the camera in a natural way (left is left, right is right, etc.). The device uses very small 5mm diameter servo motors to move the camera. Software has been developed to control the cameras, and using feedback from the images, the camera can automatically track organs and instruments during a procedure, freeing the surgeon to concentrate on the task at hand instead of positioning the cameras. The device has been tested in animals, where it has been successfully used for surgical procedures including gall bladder removal and appendectomies *(see Figure 2)*. The pan, tilt, and zoom axes allow a larger field of view for the surgeon, and this can be achieved with only one insertion port as opposed to moving a traditional laparoscope from port to port. Recently, a stereo version of the device was tested that provides true 3D imaging using a head mounted display *(see Figure 3)*.

Figure 2: Image from the new device during a suturing procedure.

The device has great market potential. Current robotic surgery systems are extremely large and very expensive, often costing over a million dollars. The new device is designed to be potentially disposable, which eliminates the need for sterilization between uses. The success of this device has led to new research into creating both an imaging and tooling platform inside the body, funded by a new grant from the National Institutes of Health. Working with Prof. Nabil Simaan from Mechanical Engineering, a new insertable robot surgical platform is being developed that has small snake-like robot end effectors attached and can be controlled by the new cameras which will be part of the platform. A device like this is needed for NOTES (Natural Orifice Transluminal Endoscopic Surgery) surgical procedures, in which natural body openings are used to insert the devices rather than through external incisions. NOTES surgery promises even greater reductions in pain, scarring and healing time for patients.

Knarig Arabshian

Advisor: Professor Henning Schulzrinne

Ontology-based Context-aware Service Discovery in a Globally Distributed Network

In order to accomplish context-aware service discovery, static or dynamic service information must be made available globally and represented in such a way that registering and querying for them can be more complicated than what is available to us now. Imagine a service discovery system where a user can enter the following searches: "Find me a nearby Italian restaurant that has available seating for 5 people" or "Find me a Chinese restaurant or something similar to it that has a nearby movie theatre playing an action movie". This requires efficiently distributing dynamic service information which require frequent updates as well as enhancing querying and registration beyond text and key word search so that specific services as well as logically similar services can be searched for.

This dissertation introduces GloServ, an ontology-based global service discovery system. Two of its main contributions are in its scalable network architecture and intelligent querying and registration of services. GloServ aggregates different types of services in a globally distributed network. A partial list of such services include: real-time event-based services, location-based services, communication, e-commerce or web services. GloServ is a service discovery architecture that uses a description logic ontology, such as the Web Ontology Language Description Language (OWL DL), for classifying service information and mapping them onto a physical hierarchical peer-to-peer network. GloServ operates on a wide as well as local area network and supports the discovery of all types of services that are described in an ontology. This classification ontology defines service classes and their relationships with other services and properties. Thus, Gloserv aggregates and classifies service information in a distributed network.

Due to the use of description logic ontologies, GloServ can perform semantic matching of queries to return results that are logically related to the user's request, rather than searching only for exact information using attribute-value matching as current service discovery protocols do. Initially, the GloServ query matchmaking engine represents the query as a temporary class within the service ontology. This class is essentially a first order logic statement which has a number of its properties restricted. The matching engine then uses an ontology reasoner to classify the temporary query class within the ontology. Service information within equivalent classes and subclasses are considered matches. For similarity matches, service information within the non-disjoint sibling classes are looked into.

Additionally, GloServ further refines the results obtained by combining both ontology-based querying with key word matching. Since the service ontology may not capture all parts of a particular service, it allows each ontology to have a keywords property which service providers can populate with keyword terms that describe the specifics of their service. Furthermore, GloServ allows querying of more than one service in a single query by supporting subquerying between its servers, which allows services that share common properties to be composed into a single query. This enhancement causes fairly complicated queries to be issued in a single search such as the one given in the above example of searching for a restaurant and a nearby theatre.

We have built a prototype implementation of GloServ and the results have shown that the underlying distributed architecture improves the query latency and load distribution of query and registration messages. Also, because GloServ uses an ontology to map a network, different service classes are added to the network quite easily as only an ontology configuration file is needed for it to be generated and mapped to a network of servers. Furthermore, we have built a web-based front-end for a Restaurant and Theatre service search which demonstrates the different querying methods GloServ allows.

Marcio Buss

Advisor: Professors Stephen Edwards and Alfred Aho

Summary-Based Pointer Analysis Framework for Modular Bug-Finding

Modern society is irreversibly dependent on computers and, consequently, on software. However, as the complexity of programs increase, so does the number of defects within them. To alleviate the problem, automated techniques are constantly used to improve software quality. Static analysis is one such approach in which violations of correctness properties are searched and reported. Static analysis has many advantages, but it is necessarily conservative because it symbolically executes the program instead of using real inputs, and it considers all possible executions simultaneously. Being conservative often means issuing false alarms, or missing real program errors.

Pointer variables are a challenging aspect of many languages that can force static analysis tools to be overly conservative. It is often unclear what variables are affected by pointer-manipulating expressions, and aliasing between variables is one of the banes of program analysis. To alleviate that, a common solution is to allow the programmer to provide annotations such as declaring a variable as unaliased in a given scope, or providing special constructs such as the "never-null" pointer of Cyclone. However, programmers rarely keep these annotations up-to-date.

The solution is to provide some form of pointer analysis, which derives useful information about pointer variables in the program. An appropriate pointer analysis equips the static tool so that it is capable of reporting more errors without risking too many false alarms. This dissertation proposes a methodology for pointer analysis that is specially tailored for "modular bug finding." It presents a new analysis space for pointer analysis, defined by finer-grain "dimensions of precision," which allows us to explore and evaluate a variety of different algorithms to achieve better trade-offs between analysis precision and efficiency. This framework is developed around a new abstraction for computing points-to sets, the Assign-Fetch Graph, that has many interesting features. Empirical evaluation shows promising results, as some unknown errors in well-known applications (such as the Linux Kernel) were discovered.

Marc Eaddy

Advisor: Professor Alfred Aho

An Empirical Assessment of the Crosscutting Concern Problem

Modularity is essential for creating evolvable software. Traditional programming languages limit the level of modularity that can be achieved because they provide only one way to structure the program, e.g., as a hierarchy of types. Regardless of how a program is structured, some features, requirements, or other types of concerns of the program cannot be modularized; that is, they cut across the program's structure. These so-called crosscutting concerns result in programs that are difficult to understand and reason about, to divide into manageable pieces, to reuse, and to evolve.

In this dissertation, we assess the impact of crosscutting

concerns on software quality. Little is known about the impact because determining the concerns of a program and how they are implemented is difficult, existing metrics for quantifying crosscutting are inadequate, and empirical evidence is lacking.

To locate the source code that implements a concern, i.e., concern location, we present a new technique called prune dependency analysis, which can be combined with existing concern location techniques to dramatically improve accuracy. We developed Cerberus, a potent hybrid technique for locating concerns that combines information retrieval, execution tracing, and prune dependency analysis. Our experiments show that prune dependency analysis can boost the recall of information-retrieval-based concern location by 155% and execution tracing by 104%. In addition, we show that our combined technique outperformed other techniques when run individually or in pairs.

After using our techniques to locate all the code that implements a concern, our novel concern metrics quantify the amount of crosscutting that exists. We assessed the crosscutting concern problem by performing rigorous empirical studies of five medium sized programs. We found a moderate to strong statistically significant correlation between the amount of crosscutting and the number of defects. That is, the more a concern crosscuts the program, the more likely it is to have defects. We also found that the crosscutting concern problem was prevalent: 86% of the concerns we analyzed were crosscutting, concerns are implemented by 6 classes on average, and classes implement 10 concerns on average. We propose unique solutions to the crosscutting concern problem, including a novel language, Wicca#, which provides direct support for modularizing crosscutting concerns.

Alexander Haubold

Advisor: Professor John Kender

Indexing and Browsing Unstructured Videos using Visual, Audio, Textual, and Facial Cues

In the domain of highly unstructured and unedited videos, we introduce novel approaches for automatic indexing and browsing of audio, visual, and textual contents. Our video browser is one of the first of its kind for large libraries of unstructured video data, having been tested on more than 500 university lecture and student presentation videos. Unique multi-modal indices sensitive to the raw video data and to user interaction have been measured and refined in extensive user studies with more than 1,000 students. We report on significant improvements in typical video retrieval tasks, such as search by text and visual contents. We also report on statistically significant improvements in student performance when the video browser was used for exam preparation.

We introduce the prototype platform VAST MM (Video Audio Structure Text MultiMedia), which encompasses a video indexer and a platform independent video browser, featuring visual cues, a custom-designed MPEG1 video stream player, text search, and comparison tools based on key phrases. Relevant indices shared among video genres include visual-based segmentation of likely scenes, the boundary of which can be adjusted interactively in the browser to increase or decrease the quantity of displayed visual information; automatic speech recognition (ASR) transcripts; and filters for ASR text derived from external material, such as presentation slides and textbook indices to extract key phrases. Furthermore, indices applicable to student presentation videos include audio-based speaker segmentation; a graphical face index of speakers in the video, displaying frontal and profile face shots linked to the video clips in which they occur; and audio-based speaker clustering of recurring speakers used to illustrate a visual flow of speakers throughout a video. Indices unique to lecture videos include visual-based clustering of scenes into teaching units, characterized by similar blackboard content; and visual-based classification of media used during the lecture, e.g. snapshots of blackboard, computer, illustration, students, lecturer, etc.

We note the following contributions: The VAST MM browser demonstrates how multi-modal information can be combined to provide a rich set of indices which improve the video browsing experience for users even when the video is unedited. The timeline ties together these modalities to maintain context, while individual visual indices can be manipulated interactively by the user to adjust the amount of displayed information. Instead of producing user interfaces that provide a high-level summary of a video, content retains its temporal context. In lieu of a tag cloud we show ranked words and phrases tied to their temporal occurrence while clustering their recurrences visually. The streaming video player is augmented with "keyframes," redefined here as visually distinct images representative of the video's content. The player's location slider serves as a very-fast-forward and reverse feature, with which visual content can be viewed in short amount of time, e.g. 10 seconds. We introduce face indices to replace otherwise expensive and inaccurate person recognition. Finally, with a database of more than 500 hours of video, VAST MM demonstrates a tested application with large amounts of data.

Observations have been collected from more than 1,000 students in the course of their studies to measure usefulness of video indices, ease of use of browsing tools, and impact of the availability of such resources. Seven differential experiments were conducted to better quantify selected features: the same search and retrieval tasks were completed by groups of students with and without the availability of streaming video, text cues, visual cues, etc. We find, among other results, that the presence of streaming video negatively impacts search performance by requiring significantly more time (up to 74%) for task completion without providing higher task completion rates (which remain at 90%). In general, in the course of 11 experiments in a three-year period over which index cues and user interfaces were refined, we were able to significantly improve search performance for difficult tasks by decreasing required time from 436 seconds to 128 seconds while increasing completion rates from 57% to 97%.

In the domain of presentation videos made available through VAST MM, students not only gain insight into professional presentation skills from prior semesters' examples, but also assess their own progress on skill development. In the domain of instructional videos, availability of recorded lectures through VASTMM for exam preparation strongly suggests that it should be a mainstream resource. Statistical results from studies conducted on two core computer science courses show that students who use video-recorded lectures experience an average improvement of one third of a standard deviation in midterm-to-final exam scores.

Abhinav Kamra

Advisor: Vishal Misra

Distributive Adaptive Algorithms or Data Resilience in Sensor Networks

A central idea in sensor networks is collecting data from the various sensor nodes at one or more access points (located typically at the edge of the network). Data aggregation schemes take advantage of the implicit routing to aggregate data while it is being sent from the nodes to the access point(s). In such data aggregation, sensor networks are treated as distributed database systems which can be queried by the access point(s).

We focus on sensor networks deployed in disaster scenarios such as floods, fires, earthquakes or any region where the sensor nodes are failure-prone. Such scenarios present an interesting design challenge since the sensor nodes used to collect and communicate data may themselves fail suddenly and unpredictably, resulting in the loss of valuable data. Furthermore, the nodes cannot rely on an underlying routing mechanism being always present since the routing setup is frequently disrupted as nodes fail. In such scenarios, more sophisticated data collection/aggregation schemes are needed to preserve as much data as possible from the network and forward it to the access point(s). Such methods for Data Persistence need to be able to adapt to changing network scenarios, whether it is taking advantage of the underlying routing setup whenever available, replicating data across the nodes when the nodes are failing, or intelligently encoding data to maximize its recovery at the access point(s).

We propose algorithms for data resilience in sensor networks which are failure-prone or where the underlying routing mechanism is either non-existent or only partially available. Such networks include mobile sensor networks where the natural connection topology of the network changes continuously due to the movement of the nodes. Our algorithms are fully distributed in that each sensor node needs to make decisions using only local information. Furthermore, our algorithms are adaptive in the sense that as network conditions change the algorithms seamlessly adapt to maximize data resilience according to the changed environment.

We propose Growth Codes: a distributed channel coding scheme to maximize data recovery at the access point(s). Growth Codes provide maximal recovery of sensor data at the access point(s) irrespective of how much data is actually received from the sensor nodes. These codes do not require an underlying routing setup but can take advantage of any routing information whenever available to speed up the data dissemination process. We also propose CountTorrent: a technique for efficient data aggregation over sensor networks where the network topology is continuously evolving. The data aggregation scheme remains efficient and accurate even as nodes move, join or leave the network. The accuracy of the data aggregation degrades gracefully as network conditions deteriorate.

The contributions of this thesis include a novel distributed channel coding mechanism as well as an efficient and adaptive data aggregation mechanism for sensor networks. Both these algorithms facilitate efficient data resilience in sensor networks deployed in adverse network conditions. On the other hand, when network conditions do improve, both the mechanisms can adapt to take advantage of the changed circumstances for increased efficiency. This automatic adaptability distinguishes our algorithms from most of the previous approaches for data collection in sensor networks.

Maryam Kamvar

Advisor: Professor Steven Feiner

Using Context to Improve Query Formulation and Entry from Mobile Phones

The goal of this thesis is to improve the query formulation and entry step of the web search process from mobile phones. We achieve this by providing query recommendations (queries displayed to the user before she has begun query input) and query predictions (queries displayed to the user as she is entering the query) as part of the mobile search interface.

We derived the motivation for our research from a comprehensive overview of the state of mobile web search. After analyzing millions of requests made to the Google mobile search interface, we were able to identify areas for improvement. We discovered that the first step of the search process—formulating and entering the query—is time-consuming and cumbersome for mobile users.

To address this, we present two approaches for improving mobile query formulation and entry: relevant query recommendations and accurate query predictions. We build recommendation and prediction models that consider a user's context when she is interacting with the search interface. We restrict context to the set of circumstances captured in the Google search logs at the time of query. We find that the two contextual signals that make the biggest impact in improving the accuracy of predictions, thus reducing the number of key presses needed to enter a query, are knowledge of the application being used (in this case a search engine), and of the location of the user. Although knowledge of the day of week did not significantly impact the prediction model, it does improve the recommendation model. Knowledge of the time of day did not significantly improve either model.

The query recommendation and prediction models presented in this thesis are general and can be used in conjunction with any search interface. However, we focus our research on search interfaces for mobile phones, where we believe the greatest need for improvement exists. We design, implement, and evaluate interfaces for the prediction and recommendation models in order to improve the mobile search experience. Our findings can be easily incorporated in existing mobile search systems.

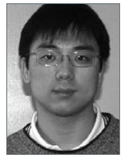

Fei Li

Advisors: Professors Jay Sethuraman and Clifford Stein

Competitive Queuing Policies for Packet Scheduling

In the Internet, all information is aggregated into packets. The computers connected to the Internet communicate with each other by means of exchanging packets. All packets travel through communication links and network switches. If a burst of packets arrives at the same time, a network switch cannot transmit all of them on the fly. Inside a network switch, there are some output buffer(s). Arriving packets are queued in the output buffer(s), waiting to be delivered.

Most current Internet switches adopt the First-In-First-Out (FIFO) buffering policy. Using the FIFO buffering policy, network switches send packets in the same order as they arrive. FIFO buffering policy cannot provide assured data transmission for time-critical applications or mission-critical applications due to unpredictable packet loss, end-to-end delay, out-of-order delivery, and jitter.

In the past ten years, there has been rapid growth of network traffic and time-critical applications. The diversity of applications results in unpredictable packet flows and heterogeneous network traffic, and motivates us to study buffer management at the switch levels for providing better Quality of Service (QoS). The difficulty of achieving better QoS in the existing Internet infrastructure without sacrificing high resource utilization remains open.

In this thesis, we study a model in the Differentiated Services (DiffServ) QoS infrastructure, which is called the "bounded-delay model". We characterize packets by their deadlines by which they should be sent, and their payoffs when they are delivered on time. Our goal is to maximize the total value of packets sent by their deadlines. We design

deterministic online packet scheduling algorithms for QoS queuing policies. We also develop novel and effective analysis techniques. Our algorithms perform better over existing solutions in terms of competitive ratio, which provides a worst-case performance guarantee for all traffic patterns. The ideas and approaches developed in this thesis explore the insights in similar online models with deadline constraints, and can be generalized and applied to other online scheduling problems.

Wei-Jen Li

Advisor: Professor Salvatore Stolfo

SPARSE: Malcode-Bearing Document Detection

Embedding malcode within documents provides a convenient means for attackers to penetrate systems owing to the complexity of the document formats that provide ample opportunity to embed code in a myriad of ways and the multitude of document-exchange vectors which may be otherwise unreachable by network-level attacks due to the multitude of document-exchange vectors. Hence, no single approach will likely perform perfectly (reaching 100% detection accuracy and 0% false positive rate) with a minimum of computational expense. In this thesis study, I suggest a means of combining multiple detectors with various detection tactics that will likely provide improved security to detect malcode embedded in documents. To this end, I present an integrated detection system that combines both static and dynamic detection methodologies with various detection methods such as data randomization and malcode locating mechanism. In addition, I also introduce the file content difference and entropy analysis that may provide a useful forensic tool to inspect documents. I have evaluated the detection performance of a number of detectors and methods, as well as the integrated system, by utilizing thousands of both legitimate and malicious documents. Ultimately, I do not argue that the presented detection system is perfect; nevertheless, this thesis demonstrates that a hybrid detection system combining various detection strategies can enhance the level of protection, and therefore the attackers who craft documents with embedded malcode will expend far more effort to create undetectable malcode bearing documents.

Sameer Maskey

Advisor: Professor Julia Hirschberg

Automatic Broadcast News Speech Summarization

As the numbers of speech and video documents available on the web and on handheld devices soar to new levels, it becomes increasingly important to enable users to find relevant, significant and interesting parts of the documents automatically. In this dissertation, we present a system for summarizing Broadcast News (BN), ConciseSpeech, that identifies important segments of speech using lexical, acoustic/prosodic, and structural information, and combines them, optimizing significance, length and redundancy of the summary. There are many obstacles particular to speech such as word errors, disfluencies and the lack of segmentation that make speech summarization challenging. We present methods to address these problems. We show the use of Automatic Speech Recognition (ASR) confidence scores to compensate for word errors; present a phrase-level machine translation approach using weighted finite state transducers for detecting disfluency; and present the possibility of using intonational phrase segments for summarization. We also describe structural properties of BN used in determining which segments should be selected for a summary, including speaker roles, soundbites and commercials. We present Information Extraction (IE) techniques based on statistical methods such as conditional random fields and decision trees to automatically identify such structural properties. ConciseSpeech was built for handling single spoken documents, but we have extended it to handle user queries that can summarize multiple documents. For the query-focused version of ConciseSpeech we also built a knowledge resource (NENET) that can find related named entities to significantly improve the document retrieval task of query-focused summarization. We show how all these techniques improve speech summarization when compared to traditional text-based methods applied to speech transcripts.

Sangho Shin

Advisor: Professor Henning Schulzrinne

Towards the Quality of Service for VoIP traffic in IEEE 802.11 Wireless Networks

The usage of voice over IP (VoIP) traffic in IEEE 802.11 wireless networks is expected to increase in the near future due to widely deployed 802.11 wireless networks and VoIP services on fixed lines. However, the quality of service (QoS) of VoIP traffic in wireless networks is still unsatisfactory. In this thesis, I identify several sources for the QoS problems of VoIP traffic in IEEE 802.11 wireless networks and propose solutions for these problems.

The QoS problems discussed can be divided into three categories, namely, user mobility, VoIP capacity, and call admission control. User mobility causes network disruptions during handoffs. In order to reduce the handoff time between Access Points (APs), I propose a new handoff algorithm, Selective Scanning and Caching, which finds available APs by scanning a minimum number of channels and furthermore allows clients to perform handoffs without scanning, by caching AP information. I also describe a new architecture for the client and server side for seamless IP layer handoffs, which are caused when mobile clients change the subnet due to layer 2 handoffs.

I also present two methods to improve VoIP capacity for 802.11 networks, Adaptive Priority Control (APC) and Dynamic Point Coordination Function (DPCF). APC is a new packet scheduling algorithm at the AP and improves the capacity by balancing the uplink and downlink delay of VoIP traffic, and DPCF uses a polling based protocol and minimizes the bandwidth wasted from unnecessary polling, using a dynamic polling list. Additionally, I estimated the capacity for VoIP traffic in IEEE 802.11 wireless networks via theoretical analysis, simulations, and experiments in a wireless test-bed and show how to avoid mistakes in the measurements and comparisons.

Finally, to protect the QoS for existing VoIP calls while maximizing the channel utilization, I propose a novel admission control algorithm called QP-CAT (Queue size Prediction using Computation of Additional Transmission), which accurately predicts the impact of new voice calls by virtually transmitting virtual new VoIP traffic.

Stylianos Sidiroglou

Advisor: Professor Angelos Keromytis

Error Virtualization

Software errors and subsequent security vulnerabilities continue to be a thorn in the side of computer systems despite considerable efforts in both research and development strategies. They are of particular concern to high-availability systems where outages have been estimated to cost businesses billions of dollars each year. Many traditional proactive approaches have been employed to attempt to address this problem by trying to make code as dependable as possible, through a combination of safe languages, libraries, compilers, code analysis tools, and development methodologies. Unfortunately, experience has shown that it is very hard to achieve bug-free software.

Further exacerbating the problem is the fact that for existing proactive mechanisms the only available action upon detection of a fault is program termination, effectively a denial-of-service attack on the application. This situation is particularly troublesome for server applications that need to maintain high availability in the face of remote attacks, high-volume events that may trigger unrelated and possible non-exploitable bugs, or simple application-level denial of service attacks. Given these problems, we posit that in the absence of perfect software, error toleration and recovery techniques become a necessary complement to proactive approaches.

Towards this goal, this dissertation describes the development of a general autonomic self-healing framework for handling a wide variety of software failures, ranging from remotely exploitable vulnerabilities to more mundane bugs that cause abnormal program termination (e.g., illegal memory dereference) or other recognizable bad behavior (e.g., computational denial of service). Briefly, the approach of this work to self-healing systems has been to use lightweight monitoring mechanisms to observe and analyze failures (including software vulnerabilities, but also bugs leading to system crashes), develop candidate fixes by modifying the software or its environment, and validate these fixes through a combination of static and dynamic analysis, and automated testing. The goal is to automate as much of the reaction ("self-healing") process as possible, toward building a completely autonomous system-healing mechanism.

The major contribution of this thesis is the introduction of the concept of error virtualization, a mechanism for program execution recovery. Briefly, error virtualization operates under the assumption that there exists a mapping between the set of errors that could occur during a program's execution (e.g., a caught buffer overflow attack, or an illegal memory reference exception) and the limited set of errors that are explicitly handled by the program's code. Thus, a failure that would cause the program to crash is translated into a return with an error code from the function in which the fault occurred (or from one of its ancestors in the stack). Conceptually, error virtualization is a mechanism that retrofits exception-handling capabilities to legacy software. The main premise of error virtualization is that inside every complex software system there exists a well-tested core in which the system has been observed to behave acceptably. If the code inside this well-tested core can be harnessed to handle errors that occur outside this space, then error virtualization can help an application handle the failure and call upon existing code to help with data-structure cleanup.

Experimental results are presented to support our hypothesis. The results demonstrate that our techniques can recover program execution in the case of failures in 80 to 90% of the examined cases, based on the technique used. Furthermore, the results illustrate that the performance overhead induced by the techniques to protect against a specific fault can be minimized to under 10%. This dissertation also describes two deployment mechanisms, Shadow Honeypots and Application Communities, that can reduce the cost of monitoring the application and, in turn, enable efficient deployment strategies for error virtualization systems.

Bo Sun

Advisor: Professor Ravi Ramamoorthi

Analytical, Frequency and Wavelet based Mathematical Models for Real-Time Rendering

Intelligent real-time rendering techniques are critical and highly desirable in interactive graphics applications such as computer games, flight simulations, and interactive design tools. Offline rendering algorithms such as raytracing, path tracing, photon mapping and Monte Carlo methods generate high quality realistic images and capture complicated effects such as caustics, soft shadows, scattering, and near-field lighting environment but take hours or even days to run. Various real-time rendering techniques have been proposed. However, either their captured effects are too simplistic, missing intricate scattering or near-field effects, or their assumptions are too constraining, limiting them to only very specific application scenarios. In this thesis, we take analytical, frequency and wavelet based approaches to expand the real-time rendering domain to incorporate a number of much more complicated natural phenomena. While open topics are plenty, we focus on a few effects that are of the first-order importance in improving the photo-realism of current rendering techniques. In particular, we focus on four challenging problems. Firstly, we consider real-time rendering of scenes in scattering media, capturing the effects of light scattering in fog, mist and haze. We present a physically based analytic model that captures these effects while maintaining real time performance. Our method is based on an explicit analytic integration of the single scattering light transport equations for an isotropic point light source in a homogeneous scattering medium. Then we present a novel near-field relighting framework. At the core of this framework is an affine double and triple product integral theory, an important generalization of triple product wavelet integrals that enables one of the product functions to be scaled and translated. We overturn the commonly held view that affine transformation is difficult in wavelets and show that while simple analytic formulae are not easily available, there is considerable sparsity that we can exploit computationally. We demonstrate practical implementation of an intuitive lighting design system coupled with near-field relighting capabilities. We also illustrate initial examples of wavelet importance sampling with near-field area lights, and image processing directly in the wavelet domain. Finally, we introduce two frequency based approaches to normal map filtering and dynamic soft shadowing. Our main contributions are respectively formulating the frequency domain normal map filtering theory and developing spherical harmonic exponentiation techniques. Our analytical, frequency and wavelet models have laid important foundations for investigation in more general real-time algorithms. More importantly, our analysis of the theoretical characteristics of light transport and reflectance and our hybrid approach that couples numerical precomputation with compact analytic models will open up new perspectives for important real-rendering problems that are conventionally viewed as difficult.

Department News & Awards

On Friday, February 5, **Columbia University** hosted a site for the 2008 North American Olympiad in Computational Linguistics for high school students. A select group of universities across the country (Columbia, Cornell, University of Pennsylvania, CMU, University of Michigan, Brandeis University, and others) participated in this event, which was sponsored by the National Science Foundation. The February 5 event was an open competition; students who did well enough qualified to advance to the Invitational Competition. The winners of the Invitational Competition represent the United States and Canada at the International Linguistics Olympiad.

Professor **Luca Carloni** was awarded a prestigious 2008 Sloan Foundation Fellowship. The Alfred P. Sloan Foundation named 118 outstanding young scientists, mathematicians, and economists as Alfred P. Sloan Research Fellows for 2008. The winners are faculty members at 64 colleges and universities in the United States and Canada who are conducting research at the frontiers of physics, chemistry, computational and evolutionary molecular biology, computer science, economics, mathematics and neuroscience. The Sloan Research Fellowships have been awarded since 1955. Only 16 young computer scientists in the United States and Canada were honored with the fellowship this year. Prof Carloni's research is centered on the development of system-level design technologies to assist engineers in assembling high-performance integrated circuits and building distributed embedded systems.

Professor **Julia Hirschberg** was given an honorary doctorate at KTH (The Royal Institute of Technology) in Stockholm on November 16 at their graduation and promotion ceremony as "one of the leading world experts in speech communication, with experience from industrial as well as academic research." The full citation reads as follows (translated from the Swedish original): "Julia Hirschberg, professor in computer science, is active within the area of speech communications at Columbia University, USA. She belongs to the leading researchers in this field, having performed research in both industry and academia. In her work at AT&T, she contributed to the development of several voice-controlled telephone services. Julia Hirschberg has performed leading research on a variety of topics related to human-to-human and human-to-machine interaction. Specifically, within the area of prosody, she studied how people use other means than speech to communicate focus, turn-taking and emotions in a dialogue. She has also studied how this knowledge can be applied to various speech-based services. Julia Hirschberg has been president of the International Speech Communication Association (ISCA) since 2005. As such she is responsible for the yearly conference Interspeech that attracts more than 1000 attendees each year."

Professor **Angelos Keromytis** was elevated to the rank of "senior member" in the Association for Computing Machinery (ACM).

Professor **Vishal Misra** has received the IBM Faculty Award for 2007. Professor Misra's proposal seeks to develop and analyze Adaptive Sharing Mechanisms (ASMs) in which the mechanism used to share resources adapts dynamically to both the set of available resources and the current needs of the consumers, such that the system is truly autonomic. The project proposes to modularize the ASM into separate components, and then design the various components using both cutting edge novel control theoretic and scheduling analyses.

Visiting Professor **Dragomir Radev** coached the U.S. team of eight high school students that participated in the 2007 International Linguistics Olympiad in St. Petersburg, Russia in August. This year's International Olympiad featured 15 teams representing 9 different countries, including the Netherlands, Russia and Spain. Competitors were given problem sets consisting of sentences in languages most people are not familiar with, including: Tatar; Georgian; a language spoken by indigenous people in Bolivia called Movima; the Papua New Guinean language Ndom; Hawaiian; Turkish; and their English translations. With just this information, the competitors then had to translate more sentences from these languages into English. U.S. team members garnered many honors, including the highest score of all participants in the individual competition, top prize for the best solution to one of the problems, and top prize in the team competition in a tie with a Russian team.

Rajesh Ramakrishnan was selected as a CRA Undergraduate Award finalist, among only sixteen students for the whole country. The Computing Research Association (CRA) is an association of more than 200 North American academic departments of computer science, computer engineering, and related fields; laboratories and centers in industry, government, and academia engaging in basic computing research; and affiliated professional societies.

Professor **Henning Schulzrinne** was selected as an IEEE Communications Society distinguished lecturer. Distinguished lecturers visit IEEE Communications Society chapters to discuss new developments in communications and networking.

Professor **Yechiam Yemini** and Professor **Henning Schulzrinne** were honored with the lifetime entrepreneurial and innovation awards, respectively, during the Center for Advanced Technology in Telecommunications and Distributed Information Systems' 25th anniversary celebration on November 9, 2007.

According to the citation, "Prof. Yechiam Yemini is that rare individual who embodies excellence in research, innovation and entrepreneurship. He was already a successful entrepreneur before he joined CATT. He then started System Management Arts or SMARTS, a company with over 150 employees that developed network management solutions. This company was acquired by EMC Corporation. He is now working on yet another start up called Arootz. In all his ventures he brings technological innovation and an unerring vision of the market."

Henning Schulzrinne was cited as a pioneer in the development of Voice over IP technology that is supplanting circuit-switched voice, which has been the basis of phone service since the days of Alexander Graham Bell. He is a co-inventor of the Session Initiation Protocol (SIP) and the Real-Time Transport Protocol (RTP), which form the basis of VOIP, and additional standards for multimedia transport over the Internet. In addition, Verizon Communication was honored for a joint project conducted with the lab of Prof. Schulzrinne.

The Center for Advanced Technology in Telecommunications and Distributed Information Systems (CATT) is a research and education group at Polytechnic University. CATT researchers are leaders in the fields of electrical engineering and computer science. The Center also draws on the expertise of key researchers at Columbia University.

Professor **Rocco Servedio** was awarded a Phase II grant as part of the DARPA 2007 Computer Science Study Group. The Computer Science Study Group is a program that supports university research in computer science and related fields, while informing a new generation of researchers on Department of Defense (DoD) information technology needs and priorities. The CS Study Group is a multi-year program, consisting of a funded educational experience to familiarize the participants with DoD practices, challenges and risks, and up to three years of funded research to explore and develop technologies that have the potential to transition innovative and revolutionary computer science and technology advances to the government.

Professor **Sal Stolfo** has been invited to participate in the "Committee on Information Assurance for Network-Centric Naval Forces" organized by the Naval Studies Board (NSB) or the National Academies National Research Council. At the request of the Chief of Naval Operations, the Naval Studies Board of the National Academies is planning to conduct a 12-month study entitled "Information Assurance for Network-Centric Naval Forces." The study will review the Department of Defense and the Department of the Navy responsibilities for information assurance, review recent information assurance-related studies conducted by and for the Department of Defense and Department of the Navy, examine the Department of Defense and Department of Navy research, development, and acquisition process for information assurance, and recommend alternative approaches to the process that allow for greater flexibility, assess potential information assurance vulnerabilities for network-centric naval forces, review and recommend information assurance best practices, recommend investment analysis approaches for managing cyber attack risks to network-centric naval forces that address the consequences of possible cyber attacks, the likelihoods of these attacks actually occurring, and the uncertainties surrounding assumptions about these risks.

Alumni News

Halim Abbas (M.S. '06) has been in California since graduating from Columbia. He is working for a startup called Code Green Networks, which specializes in delivering data loss prevention solutions, where he is on their senior design team heading Machine Learning enabled projects. One current project is building a solid ML document classifier with a growing ontology; this may lead to joint projects with Columbia and/or Stanford.

John Andersen (B.S. '98, M.S. '00) and **Megan Pengelly** (M.S. '06) were married on April 27, 2007.

Andrew Arnold (CC '03) is in his fourth year in the Machine Learning Ph.D. program at Carnegie Mellon University in Pittsburgh. He has been working on named entity recognition and transfer learning with his advisor William Cohen, and has enjoyed internships at IBM Watson Research in Yorktown with Naoki Abe and Microsoft Research Asia in Beijing with Hang Li and Tie-Yan Liu.

Simon Metz (B.S. '01) writes, "As a Senior Technical Project Manager at IconNicholson, I am responsible for leading the team through all aspects of strategy, planning, development and implementation of projects. Much of my work begins in evaluating our client's business processes and making improvements that will affect their overall performance and technology. For a top global consulting firm I have managed a nine-month process to upgrade and improve the functionality across a custom .NET Content Management System and an XML banner feed system. I also lead Search Engine Optimization efforts across multiple teams at IconNicholson including Mass General Hospital, Forrest Labs, MasterCard, Valpsar Paints and Nestle Waters.

Previous to joining IconNicholson, I was a Senior Project Manager at Time Inc. Interactive, where I was responsible for managing Web sites and software application development efforts for some of Time Inc's premier sites such as People.com, EW.com, and InStyle.com. I led a team to migrate an internal meta data system which drove our efforts to integrate the semantic web principles into our sites.

Before joining Time Inc. Interactive, I was Software Development Manager at Omnicom Media / C2 Creative where I managed large Web application development projects in content management, digital asset management and kiosk systems development for U.S. Trust, Cablevision, Intel, Philips Lighting, Sony, Comcast, and Daimler Chrysler."

Mike Schiraldi (B.S. '00, M.S. '03) and his wife **Hilary (Gerson)** (CC '00) quit their jobs in November so they could go see the world. They invite their classmates to follow their progress at spothopping.com.

Cesar Vichido (M.S. '07) writes, "I'm working full time at an Internet startup. Sosauce.com is a social networking site that let you share quality content with other users. The site just launched last week, so everybody here is excited and happy!"

The **winner of the Fall 2007 caption contest** for the picture of Angelos Keromytis was **Ken Ross**, with the caption:

"So this is how you install python under RedHat!"

The **winner of the Fall 2007 caption contest** for the picture of various CS department faculty was **Awilda Fosse**, with the caption:

"Are you SURE this is the way you dance the Macarena?"

Two new photos are included for this issue's caption contest. The first photo features Professors Steven Bellovin, Eitan Grinspun, and Alfred Aho, and the second photo features Professors Tal Malkin and Eitan Grinspun. Please send suggestions to *newsletter-editors@lists.cs.columbia.edu*. The winners will be announced in the Fall newsletter.

PHOTO 1

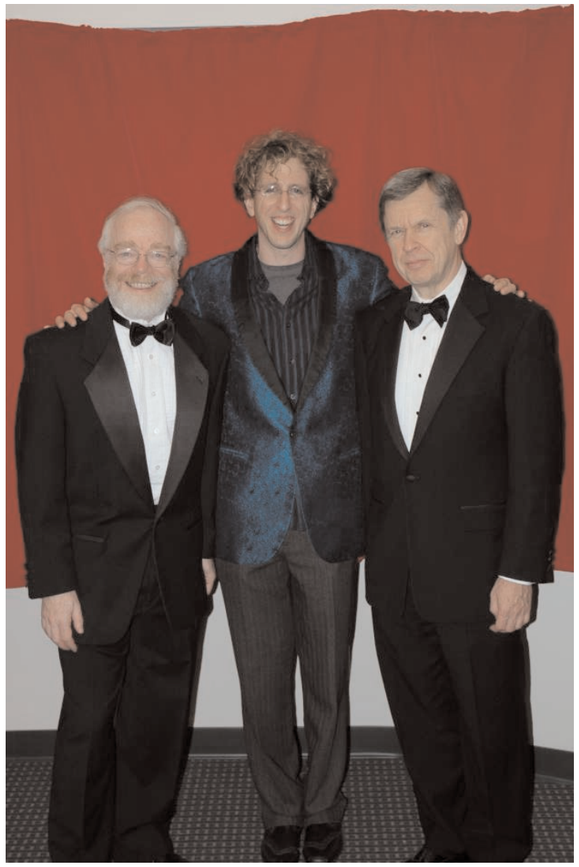

PHOTO 2

CUCS **Resources**

You can subscribe to the CUCS news mailing list at
http://lists.cs.columbia.edu/mailman/listinfo/cucs-news/

CUCS colloquium information is available at
http://www.cs.columbia.edu/lectures

Visit the CUCS alumni portal at
http://www.cs.columbia.edu/people/alumni

Columbia University in the City of New York
Department of Computer Science
1214 Amsterdam Avenue
Mailcode: 0401
New York, NY 10027-7003

ADDRESS SERVICE REQUESTED

NON-PROFIT ORG.
U.S. POSTAGE
PAID
NEW YORK, NY
PERMIT 3593