

Statement of Research

Prof. Surendar Chandra, University of Notre Dame

My research philosophy centers on choosing practical problems with experimentally demonstrable limitations and then developing techniques to address these challenges. My group is addressing future research challenges motivated by a vision of ubiquitous access to multimedia content with a light footprint on resource consumption. Frequently, my research has bucked conventional wisdom by demonstrating the limitations of state-of-the-art deployed systems. For example, we analyzed the object annotations and user generated queries for such objects in two popular peer-to-peer systems and illustrated fundamental limitations of prior P2P systems regardless of the network overlays used.

1 Prior research

Media objects are large and placed enormous resource demands on the systems that support them. I focused on two facets of this challenge: quality aware transcoding to adapt content for resource constrained environments and energy conservation mechanisms for popular streaming formats.

1.1 Quality Aware Transcoding

Users access the Internet from a wide variety of devices with different resource constraints. Transcoding objects had 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 developed applications of this quantification knowledge to dynamically allocate available resources on a per-client and per-request basis. I demonstrated 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 showed the value of this technology in a multimedia capture device to manage the local battery and storage resources. My work had generated 7 publications [1, 2] with 332 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 showed 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 [3, 4] and 147 citations.

2 Secure, ubiquitous storage

Recently, the cost of unmanaged resources such as raw storage, network capacity and computing power are plummeting. We are developing techniques to harness these resources.

2.1 Mobile collaborative storage

Resource rich laptops are finally becoming ubiquitous. In an article published in 2006, USA today described the emergence of about 30 million American nomadic users. Gartner Dataquest predicted a yearly growth of 10% of these users. There had been a rich body of prior work on mobile and wireless sharing systems. Exploiting the recent explosion in the number of wireless users in our university, we empirically investigated the behavior of various sharing systems among wireless users. Unlike prior systems, we analyzed the behavior of a large number of users at different levels including application level with explicit user control over sharing durations. We showed that the propagation behavior was far worse than predicted by mobility models assumed by prior systems. We showed that direct delivery mechanisms require an inordinately large number of replicas to provide good availability and hence are not practical. The performance of gossip based epidemic propagation was poor. However, the system behavior was limited by the availability of wireless users. We showed that the performance of a laptop based solution did not differ significantly from a system that used a storage infrastructure; lending credence to the choice of using storage from other laptops. Our work paves the way for understanding the feasibility of any particular wireless sharing system. We are developing sharing systems that are practical for the observed user behavior.

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

P2P systems have been the preferred mechanism to federate storage resources. My group analyzed the contents and queries from a popular P2P system to highlight the fundamental flaws in the assumptions used by prior systems. Since 2003, we had collected traffic on the popular Gnutella network. 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 had no matching objects in the system regardless of the overlay or search mechanism used to locate the objects. We developed a P2P middleware that transparently transforms the user queries in order to improve the query success rates. Our experimental analysis showed an success rate of over 60% for rewritten queries in Gnutella networks.

2.3 Resource Management in peer computing systems

Our work is attempting to bridge a critical incompatibility between the design canons of operating systems and peer computing systems. Popular peer computing systems include wireless ad hoc networks and peer-to-peer (P2P) applications (e.g. Skype, Gnutella and BitTorrent). Operating systems were traditionally tasked with managing the system resources among the different schedulable entities (processes). Peer computing systems federate the *spare resources* available among a set of *independen-*

dently owned and cooperating peers for the common good. 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. We are developing a scheduling class called *less than best effort* to address this problem. Under resource constrained scenarios, resource requests from this class can either be denied or delayed indefinitely even if such a request from traditional scheduling classes can be allocated resources.

2.4 Managing scalable long term storage

We are addressing the problem of continual resource management in a system that stored contents across a large number of devices for long durations. We are concerned with resources such as total storage space, object reliability, availability and security. For example, to manage the storage space, systems require the ability to continually increase the total storage space and match the storage requirements of new objects. Traditional storage systems provide persistence for all stored objects to perpetuity, relegating object reclamation to applications. However, not all of the data are equally important; less important objects could be automatically replaced by more important ones. We developed a simple and expressive temporal *importance* function. We have used this abstraction to manage object persistence. This information allows the storage system to selectively weaken the performance guarantees offered by the storage. Using extensive simulations and observations of a university wide lecture video capture and storage application, we show that our abstraction allows the users to control the amount of persistence for each individual object. We are investigating mechanisms that can weaken the performance of other resources as well.

2.4.1 Other topics

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 point coordination function (PCF) modes and catastrophically interfere with deployed wireless local area networks that are based on the distributed coordination function (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.

At Notre Dame, I manage a research lab with adequate resources that were acquired with funding support from HP, VMware, NSF and DIA-MASINT. 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 Multimedia Computing and Networking (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 my efforts with a certificate of appreciation in 2004.

References

- [1] **JPEG Compression Metric as a Quality Aware Image Transcoding.** *Surendar Chandra and Carla Schlatter Ellis.* Second Usenix Symposium on Internet Technologies and Systems (USITS '99), pages 81–92, Boulder, CO, October 1999. Citations: 76.
- [2] **Application-Level Differentiated Multimedia Web Services Using Quality Aware Transcoding.** *Surendar Chandra, Carla Schlatter Ellis and Amin Vahdat.* *IEEE JSAC - Special Issue on QoS in the Internet, Vol 18, No. 12, pages 2544–2565, December 2000.* Citations: 142.
- [3] **Application-specific network management for energy-aware streaming of popular multimedia formats.** *Surendar Chandra and Amin Vahdat.* *USENIX Annual Technical Conference*, pages 329–342, Monterey, CA, June 2002. Citations: 71.
- [4] **Wireless Network Interface Energy Consumption: Implications for Popular Streaming Formats.** *Surendar Chandra.* *ACM Multimedia Systems Journal* 9(2), pages 185–201, Springer-Verlag Heidelberg, August 2003. Citations: 56.