Statement of Research, Teaching and Service

Surendar Chandra, University of Notre Dame

A picture is worth a thousand words. Technology improvements are finally allowing end users to experience this power by creating and consuming media content from a wide variety of applications. Media objects are large and place enormous resource demands on the systems that support them. My research career has primarily focused on developing systems techniques to extend multimedia technologies to a wide range of scenarios. I focus on three facets of this challenge: quality aware transcoding to adapt content for resource constrained environments, energy conservation mechanisms for popular streaming formats and in building a large scale secure ubiquitous storage infrastructure using commodity devices for multimedia objects. My teaching goals complement my research efforts by emphasizing systems aspects in networks, storage and in multimedia.

My research philosophy centers on choosing practical problems with experimentally demonstrable limitations and then developing techniques to address these challenges. When applicable, I have publicly released the artifacts of my research for follow-on research by the broader community. Frequently, my research has bucked conventional wisdom by demonstrating the limitations of state-of-the-art deployed systems. For example, my research group analyzed the object annotations and user generated queries for such objects in two popular peer-to-peer (P2P) systems and illustrated fundamental limitations of prior P2P systems regardless of the network overlays used. I strive for minimal topic overlap between my publications, except when extending a paper for journal publication. Once I have convincingly demonstrated my techniques, I rarely continue to write papers on that same topic as it would tend to have a lesser impact. As is the convention in my field, the primary authorship of papers rests with students followed by the faculty with author names ordered by the level of contribution of the particular author to the paper.

1 Research

1.1 Quality Aware Transcoding

Users access the Internet from a wide variety of devices with different resource constraints. Transcoding objects has been a popular mechanism for content adaptation. Earlier systems blindly performed transcoding operations by choosing an arbitrary fidelity level leading to an increase in file size as well as unacceptable fidelity loss. I quantified the fidelity tradeoff characteristics and developed the ability to predict resource requirements for transcoding JPEG images. I develop applications of this quantification knowledge to dynamically allocate available resources on a per-client and per-request basis. I demonstrate how a web service can utilize informed transcoding to gracefully degrade content fidelity to different classes of users and in response to changing client demands. I also show the value of this technology in a multimedia capture device to manage the local battery and storage resources. My work has generated 8 publications with 262 citations.

1.2 Energy aware multimedia streaming

Wireless streaming media are popular and place enormous demands on the battery resources of mobile devices. I analyzed the wireless network interface energy consumption for receiving popular streaming formats through deployed infrastructure networks under varying network conditions. I have shown that the IEEE 802.11 power saving mode of operation did not provide energy savings. Analyzing the stream dynamics of these popular streaming formats, I pioneered history-based client and server side strategies to reduce the energy consumed for streaming by transitioning the WNICs to a lower power consuming *sleep* state. This work lead to 7 papers and 111 citations.

1.3 Content aware peer-to-peer (P2P) systems

Recently, P2P systems have been the focus of much research. My group analyzed the contents and queries from a popular P2P system to highlight the fundamental flaws in the assumptions used by prior systems. We have been collecting traffic on the popular Gnutella network since 2003. Based on the observed system behavior, we have developed an unstructured P2P overlay that used the peer sharing capacity and network characteristics to generate an overlay using local information. Our compact and well connected graph provides better search mechanisms using attenuated Bloom filters and random walkers. Also, an analysis of the queries and annotations of objects that are stored in show that the query terms and object annotations exhibit a Zipf like distribution. In practice, as compared to the success rates of 62% achievable using assumptions of uniform distribution, Zipf distributed objects are successfully found only 5% of the times. Also, the relative popularity in the object annotations does not correlate well with its popularity in the query workload. Almost half the queries have no matching objects in the system regardless of the overlay or search mechanism used to locate the objects. We are developing a P2P mechanism that a) quickly rejects queries that have no matching objects in the system and b) routes queries towards the objects replicated in a Zipf distribution. Our work has lead to 5 papers.

1.4 Resource Management in peer computing systems

Peer computing systems federate the *spare resources* available among a set of *independently owned and cooperating* peers for the common good. Popular peer computing systems include wireless ad hoc networks and P2P applications (e.g. Skype, Gnutella and BitTorrent). We are bridging a critical incompatibility between the design principles of operating systems and peer computing systems. Even though there was a wealth of prior research on peer federation mechanisms, the precise resource allocation responsibilities of the peer to the federation was not understood. Operating systems have been unaware of peer resource allocation assumptions and hence peer resource requests have been treated as equal or in the worst case as more important than local requests. In practice, this means that resource constrained peers continued to provide resources to other peers while jeopardizing their own resource availability for local users. To address this mismatch, we are making operating systems explicitly aware of resources consumed by peers. To conserve constrained resources, peer resource requests may go intentionally unsatisfied even though there are no other competing local resource requests. So far, this work was reported in 2 publications.

1.4.1 Mobile stow-away storage

The primary application driver for this project was to disseminate lecture videos created by the instructor as well as annotated videos from students. We have investigated the feasibility of distributing video contents from user devices. The effectiveness of this distribution mechanism depends on the total number of *voluntary* replicas, availability patterns of wireless devices as well as the level of the sharing implementation. In our study, users have been observed to explicitly control sharing durations, presumably to manage the local resource consumption; a behavior that was not considered by prior work that analyzed the viability of using desktops for storage. Using extensive analysis of the observed node behavior, we show that though laptop users were online for shorter durations, their temporal consistency can provide reasonable availability, especially at the times of the day when students are typically active. Our experiments show the inherent limits of epidemic propagation in real campus wireless network scenarios. We have also improved the message propagation rates by carefully selecting the next hop nodes based on past node behavior. We are developing mechanism that will achieve reasonable availability among student groups while still using mobile student laptops for storage. This work is published in 3 papers.

1.4.2 Other topics (each of these have already resulted in one or more publications)

We have introduced a temporal importance abstraction that allowed our storage to scale by reducing the resource allocation for *less important* objects. In our system, the end users provided these importance annotations. We have used this abstraction to manage object persistence. We are investigating its applicability in managing object reliability, availability and security.

We are developing a fully decentralized security infrastructure that provides data availability, access control, and versioning with tunable levels of certainty. The stringency of traditional security guarantees for these operations are minimally relaxed. In exchange the protocols are resilient and can continue to operate correctly despite a fraction of the storage nodes acting maliciously.

Working with two sophomores, we have analyzed the wireless behavior of the popular Nintendo DS. We have shown that they used PCF modes and catastrophically interfere with deployed wireless LAN networks that are based on the DCF mode of operation. Such interference is worse for newer technologies such as 802.11g and 802.11n. The wireless networking community has assumed that PCF modes of operation were not popular and hence the interaction of overlapping PCF and DCF systems has rarely been considered. We are continuing our investigations.

1.5 Experimental Infrastructure Development

I was the first experimental systems faculty hired by our department; building the experimental infrastructure was the critical initial task. I have successfully accomplished this task. I built, administer and maintain a experimental systems lab with adequate servers and hardware resources that were acquired with funding support from VMware, NSF and DIA-MASINT. I worked with the college to dedicate a wireless LAN channel exclusively for research purposes. I am the site PI for Planetlab, a world-wide research infrastructure that is available to the university community. I maintain the hardware as well as work with the general counsels office and OIT to address any legal as well as security concerns associated with operating a Planetlab site. I am the representative for USENIX, the advanced computer systems association. I maintain the university-wide digital library for the proceedings of USENIX systems conferences. Some of our students have also received scholarships from USENIX to attend top conferences.

It is important for undergraduate students to have reasonable familiarity with the hardware and software systems that they will encounter in industry if they are to be employable upon graduation. I built, administer and manage the educational systems lab which consists of six desktops (provided by the department) and a dual and quad processor Itanium2 servers (valued at $40K and $80K and provided by HP Teaching initiative). Students enjoy full freedom to install and modify any software on these machines. I use this lab extensively for my course projects.

1.6 Research supervision

A good systems research program requires the support of qualified student researchers. I work with three excellent students who are in their sixth, third and first year of the Ph.D. program. Each student has taken principal responsibility for developing the ideas and carrying out the experiments for a different research project as part of their Ph.D. dissertation. Recruiting new students is challenging; we receive few qualified graduate applicants who express interest in working in systems related topics. Also, we do not offer a terminal Masters degree. Hence, I have taken a pro-active role in directly recruiting students. This is a resource intensive and sometimes frustrating approach. Three of the students that I recruited rescinded at the last moment (two to top CS programs and one to a fatal car accident). Of the four students who came to Notre Dame, there has been attrition to other attractive opportunities. Fortunately, one proactively recruited student is making excellent progress toward a degree. Recently, I successfully recruited into graduate school a student who took my undergrad OS course in Spring ’06 and an independent study in Summer ’06. I am working with two sophomores whose independent study as led to a publication in Netgames’07.

1.7 A note about experimental systems area

My research methodology is experimental in nature. *Academic Careers for Experimental Computer Scientists and Engineers* [1] describes some of the challenges that are encountered within the experimental systems research community. Programming and creating a novel computational artifact, while not by themselves synonymous with conducting experimental systems research, are often required to demonstrate the practical implications and measured performance of research innovations. Construction of artifacts is labor-intensive because they involve a great deal of low-level detail and they require technically sophisticated system builders capable of working with totally new concepts rather than detailed specifications. It takes several semesters to attract talented systems students and train them in experimentally oriented systems courses. The best available data [2] shows that the average time to complete a systems Ph.D. was over 6.4 years; it takes longer to complete a systems prototype to be used in an extensive evaluation.

Systems research papers are typically published in refereed conferences with low acceptance rates rather than journals. A substantial majority of respondents to the CRA-CSTB survey of systems faculty preferred conferences, as compared to journals, as the means of dissemination by which to achieve maximum intellectual impact because of timeliness and the better focused feedback offered by conference audiences. The reports also warn against the dangers of a less influential approach to maximize publication count by adopting a “least publishable unit” strategy in which the smallest possible increments of progress are published at frequent intervals. As I continue to build a quality systems program, I have tried to avoid the pitfalls mentioned in the report while continuing to publish with meaningful productivity.

2 Teaching

I deeply value teaching and wish to convey my enthusiasm for computer science in general and experimental systems in particular to all of my students. As a fast-moving and technologically sensitive area, I relish the challenge of maintaining the vitality of curricula using constant updates and drawing on state-of-the-art examples from topical research domains. Also, earlier research had shown the value of lecture videos as an effective review tool. For the past three semesters, I have captured my lectures on video and made them publicly available on the web, podcast as well as made available in Google video. I have shown that technology improvements can allow any instructor to perform these operations with reasonable effort. I presented my positive educational experiences as well as the pitfalls of video capture in a colloquium at the Kaneb Center for teaching and learning at Notre Dame. These video access traces were the workload for my storage research.

At Notre Dame, I have developed a new grad/ugrad and ugrad course in *Computer Networks*, redesigned the core *Graduate Operating Systems* and ugrad *Operating Systems* course. I also developed a new advanced grad course in *Distributed storage* and grad/ugrad course in *Multimedia Systems*. I am currently developing a grad/ugrad course on *Networked Sensor Systems* for Fall ’07. Since our department is small, student enrollment requires that these specialized courses be offered every three years or so. My courses are project oriented with a significant course project. Some students who can commit to significant time and effort beyond the course have widened the scope of their course projects into research publications. Two course projects led to publications and two other projects have formed the basis for continuing dissertation research. Next, I briefly summarize a core ugrad course that I taught at Notre Dame.

2.1 Junior level core Operating Systems

I wish to strike a balance between describing fundamental concepts while introducing variants that are more appropriate for contemporary hardware (e.g., modern CPUs are multi-core, making multi-threading an indispensable concept). Operating systems manage system resources such as CPU, memory and storage. Each of these subsystems uses variants of the same fundamental techniques. In order to make these similarities and differences clear, I organize my course around five different modules: *Process Management*, *Process Synchronization*, *Memory Management*, *Storage Management* as well as a module that focuses on understanding the interplay between these various concepts, especially in the context of servers, desktops, laptops and PDAs. Student understanding in each module is evaluated using a quiz, homework assignment, homework projects and a module exam for each module as well as a final exam.

3 Service

Besides serving in departmental committees and as external dissertation chairs, I took an active role in advancing systems related artifacts within the university (e.g., USENIX, Planetlab). I lead the design of an advertisement poster to improve our visibility among potential student applicants. I coordinate the department colloquium and help improve our visibility among research peers.

My stature in the field is indicated by invitations to serve on nine NSF review panels and over fifty conference program committees. I have served on the program committee for a number of prestigious conferences such as ACM Multimedia, IEEE ICDCS and IEEE INFOCOM. Twice, I served as the technical program committee co-chair of ACM/SPIE MMCN. In this role, I recruited and worked with a strong committee and significantly improved the number of quality submitted papers; much more than was typical for this conference. IEEE Computer society recognized some of my efforts with a certificate of appreciation in 2004.

4 Summary

I have made significant contributions in the areas of multimedia systems, including the storage and energy-efficient transmission of multimedia objects. I have also provided important insights in peer-to-peer resource management. My group is addressing future research challenges motivated by a vision of ubiquitous access to multimedia content with a light footprint on resource consumption. I achieve impact by disseminating results in excellent, highly visible venues. I have laid the foundations for continuing contributions by building an environment for innovative systems research at Notre Dame. A strong program requires a group of qualified student researchers who are working on research topics that are well within their capabilities along with adequate research infrastructure. I have built a group of enthusiastic graduate students who are productive and in various stages of their career (one about to graduate). I manage a research lab with adequate resources that were acquired with funding support from HP, VMware, NSF and DIA-MASINT.

References

[1] **A**cademic Careers for Experimental Computer Scientists and Engineers. *Committee on Academic Careers for Experimental Computer Scientists*, *Computer Science and Telecommunications Board*, *Commission on Physical Sciences, Mathematics, and Applications*, and *National Research Council*. National Academy Press, Washington, D.C., 1994. Online for free: `http://books.nap.edu/catalog.php?record_id=2236`.

[2] **C**omputing the Future: A Broader Agenda for CSE. *Computer Science and Telecommunications Board*, National Academy Press, Washington, D.C., 1992. `http://books.nap.edu/catalog.php?record_id=1982`.