Research Statement of Abhigyan Sharma

My research studies networked systems with a focus on enhancing delivery of content and services over the Internet. The two broad topics of my current research are convergent network and content delivery infrastructure and geo-distributed cloud computing platforms.

Convergent network and content delivery infrastructure: To facilitate content delivery on the Internet, the idea of a convergent network and content delivery infrastructure is gaining popularity. Content-oriented network architectures have been proposed that promise enhancements in content delivery, security and mobility. Network operators have deployed content delivery infrastructure on their network to leverage the benefits of content caching and to monetize content delivery services. These convergent infrastructures represent a paradigm shift in which the role of content delivery is entrusted to network operators.

My research investigates schemes for network management in these convergent infrastructures, i.e., techniques for content delivery and network routing towards optimizing cost and performance objectives. The challenge is that these objectives depend not only on the routing and the content delivery techniques that are deployed but also on the interaction between them. The distinguishing aspect of my work is to study the role of content placement schemes on this interaction, an aspect largely overlooked in prior work. There are two main findings of my work. First, I show that combinations of simple placement and routing schemes are effective in optimizing cost and performance objectives, thereby simplifying the task of network management for the operators. Second, I show that content placement matters more than routing: cost and performance objectives depend much more on the content placement scheme compared to the routing scheme. The implication is that effective content placement is the key to network management in a convergent network and content delivery infrastructure.

Geo-distributed cloud computing platforms: Geo-distributed cloud computing platforms make it possible to deploy applications across global datacenters and to dynamically provision resources within a datacenter. These platforms can potentially reduce user-perceived latency via geo-distributed deployment and reduce cost by "right-sizing" the provisioned resources to the current demand. But, existing applications do not always realize these benefits due to several reasons, which include lack of controls to dynamically provision resources, and sub-optimal data placement across geo-distributed locations. Addressing these limitations of existing applications is important given the growing popularity of cloud computing.

In my work, I design systems that use the flexibility available in these platforms to provide better cost-performance tradeoffs than existing designs. Specifically, I design *controllers* that reconfigure a system in response to changes in demand and reduce the overhead and the performance impact of reconfiguration to provide significant improvements in cost and performance metrics. I have used this approach to design Auspice, a geo-distributed data store, and Shrink, a system to reduce energy use for of content datacenters. Auspice is the first system that automatically places dynamic data items at geo-distributed locations based on demand locality while maintaining consistency. Shrink is the first effort at characterizing energy-performance tradeoff for content datacenters, and proposes novel techniques that jointly optimize server and network energy use and its cost for these datacenters. My research demonstrates that leveraging the flexibility of cloud computing can improve networked systems in several aspects such as cost, performance and energy efficiency.

My research takes an empirical approach with a strong emphasis on implementation and realistic performance evaluation. I have found that the implementation process often raises new research questions. For example, providing strong consistency guarantee in the Auspice data store required a memory-efficient implementation of the Paxos consensus algorithm, which has spawned off into an independent research project in our lab. A realistic performance evaluation of Internet-scale systems is a daunting task, but I have found that it has potential to yield new insights. One such instance is from my research comparing several proposed network routing schemes. For a realistic comparison, I designed a large-scale simulation of the traffic in an Internet Service Provider network. Surprisingly, I found that despite variation in network link utilization, all schemes provide near-identical end-user performance. In many of my projects, I have collected and used extensive real world datasets of content demand, network topologies, and network traffic matrices from commercial CDNs and network operators. The use of real-world datasets helps capture the reality of end-user behavior and strengthens my conclusions.

Research contributions

Convergence of network and content delivery

My research addresses questions that are central to the management of a convergent network and content delivery infrastructure. How does the interaction between placement and routing schemes used today affect operators' cost and performance objectives? How do existing schemes compare to the best-possible placement and routing schemes? What is the relative important of placement vs. routing? I have studied these questions in two settings with varying degrees of flexibility in placing content as explained below.

Content location diversity: Consider a network in which content is placed at a fixed number of randomly chosen locations. I have conducted a large-scale experimental study based on real network topologies and traffic matrices that compares the benefit of placing content at multiple locations compared to a single location. I have shown that adding a small number (2-4) of additional content locations increases effective network capacity by up to 2X, reduces the difference between the optimal routing and a *static* shortest-path routing to at most 30% and improves performance of sub-optimal *dynamic* routing schemes so that they yield near-optimal results. Further, my work measures the user-perceived impact of network routing schemes in detail and represents a significant departure from the existing trend of evaluating network routing schemes based on link utilization metrics. A surprising finding is that, for current traffic demands, there is little difference in end-user performance irrespective of which network routing scheme is used.

Network CDN: A *network CDN* (NCDN) is a network provider that operates a CDN deployed across its point of presence locations and enjoys full control over placement and routing on its network. Of particular interest to NCDN operators are *demand-oblivious* schemes, as they make placement and routing decisions without measurement of content demand and simplify network management for the operators. My work has demonstrated the effectiveness of a demand-oblivious scheme – a combination of distributed caching and static shortest-path routing – based on experiments with extensive traces from Akamai CDN. First, the demand-oblivious scheme significantly outperforms a scheme that optimizes placement and routing *jointly* based on historical content demand and performs close to an ideal joint-optimization strategy with knowledge of future content demand. Second, optimizing routing matters little in NCDNs: whether a demand-oblivious routing (a static-shorted path routing) or a routing computed based on traffic matrix is used along with a demand-oblivious placement scheme, the network cost differs by less than 10%.

To summarize, operators have traditionally focused on optimizing routing to achieve their objectives, but my research shows that in a convergent network and content delivery infrastructure, they are better served by leveraging placement flexibility using simple demand-oblivious placement schemes.

Geo-distributed cloud computing platforms

Auspice: Auspice is a data store for dynamic content deployable across geo-distributed locations. Auspice's design addresses a fundamental cost-performance tradeoff in the placement of dynamic content: placing content at multiple locations improves latency of content accesses but increases update propagation costs. In comparison to existing data stores that require manual placement configuration or support only a static placement policy, Auspice automatically places content by inferring pockets of high demand for a name and using a heuristic placement scheme that provides low request latency with a low update cost of maintaining consistent data replicas. My experimental evaluation of Auspice in the context of a representative application—a global name service that provides network address of mobile entities (simplistically, a DNS for mobile devices)—demonstrates Auspice's latency and cost improvements over commercial managed DNS services as well as DHT-based designs. As a global name service, Auspice is being used in multiple research projects within our university including an end-user socket library with mobility and multipath support (MSocket) and a service to provide context-aware network communication.

Shrink: Shrink is a system that reduces the energy use for *content datacenters*, an important class of datacenters that are used primarily for storing and serving content to end users. My approach dynamically adjusts the set of servers and switches to be kept active with a modest performance penalty over static resource provisioning based on peak demand. While dynamic provisioning has been explored before, Shrink's novel aspect is that it provisions servers and switches using a joint optimization approach that further reduces datacenter energy use. My research is the first to accurately quantify the impact of energy optimization on end-user performance as it models the effect of server load balancing, cache hits and misses and non-steady state behavior due to on-off transitions for servers and switches; the use of detailed request-level workloads collected from a commercial CDN datacenter (Akamai) is also the first. My findings encourage deployment of energy optimization schemes in content datacenters and motivate future research in techniques to lessen the impact of energy optimization on end-user performance.

Future research

My research vision is to architect a cost-effective and high performance infrastructure for delivering content and services on the Internet. Below, I outline specific research topics towards accomplishing my goal that build upon my previous research.

Architecting network edge for content delivery: Recent research (including my work on NCDNs) has shown that much of the benefits due to in-network caching can be achieved by caching content on the network edge. Still, many questions remain unanswered regarding the design of caching infrastructure on the network edge. What should be the granularity of cache placement - individual homes or a neighborhood, or a city? Which algorithms for content search are effective in locating content from among nearby caches? What would be the computational and storage requirements for these caches? What, if any, changes do we need to the design of switches and routers to interface with cache servers? Addressing these questions presents avenues for novel measurement studies, prototype development and deployment studies in organizations.

Unifying edge and cloud resources to optimize content delivery: I am keenly interested in designing a unified content delivery infrastructure that leverages resources on the network edge along with large-scale cloud datacenters. A critical question is to decide which functionality to keep at the network edge and what to process at cloud datacenters towards optimizing end-to-end performance? For example, more resource-intensive processing could be handled by the cloud datacenters whereas edge resources can handle lightweight computation. While commercial CDNs today do deploy a combination of edge and cloud computing to accelerate content delivery, the design challenge is to expose this functionality as a common network service available to all applications.

Datacenter energy efficiency: Given that energy is an important factor in the operational costs of datacenter, an important design goal for networked systems in future would be energy-efficiency. Recent research (including my work on content datacenters) has shown that software-level solutions based on consolidation reduce datacenter energy use. But, many popular networked systems running inside datacenters were not designed to be energy-efficient originally and hence may not interoperate well with consolidation-based energy optimization. A key open question is to evaluate the performance impact of energy optimization on various classes of networked systems such as network file systems, distributed data stores and distributed memory caches, and if needed, to re-architect these systems keeping energy efficiency as a primary design goal. I hope that this line of research would yield design principles that could be broadly applied to the design of energy-efficient networked systems.