

UTKARSH GOEL

ugoel@akamai.com +1 (406) 600-5170 <https://scholar.google.com/citations?user=RWimcU8AAAAJ&hl=en>
799 Market Street, 6th Floor, San Francisco, CA 94103

PROFESSIONAL PREPARATIONS

Doctor of Philosophy (Ph.D.) in Computer Science

Montana State University

Outstanding Ph.D. Researcher Award

Research Advisors: Professor Mike P. Wittie (MSU) and Dr. Moritz Steiner (Akamai)

May 2017

Bozeman, MT, USA

Master of Science in Computer Science

Montana State University

GPA: 3.77 on a scale of 4.00

May 2017

Bozeman, MT, USA

Bachelor of Technology in Computer Engineering

Aligarh Muslim University

GPA: 8.3 on a scale of 10.0

May 2012

Aligarh, India

RESEARCH INTERESTS

My research interests span the areas of Mobile Web performance optimization, real user measurements of client perceived performance, closest content server selection techniques, detecting and understanding the performance impacts of cellular proxies on mobile Web experience, and techniques to accelerate third party content.

PROFESSIONAL APPOINTMENTS

Architect, Akamai Technologies, San Francisco, CA

July 2017 - Present

- Member of the Akamai's cutting-edge arm of Web Performance business unit, the Foundry Web Performance team.
- Responsible for exploring and prototyping of new and faster Web technologies for improved performance.
- Participate and collaborate with field experts at various industrial and academic institutions through networking conferences and technical summits.
- Perform extensive investigation on Akamai's global infrastructure of content delivery to diagnose and address challenging Web performance problems.
- Explore opportunities to extend Akamai's infrastructure for new and innovative technologies, such as Internet of Things and Augmented Reality.

Web Performance Research Engineer (Intern), Akamai Technologies, San Francisco, CA

Summers of 2015 and 2016

Mentors: Dr. Moritz Steiner, Stephen Ludin, Erik Nygren, Wontaek Na, and Martin Flack

- Led investigations on Web performance bottlenecks pertaining to HTTP/2 and IPv6 protocols in a large-scale content delivery network.
- Performed measurement studies to understand impact of cellular network proxies and third party scripts on Web performance.

- Collaborated with several architects and performance engineers to refine my research directions and develop new techniques to improve mobile Web experience.
- Published and presented novel research at top-tier mobile networking and systems conferences.

Graduate Research Assistant, Networks Lab, Montana State University, Bozeman, MT

August 2012 to May 2017

- Pursued several research problems in area of content server selection to improved client perceived Web performance.
- Collaborated with well-reputed researchers at various academic institutions on several research projects.
- Mentored several graduate and undergraduate students in their research projects.
- Invited for presenting novel research at departmental seminars and advanced graduate classes.

Graduate Teaching Assistant, Montana State University, Bozeman, MT

August 2012 to May 2017

- Assisted undergraduate students with in-lab assignments.
- Conducted occasional lectures to discuss certain networking topics in detail.
- Provided recommendations to the faculty for refining the structure of course material to promote student participation in the class.

Summer Associate, LetsLiveWire Inc., Houston, TX

Summer of 2013

- Worked with a startup to develop a Tele-Presence Web application for Tax and Health related services.
- Developed and launched incremental product functionality through scheduled deployments.
- Applied agile software development methodologies to increase both personal and team efficiency.

PEER REVIEWED JOURNAL AND CONFERENCE CONTRIBUTIONS

1. Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. "Measuring What is Not Ours: A Tale of 3rd Party performance", in the proceedings of ACM Passive and Active Measurements Conference (PAM), Sydney, Australia, March 2017.
2. Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. "A Case for Faster Mobile Web in Cellular IPv6 Networks", in the proceedings of ACM Conference on Mobile Computing and Networking (MobiCom), New York, October 2016.
3. Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. "Poster: HTTP/2 Performance in Cellular Networks", in the proceedings of ACM Conference on Mobile Computing and Networking (MobiCom), New York, October 2016.
4. Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. "Detecting Cellular Middleboxes using Passive Measurement Techniques", in the proceedings of ACM Passive and Active Measurements Conference (PAM), Heraklion, Greece, March 2016.
5. Utkarsh Goel, Mike P. Wittie, KC Claffy, and Andrew Le. "Survey of End-to-End Mobile Network Measurement Testbeds, Tools, and Services", in the proceedings of IEEE Communications Surveys and Tutorials, Volume: 18, Issue: 1, First quarter 2016.

6. Utkarsh Goel, Mike P. Wittie, and Moritz Steiner. "Faster Web through Client-assisted CDN Server Selection", in the proceedings of IEEE International Conference on Computer Communications and Networks (ICCCN), Las Vegas, NV, August 2015.
7. Samuel Micka, Utkarsh Goel, Hanlu Ye, Mike P. Wittie, and Brendan Mumey. "PCP: Internet Latency Estimation Using CDN Replicas", in the proceedings of IEEE International Conference on Computer Communications and Networks (ICCCN), Las Vegas, NV, August 2015.
8. Utkarsh Goel, Clint Cooper, Brittany T. Fasy, and Mike P. Wittie. "A First Look at Web Browsing Predictions using DNS Logs", in the proceedings of Conference on Software Engineering and Data Engineering (SEDE), Denver, CO, September 2016.
9. Utkarsh Goel, James Espeland, Upulee Kanewala, and Mike P. Wittie. "Quality Assurance of a Mobile Network Measurement Testbed Through Systematic Software Testing", in the proceedings of Conference on Software Engineering and Data Engineering (SEDE), Denver, CO, September 2016.
10. Utkarsh Goel, Clemente Izurieta, and Mike P. Wittie. "Understanding Factors Influencing the Citation Count of Networking Conference Papers", in the proceedings of Conference on Software Engineering and Data Engineering (SEDE), San Diego, CA, October 2015.
11. Utkarsh Goel, Ajay Miyapuram, Mike P. Wittie, and Qing Yang. "MITATE: Mobile Internet Testbed for Application Traffic Experimentation", in the proceedings of Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS), Tokyo, Japan, December 2013.
12. Utkarsh Goel, Ajay Miyapuram, Mike P. Wittie, and Qing Yang. "MITATE: Mobile Internet Testbed for Application Traffic Experimentation", (poster paper) in the proceedings of USENIX Symposium on Networked Systems Design and Implementation (NSDI), Lombard, IL, April 2013.
13. Mohammed Abdul Qadeer, Kanika Shah, Utkarsh Goel. "Voice - Video Communication on Mobile Phones and PCs' using Asterisk EPBX", in the proceedings of IEEE Conference on Communication Systems and Network Technologies (CSNT), Rajkot, India, May 2012.
14. Utkarsh Goel, Kanika Shah, Mohammed Abdul Qadeer. "The Personal SMS Gateway", in the proceedings of IEEE Conference on Communication Software and Networks (ICCSN), Xi'an, China, May 2011.

OTHER SIGNIFICANT CONTRIBUTIONS

1. Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Stephen Ludin, and Martin Flack. "Domain-Sharding for Faster HTTP/2 in Lossy Cellular Networks", arXiv preprint arXiv:1707.05836, July 2017.
2. Utkarsh Goel. "Internet Measurements and Application Layer Optimizations for Faster Web Communications", Doctoral Thesis, Montana State University, Bozeman, MT, April 2017.
3. Utkarsh Goel, Moritz Steiner, Wontaek Na, Mike P. Wittie, Martin Flack, and Stephen Ludin. "Are 3rd Parties Slowing Down the Mobile Web?" in the proceedings of ACM S3 Workshop, New York, October 2016.
4. USENIX; login: issue: August 2013, Volume 38, Number 4, in the technical summary of USENIX Symposium on Networked Systems Design and Implementation (NSDI), Lombard, IL, April 2013.

RECENT RESEARCH PROJECTS

1. Understanding page abandonment for slow loading websites

Content providers (CPs) make contracts with CDN operators to improve website performance. However, when webpages do not load fast, users tend to abandon the page and move to a competitor website. Understanding users' expectations in terms of how fast the page must load is very important for both CPs and CDN operators. I am investigating user behavior with slow loading websites, considering different factors such as time of day, access network type, and the website category. Preliminary results suggest that users tend to

retain on pages whose first paint happens within 2 seconds, even when the complete page loads several seconds later. Other results indicate faster abandonments on laptops, compared to mobile devices.

2. **HTTP/2-aware domain-sharding in lossy cellular networks**

HTTP/2 (h2) is a new standard for Web communications that already delivers a large share of Web traffic. Unlike HTTP/1, h2 uses only one underlying TCP connection. In a cellular network with high loss and sudden spikes in latency, which the TCP stack might interpret as loss, using a single TCP connection can negatively impact Web performance. In this work, I performed an extensive analysis of real world cellular network traffic and design a testbed to emulate loss characteristics in cellular networks. I used the emulated cellular network to measure h2 performance in comparison to HTTP/1.1, for webpages synthesized from HTTP Archive repository data. My results show that, in lossy conditions, h2 achieves faster page load times (PLTs) for webpages with small objects. For webpages with large objects, h2 degrades the PLT. I devised a new domain-sharding technique that isolates large and small object downloads on separate connections. Using domain-sharding, I demonstrated that under lossy cellular conditions, h2 over multiple connections improves the PLT compared to h2 with one connection and HTTP/1.1 with six connections. Finally, I made recommends to content providers and content delivery networks to apply h2-aware domain-sharding on webpages currently served over h2 for improved mobile Web performance.

3. **Application Layer Optimizations for Faster Web Communications**

The evolution of Web technologies enables interactive Web communications and makes the Web ecosystem more complex. To ensure timely delivery of Web content, the Web Performance Community (WPC) - comprised of browser vendors, content providers, content delivery networks (CDNs), and network regulators - develops new protocols and optimization techniques. However, new protocols suffer from insufficiently wide adoption and the optimization techniques often require ISP support. To cope with these challenges, I presented several measurement techniques through which WPC could better understand the current state of the Web performance. I also present several application-layer optimizations that enable applications to control how content is delivered in different networks. This work summarizes several best-practices, which have been extensively evaluated on production infrastructure, to which the WPC could and should transition to achieve faster Web communications.

4. **Speeding 3rd party downloads for faster Web experiences.**

Content Providers make use of, so called 3rd Party (3P) services, to attract large user bases to their websites, track user activities and interests, or to serve advertisements. In this work, I performed an extensive investigation on how much such 3Ps impact the Web performance in mobile and wired last-mile networks. I developed a new Web performance metric, the 3rd Party Trailing Ratio, to represent the fraction of the critical path of the webpage load process that comprises of only 3P downloads. My results show that 3Ps inflate the webpage load time (PLT) by as much as 50% in the extreme case. Using URL rewriting to redirect the downloads of 3P assets on 1st Party infrastructure, I demonstrated speedups in PLTs by as much as 25%.

5. **Measuring and improving Web performance in cellular IPv6 networks**

The transition to IPv6 cellular networks creates uncertainty for content providers (CPs) and content delivery networks (CDNs) of whether and how to follow suit. Do CPs that update their CDN contracts to allow IPv6 hosting achieve better, or worse performance in mobile networks? Should CDNs continue to host mobile content over IPv4 networks, or persuade to their CP customers the performance benefits of IPv6 content delivery? In this work, I answered these questions through a comprehensive comparison of IPv4 and IPv6 mobile Web performance in cellular networks in the US from the point of view of Akamai's content delivery infrastructure. My data show that IPv6 hosting outperforms legacy IPv4 paths in mobile Web. My analysis leads to clear recommendations for CPs to transition to IPv6-hosted mobile Web. Finally, I proposed new mechanisms, through which CDNs can safely transition mobile content to IPv6-enabled servers for improved content delivery.

6. Detecting transparent Web proxies in mobile networks

The Transmission Control Protocol (TCP) follows the end-to-end principle – when a client establishes a connection with a server, the connection is only shared by two physical machines, the client and the server. In current cellular networks, a myriad of middleboxes disregard the end-to-end principle to enable network operators to deploy services such as content caching, compression, and protocol optimization to improve end-to-end network performance. If server operators remain unaware of such middleboxes, TCP connections may not be optimized specifically for middleboxes and instead are optimized for mobile devices. I argue that without costly active measurement, it remains challenging for server operators to reliably detect the presence of middleboxes that split TCP connections. In this work, I presented three techniques (based on latency, loss, and characteristics of TCP SYN packets) for server operators to passively identify Connection Terminating Proxies (CTPs) in cellular networks, with the goal to optimize TCP connections for faster content delivery. Using TCP and HTTP logs recorded by Content Delivery Network (CDN) servers, I demonstrated that my passive techniques are as reliable and accurate as active techniques in detecting CTPs deployed in cellular networks worldwide.

7. Survey of End-to-End Mobile Network Measurement Testbeds, Tools, and Services

Mobile (cellular) networks enable innovation, but can also stifle it and lead to user frustration when network performance falls below expectations. As mobile networks become the predominant method of Internet access, developer, research, network operator, and regulatory communities have taken an increased interest in measuring end-to-end mobile network performance to, among other goals, minimize negative impact on application responsiveness. In this work, I examined current approaches to end-to-end mobile network performance measurement, diagnosis, and application prototyping. I compared available tools and their shortcomings with respect to the needs of developers, researchers, network operators, and regulators. This work provides a comprehensive view of currently active efforts and some auspicious directions for future work in mobile network measurement and mobile application performance evaluation.

8. Improving content server selection via client-side assistance

Modern websites use Content Delivery Networks (CDNs) to speed up the delivery of static content. However, I demonstrated that DNS-based selection of CDN servers can be refined to fully deliver on the speedup of CDNs. I proposed DNS-Proxy (*dp*), a client-side process that shares load-balancing functionality with CDNs by choosing from among resolved CDN servers based on last mile network performance. My measurement study of CDN infrastructure deployed by five major CDN providers shows that *dp* reduces webpage load time by 29% on average. If *dp* has already resolved the domain, the reduction in webpage load time is as much as 40%. Finally, *dp* reduces the load time of individual static Web objects by as much as 43%. I argue that *dp* enables an effective use of existing content delivery infrastructure and represents a complementary strategy to a continual increase of geographic content availability.

9. Closest content server selection tool using traceroutes

Interactive mobile Internet applications have become increasingly common and necessary to perform everyday tasks such as augmented reality, online gaming, video chat, and cloud-based voice recognition. These applications speed up their communications by connecting users to the closest servers and clustering nearby users together. In replica server selection and client clustering, determination of the closest host quickly and accurately is crucial to interactive application responses and the satisfaction of users' expectations. Researchers commonly use latency as the primary metric of network proximity and have developed various latency approximation tools. However, these tools do not yet offer an attractive balance of measurement accuracy, scalability, and maintainability. In this work, I proposed a new latency estimation system for arbitrary hosts using host-to-CDN latency measurements. Compared to existing latency estimation tools, my technique offers superior coverage of the IP address space and latency estimation accuracy. With improved coverage and accuracy of latency estimation it will become easier to establish low latency connections between hosts in a network, improving the responsiveness of interactive Internet applications.

10. **A First Look at Web Browsing Predictions using DNS Logs**

Despite several decades of research towards improving Web experience, users remain dissatisfied. Content Providers (CPs) make contracts with Content Delivery Networks (CDNs) to speed up webpage load times. CPs also make contracts with Online Advertisement providers (OAPs) to improve relevancy of Web content shown on the website. However, both CDNs and OAPs remain unaware of what Web content users will request next. As a result, CDNs cannot pre-fetch data from servers hosted by CPs, which prevents reductions in webpage load times. Also, OAPs display advertisements based on stale and incomplete Web browsing histories. In this work, I introduced a new technique to analyze passively collected DNS logs for predicting users' Web browsing behaviors. In my experience with using the technique on a large-scale DNS dataset, I identified several sequences of websites that many users often tend to follow. Therefore, I made suggestions to CDNs and OAPs to predict users' Web browsing interests for improved quality of Web experience.

11. **Quality Assurance of a Mobile Network Measurement Testbed Through Systematic Software Testing**

The popularity of innovative mobile applications that offer services such as Web browsing, video streaming, online gaming, and collaborative communication puts utmost pressure on mobile application developers to ensure a high-quality user experience. As such, the research and development communities have developed several networking testbeds that measure the performance of application traffic in production cellular networks. In this work, I evaluated the quality of the source code of one of the mobile network measurement testbeds (MITATE) using multiple software testing techniques. My extensive testing experience with MITATE's source code indicates that network measurement testbeds are complex in their functionality and require multiple software components to interact with each other for any given operation. I demonstrated that using multiple testing techniques results in different types of issues with the code under test. Finally, based on the results, I argue that MITATE, in production, could offer high reliability and accuracy in executing network experiments.

12. **Mobile network measurement testbed for rapid application prototyping**

In this work, I designed and developed a mobile network measurement testbed, formally known as Mobile Internet Testbed for Application Traffic Experimentation (MITATE). MITATE is the first programmable testbed to support the prototyping of application communications between mobiles and cloud datacenters. MITATE incorporates novel solutions to device security and resource sharing behind MITATE and can help answer network performance questions crucial to mobile application design.

13. **Understanding Factors Influencing the Citation Count of Networking Conference Papers**

The impact of a journal on its community is usually judged based on its impact factor. Researchers often aim to publish their research findings at journals with high impact factors. Conversely, researchers tend to judge the quality of a conference based on its acceptance rate. However, it remains unknown if the low acceptance rate of a conference reflects a high impact on its community compared to a conference with high acceptance rate. I performed an empirical study to analyze the influence of conference rankings on the number of citations a paper may receive over time. My results show that papers published at highly reputed conferences tend to have a higher number of citations on average, than low ranked conferences. Additionally, I found that the title of the paper also influences the citation count.

14. **Investigating Human Brain Waves in Shared Experiences**

Artists desire to bring new experiences to viewers of their art. While modern technology, such as Neurosky Mindwave headset, allows artists to create such experiences, artists often strive to appropriately utilize such technologies. In other words, while the technology does exist, the techniques to utilize it does not. In this paper, I developed a technique that investigates similarity between the brain waves of two participants in a shared experience. Results indicate that there exist multiple brain waves that offer a high similarity when subject to a given external stimuli. Based on the observations I argue to artists to leverage the techniques and findings to better generate new experiences for all viewers of their art.

HONORS AND AWARDS

1. Outstanding Ph.D. Researcher Award by Montana State University for excellent scientific contributions made in the field of computer networking, April 2016.
2. Winner of the Three Minute Thesis Competition at Montana State University, March 2017.
3. Best Technical Solution for the Data Competition at Montana State University, October 2016.
4. Best Question of the Session Award by ACM MobiCom, October 2016.
5. Enhanced Graduate Teaching Assistant position by Montana State University, 2016.
6. Runner-up of the Three Minute Thesis Competition at Montana State University, March 2016.
7. Travel award from ACM to attend and present research at ACM MobiCom conference, October 2016.
8. Travel award from NSF to attend PhantomNet Workshop, 2015.
9. Travel award from IEEE to attend and present research at ICCCN conference, August 2015.
10. Travel award from ACM to attend CoNEXT conference, October 2014.
11. Travel award from ACM to attend IMC, August 2014.
12. Travel award from ACM to attend and present research at IMC, August 2013.
13. Travel award from USENIX to attend and present research at NSDI, March 2013.
14. Sir Syed Excellence in Science Award to pursue graduate studies in the USA, April 2012.
15. Merit Research Scholarship by Aligarh Muslim University for encouraging undergraduate research, 2011.

OTHER PRESENTATIONS AND INVITED TALKS

1. Changing Perception with Augmented Reality, at the Three Minute Thesis Competition, Montana State University 2017.
2. Towards a Faster Mobile Web, at Gianforte School of Computing Monday Seminar 2016.
3. Web Performance in Cellular Networks, at Akamai Summer Intern Showcase 2016.
4. Understanding Cellular Networks from Akamai's Perspective, at Akamai Summer Intern Showcase 2015.
5. Making the Internet Faster, at the Three Minute Thesis Competition, Montana State University 2016.
6. MITATE: Mobile Internet Testbed for Application Traffic Experimentation, at News & Community Feedback Session, ACM Internet Measurement Conference 2013.
7. MITATE: Mobile Internet Testbed for Application Traffic Experimentation, at Student Research Celebration, Montana State University, 2013.

SYNERGISTIC ACTIVITIES

- PlanetLab Consortium Member 2012 - Present
- Association for Computing Machinery Member 2016 - Present
- TPC member: IWQoS 2017; IWQoS 2016; IPCCC 2015; ICC 2015; SocialCom 2011.
- Graduate Representative/Mentor: ZHCET - North America 2013; Indian Student Association MSU 2012, 2013.

KEY TECHNICAL SKILLS

- Programming: Java, NodeJs, Python, PHP, R, SQL, Shell, Latex
- Tools: Wireshark, Web browser DevTools, Navigation and Resource Timing APIs, HTTP Archives
- Techniques: RUM, active and passive measurements, sharding, caching
- Services: AWS EC2, Planet-Lab, S3, Route 53, HTTP/DNS/TCP logs

RECENT COLLABORATORS

I have had the pleasure of working with several Internet experts. Here is a complete list of my recent collaborators from various industrial and academic institutions, for both published and unpublished literature.

Akamai Technologies: Dr. Moritz Steiner, Stephen Ludin, Erik Nygren, Martin Flack, and Wontaek Na.

Montana State University: Prof. Mike P. Wittie, Prof. Qing Yang, Prof. Brendan Mumey, Prof. Clemente Izurieta, Prof. Brittany T. Fasy, Prof. Upulee Kanewala, Samuel Micka, and James Espeland.

Other Institutions: Kimberly C. Claffy (CAIDA), Lara Deek (Netflix), Ajay Miyyapuram (Cerner), Kanika Shah (Cisco), Clint Cooper (SRI), Eben Howard (Bandi Namco), and Shantharaju Jayanna (Verizon!)