Nick Feamster

MIT Computer Science and Artificial Intelligence Laboratory
32-G982, The Stata Center
32 Vassar Street
Cambridge, MA 02139

Phone: (617) 253-7341
Fax: (617) 253-8460
feamster@csail.mit.edu
http://nms.csail.mit.edu/~feamster/

Research Interests

Networked computer systems: network protocols, routing, security, and management

Education

MASSACHUSETTS INSTITUTE OF TECHNOLOGY — *Cambridge, MA*

Ph.D. candidate in Computer Science. (Expected Summer 2005)
Dissertation: Robust and Predictable Internet Routing
Advisor: Hari Balakrishnan
Minor in Game Theory

M.Eng. in Computer Science, 2001
Thesis: Adaptive Delivery of Real-Time Streaming Video
Advisor: Hari Balakrishnan
William A. Martin Memorial Thesis Award (MIT M.Eng. thesis award)

S.B. in Electrical Engineering and Computer Science, 2000
Concentration in Economics

Professional Experience

2000– **Research Assistant** — MIT, *Cambridge, MA*

Research assistant at the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). Projects include work on interdomain routing robustness, circumventing Web censorship (Infranet), and the Congestion Manager project. More details are available on Page 5.

2001– **Technical Intern and Consultant** — AT&T Labs–Research, *Florham Park, NJ*

Research on interdomain traffic engineering and modeling. More details are available on Page 5.

1999 **Technical Associate** — Bell Laboratories, Lucent Technologies, *Murray Hill, NJ*

Designed and implemented a JavaBeans-based call filtering/disposition system which allows end users to easily design a call flow based on various criteria.

1998–2000 **Intern** — Hewlett-Packard Laboratories, *Palo Alto, CA*

Designed and implemented a transcoding algorithm for real-time conversion of MPEG-2 to H.263 bitstreams. More details are available on Page 6.

1997 **Technical Staff** — LookSmart Ltd., *San Francisco, CA*

Designed and implemented Web crawler, as well as monitoring and testing scripts for production search engine system.

Teaching Experience

2002 Teaching Assistant, MIT Course 6.829, Computer Networks.
Contributed new problems to problem sets and quizzes, gave two lectures, and taught recitations covering advanced topics.

2002–2003 M.Eng. research supervisor, MIT.
With Hari Balakrishnan, supervised Winston Wang, whose thesis on an implementation of the Infranet anti-censorship system received MIT's Charles and Jennifer Johnson Thesis Prize.

Refereed Publications

Note: Papers are listed in reverse chronological order by topic area.

Internet Routing

[1] Nick Feamster and Hari Balakrishnan. Detecting BGP Configuration Faults with Static Analysis. In *Proc. 2nd Symposium on Networked Systems Design and Implementation*, Boston, MA, May 2005. *Best paper award.*

[2] Matthew Caesar, Nick Feamster, Jennifer Rexford, Aman Shaikh, and Kobus van der Merwe. Design and Implementation of a Routing Control Platform. In *Proc. 2nd Symposium on Networked Systems Design and Implementation*, Boston, MA, May 2005.

[3] Nick Feamster, Hari Balakrishnan, and Jennifer Rexford. Some foundational problems in interdomain routing. In *Proc. 3nd ACM Workshop on Hot Topics in Networks (Hotnets-III)*, San Diego, CA, November 2004.

[4] Nick Feamster, Hari Balakrishnan, Jennifer Rexford, Aman Shaikh, and Kobus van der Merwe. The case for separating routing from routers. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Portland, OR, September 2004.

[5] Nick Feamster, Jared Winick, and Jennifer Rexford. A model of BGP routing for network engineering. In *Proc. ACM SIGMETRICS*, New York, NY, June 2004.

[6] Nick Feamster. Practical verification techniques for wide-area routing. In *Proc. 2nd ACM Workshop on Hot Topics in Networks (Hotnets-II)*, Cambridge, MA, November 2003.

[7] Nick Feamster, Jay Borkenhagen, and Jennifer Rexford. Guidelines for interdomain traffic engineering. *ACM Computer Communications Review*, 33(5), October 2003.

[8] Nick Feamster and Hari Balakrishnan. Towards a logic for wide-area Internet routing. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Karlsruhe, Germany, August 2003.

Internet Measurement

[9] Nick Feamster, Jaeyeon Jung, and Hari Balakrishnan. An empirical study of "bogon" route advertisements. *ACM Computer Communications Review*, November 2004.

[10] Nick Feamster, Zhuoqing Morley Mao, and Jennifer Rexford. BorderGuard: Detecting cold potatoes from peers. In *Proc. ACM SIGCOMM Internet Measurement Conference*, Taormina, Sicily, Italy, October 2004.

[11] Nick Feamster, David Andersen, Hari Balakrishnan, and M. Frans Kaashoek. Measuring the effects of Internet path faults on reactive routing. In *Proc. ACM SIGMETRICS*, San Diego, CA, June 2003.

[12] David G. Andersen, Nick Feamster, Steve Bauer, and Hari Balakrishnan. Topology inference from BGP routing dynamics. In *Proc. ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, November 2002.

Network Security

[13] Nick Feamster and Roger Dingledine. Location diversity in anonymity networks. In *ACM Workshop on Privacy in the Electronic Society*, Washington, DC, October 2004.

[14] Nick Feamster, Magdalena Balazinska, Winston Wang, Hari Balakrishnan, and David Karger. Thwarting Web censorship with untrusted messenger discovery. In *3rd Workshop on Privacy Enhancing Technologies*, Dresden, Germany, March 2003.

[15] Nick Feamster, Magdalena Balazinska, Greg Harfst, Hari Balakrishnan, and David Karger. Infranet: Circumventing Web censorship and surveillance. In *Proc. 11th USENIX Security Symposium*, San Francisco, CA, August 2002.
Best student paper award.

[16] Kevin Fu, Emil Sit, Kendra Smith, and Nick Feamster. Dos and don'ts of client authentication on the Web. In *Proc. 10th USENIX Security Symposium*, Washington, DC, August 2001. *Best student paper award.*

Adaptive Streaming Media Protocols

[17] Nick Feamster and Hari Balakrishnan. Packet loss recovery for streaming video. In *Proc. 12th International Packet Video Workshop (PV 2002)*, Pittsburgh, PA, April 2002.

[18] Nick Feamster, Deepak Bansal, and Hari Balakrishnan. On the interactions between congestion control and layered quality adaptation for streaming video. In *11th International Packet Video Workshop*, Kyongju, Korea, May 2001.

[19] Susie Wee, John Apostolopoulos, and Nick Feamster. Field-to-frame transcoding with temporal and spatial downsampling. In *IEEE International Conference on Image Processing*, October 1999.

[20] Nick Feamster and Susie Wee. An MPEG-2 to H.263 transcoder. In *SPIE Voice, Video, and Data Communications Conference*, Boston, MA, September 1999.

Submitted Publications and Works-in-progress

[21] Feng Wang, Nick Feamster, and Lixin Gao. Quantifying the causes of end-to-end Internet path failures, February 2005. Submitted for publication.

[22] Nick Feamster and Jennifer Rexford. Modeling BGP route selection within an AS. *IEEE/ACM Transactions on Networking*, July 2004. Earlier version appears in *ACM SIGMETRICS 2004*.

[23] Claire Monteleoni, Hari Balakrishnan, Nick Feamster, and Tommi Jaakkola. Managing the 802.11 energy/performance tradeoff with machine learning. Technical Report MIT-LCS-TR-971, Massachusetts Institute of Technology, 2004.

Unrefereed Papers and Technical Reports

[24] Nick Feamster, Jennifer Rexford, and Jay Borkenhagen. Controlling the impact of BGP policy changes on IP traffic. Technical Report 011106-02, AT&T Labs–Research, Florham Park, NJ, November 2001.

[25] Nick Feamster and Jennifer Rexford. Network-wide BGP route prediction for traffic engineering. In *Proc. SPIE ITCom*, Boston, MA, August 2002.

[26] Nick Feamster. Rethinking routing configuration: Beyond stimulus-response reasoning. In *Workshop on Internet Routing Evolution and Design (WIRED)*, Mt. Hood, OR, October 2003.

[27] Nick Feamster and Hari Balakrishnan. Verifying the correctness of wide-area Internet routing. Technical Report MIT-LCS-TR-948, Massachusetts Institute of Technology, May 2004.

[28] Russ White and Nick Feamster. *Considerations in Validating the Path in Routing Protocols*. IETF, April 2004. Internet Draft. Expires October 2004.

Invited talks at the North American Network Operators Group (NANOG), Cooperative Association for Internet Data Analysis (CAIDA), Boston University, Carnegie Mellon SDI seminar, New York University, Harvard University, University Catholique de Louvain (Belgium), AT&T Research, Hewlett-Packard Laboratories, and Agilent Technologies.

Awards and Honors

2005	Best Paper, 2nd Usenix Symposium on Networked Systems Design and Implementation
2004	Cisco URP Grant Recipient
2002–	NSF Graduate Research Fellow
2002	Best Student Paper, 11th Usenix Security Symposium
2001	Best Student Paper, 10th Usenix Security Symposium

2001	MIT William A. Martin Memorial Thesis Award
1999–	Tau Beta Pi Engineering Honor Society
1999–	Eta Kappa Nu Honor Society
1999	Letter of Commendation for Outstanding Performance, MIT Digital Design Laboratory
1998–1999	Phi Sigma Kappa Scholarship Award
1997	National Merit Scholar
1997	Rotary Club Scholarship
1996	AP Scholar with Distinction

Service and Other Activities

Reviewer for *IEEE/ACM Transactions on Networking*, *SIGCOMM* (2002, 2003, 2004), *SOSP* (2001, 2003), *Infocom* (2004), *HotNets* (2003), *HotOS* (2001), *USENIX Security Symposium* (2002), *ACM Computer Communication Review*, *IEEE Network Magazine*, *Image Communication* (EURASIP), *ASPLOS* (2004), *MobiSys* (2004), *USENIX* (2005), *NSDI* (2005), *IPTPS* (2005).

References

Prof. Hari Balakrishnan
MIT Computer Science & AI Lab
32 Vassar Street, 32G-940
Cambridge, MA 02139
(617) 253-8713
hari@csail.mit.edu

Prof. M. Frans Kaashoek
MIT Computer Science & AI Lab
32 Vassar Street, 32G-992
Cambridge, MA 02139
(617) 253-7149
kaashoek@csail.mit.edu

Prof. Jennifer Rexford
Princeton University
Department of Computer Science
35 Olden Street, CS 306
Princeton, NJ 08544
(609) 258-5182
jrex@cs.princeton.edu

Prof. Ramesh Johari
Stanford University
Department of Management Science and Engineering
380 Panama Mall
Stanford, CA 94305
(650) 723-0937
ramesh.johari@stanford.edu

Prof. Lixin Gao
Department of Electrical and Computer Engineering
Knowles Engineering Building
University of Massachusetts
Amherst, MA 01002
(413) 545-4548
lgao@ecs.umass.edu

Internet Routing

The Internet is composed of more than 17,000 independently operated networks, or autonomous systems (ASes), that exchange routing information using the Border Gateway Protocol (BGP). Network operators in each AS configure routers to control the routes that the routers learn, select, and propagate. Configuring a network of BGP routers is like writing a distributed program where complex feature interactions occur both within one router and across multiple routers. This complex process is exacerbated by the number of lines of code, by the absence of useful high-level primitives in today's router configuration languages, by the diversity in vendor-specific configuration languages, and by the number of ways in which similar high-level functionality can be expressed in a configuration language. As a result, router configurations tend to have faults. Faults in BGP configuration can cause forwarding loops, packet loss, and unintended paths between hosts. Operators must be able to evaluate the effects of a configuration and be assured that the configuration is correct before deploying it. My dissertation advances the state of the art in Internet routing by devising fault detection and modeling tools for today's Internet routing protocols and proposing a new Internet routing architecture that alleviates many of the problems we uncovered in our work on fault detection and modeling.

Detecting Faults in BGP Configuration with Static Analysis MIT

rcc, the *router configuration checker*, detects faults in the BGP configurations of routers in an AS using static analysis. **rcc** detects two broad classes of faults that affect network reachability: route validity faults, where routers may learn routes that do not correspond to usable paths, and path visibility faults, where routers may fail to learn routes for paths that exist in the network. **rcc** enables network operators to test and debug configurations before deploying them in an operational network, improving on the status quo where most faults are detected only during operation. **rcc** has been downloaded by more than sixty network operators to date. I presented **rcc** to the North American Network Operators Group (NANOG), and the tool has been used by several large backbone Internet Service Providers (ISPs) to successfully detect faults in deployed configurations. This work was inspired by my work on the *routing logic* that I presented at the 2003 *ACM SIGCOMM Workshop on Future Directions in Network Architecture* and appears at the *2nd USENIX Symposium on Networked Systems Design and Implementation*. We have also studied configuration faults as part of several measurement studies. We presented an algorithm to detect route advertisements that violate peering contracts and an empirical study of their prevalence at the 2004 *ACM Internet Measurement Conference*.

Modeling Internet Routing for Network Engineering MIT/AT&T Labs–Research

Since interdomain route selection is distributed, indirectly controlled by configurable policies, and influenced by complex interactions with *intra*domain routing protocols, operators cannot predict how a particular BGP configuration would behave in practice. We devised an algorithm that computes the outcome of the BGP route selection process for each router in a *single* AS, given only a static snapshot of the network state, without simulating BGP's complex dynamics. Using data from a large ISP, I demonstrated that the algorithm correctly computes BGP routing decisions and has a running time that is efficient and accurate enough for many tasks, such as traffic engineering and capacity planning. Studying the general properties and computational overhead of modeling the route selection process in each of these cases provides insight into the unnecessary complexity introduced by various aspects of today's interdomain routing architecture. I used these insights to propose improvements to BGP that avert the negative side effects of various artifacts without limiting functionality. This work appeared in *ACM SIGMETRICS* 2004 and has also been submitted to *IEEE/ACM Transactions on Networking*.

Internet Routing Architecture: Routing Control Platform MIT/AT&T Labs–Research

The limitations in today's routing system arise in large part from the fully distributed path-selection computation that the IP routers in an AS must perform. We proposed that interdomain routing should be separated from today's IP routers, which should simply forward packets (for the most part). Instead, a separate *Routing Control Platform (RCP)* should select routes on behalf of the IP routers in each AS and exchange reachability information with other domains. RCP could both select routes for each router in a domain (*e.g.*, an AS) and exchange routing information with RCPs in other domains. By selecting routes on behalf of *all* routers in a domain, RCP can avoid many internal BGP-related complications that plague today's mechanisms for disseminating and computing routes within an AS. RCP facilitates traffic engineering, simpler and less error-prone policy expression, more powerful diagnosis and troubleshooting, more rapid deployment of protocol modifications and features, enforceable consistency of routes, and verifiable correctness properties. The architectural proposal for RCP appeared at the 2004 *ACM SIGCOMM Workshop on Future Directions in Network Architecture*; the design and implementation of an RCP prototype won the best paper award at the *2nd USENIX Symposium on Networked Systems Design and Implementation (NSDI)*.

Internet Measurement

Understanding End-to-End Internet Path Failures MIT

Empirical evidence suggests that reactive routing systems, which detect and route around faulty paths based on measurements of path performance, improve resilience to Internet path failures. We studied *why* and under *what circumstances* these techniques are effective by correlating end-to-end active probes, loss-triggered traceroutes of Internet paths, and BGP routing messages. This work was the first known study to correlate routing instability with degradations in *end-to-end* reachability. We found that most path failures last less than fifteen minutes. Failures that appear in the network core correlate better with BGP instability than failures that appear close to end hosts. Surprisingly, there is often increased BGP traffic both before and after failures. Our findings suggest that reactive routing is most effective between hosts that have multiple connections to the Internet and that reactive routing systems could pre-emptively mask about 20% of impending failures by using BGP routing messages to predict these failures before they occur. This work appeared at *ACM SIGMETRICS* 2003.

End-to-end path failures are typically attributed to either congestion or routing dynamics. Unfortunately, the extent to which congestion and routing dynamics cause end-to-end failures, and the effect of routing dynamics on end-to-end performance, are poorly understood. In a follow-up study, we used similar techniques to find that routing dynamics contribute significantly to end-to-end failures and, in particular, routing dynamics are responsible for most long-lasting path failures. The study also finds that long-lived end-to-end path failures that involve routing dynamics are typically caused by BGP convergence or instability. This work is the first to quantify the impact of routing dynamics on end-to-end path availability; it was submitted to *ACM SIGMETRICS* 2005.

Network Security

Infranet: Circumventing Web Censorship MIT

An increasing number of countries and companies routinely block or monitor access to parts of the Internet. To counteract these measures, we designed and implemented *Infranet*, a system that enables clients to surreptitiously retrieve sensitive content via cooperating Web servers distributed across the global Internet. These Infranet servers provide clients access to censored sites while continuing to host normal uncensored content. Infranet uses a tunnel protocol that provides a covert communication channel between its clients and servers, modulated over standard HTTP transactions that resemble innocuous Web browsing. In the upstream direction, Infranet clients send covert messages to Infranet servers by associating meaning to the *sequence* of HTTP requests being made. In the downstream direction, Infranet servers return content by hiding censored data in uncensored images using steganographic techniques. This work appeared at the *11th USENIX Security Symposium*.

Adaptive Streaming Media Protocols

Reliable, Adaptive Video Streaming MIT

Video compression exploits redundancy between frames to achieve higher compression, but packet loss can be detrimental to compressed video with interdependent frames because errors potentially propagate across many frames. In my Master's thesis, I quantified the effects of packet loss on the quality of MPEG-4 video, developed an analytical model to explain these effects, and presented an RTP-compatible protocol, called *SR-RTP*, that *adaptively* delivers higher quality video in the face of packet loss. This work appeared at the *12th International Packet Video Workshop* and was later implemented as part of a streaming video server for MIT Project Oxygen. We also designed a scheme for performing quality adaptation of layered video for a general family of congestion control algorithms called *binomial congestion control*. This work appeared at the *11th International Packet Video Workshop*.

Video Transcoding Hewlett-Packard Laboratories

We designed and implemented an algorithm that transcoded MPEG video input to a lower-bitrate H.263 progressive bitstream, facilitating the transmission of a digital television signal over a wireless medium. This algorithm was the first to use both spatial and temporal downsampling in an MPEG-2 to H.263 field to frame transcoder to achieve substantial bitrate reduction. The proposed algorithm exploits the properties of the MPEG-2 and H.263 compression standards to perform interlaced to progressive (field to frame) conversion with spatial downsampling and frame-rate reduction in a CPU and memory efficient manner, while minimizing picture quality degradation. This work appeared at the *IEEE International Conference on Image Processing* in 1999.

Research Statement

Nick Feamster

I am interested in designing, building, analyzing, and measuring networked systems that are composed of multiple autonomous, potentially untrusted entities. There are numerous examples of such systems: the Internet routing infrastructure, public wireless networks, peer-to-peer systems, large-scale distributed computing infrastructures (*e.g.*, grid computing), Web services, content delivery networks, and enterprise networks.

These systems present three challenges to system designers:

1. Developing communication protocols that ensure correct, predictable, and robust operation;
2. Designing practically deployable techniques for secure operation;
3. Supporting fault diagnosis, troubleshooting, and monitoring.

Some fields of engineering have reached a level of maturity where there are concrete design principles that ensure a level of correctness and robustness: for example, it is reasonably well understood how to build bridges and buildings that withstand earthquakes and high winds. Engineers can apply such principles to help them design robust, reliable systems. Unfortunately, network system designers and engineers lack such a rubric.

My research uses four techniques towards the ultimate goal of developing sound methods for designing and implementing networked systems: measurement; modeling; design and implementation; and deployment. Measurement provides evidence and intuition for the severity of a problem. Modeling—the process of deriving a simpler representation that abstracts irrelevant details and concisely describes the aspects that affect the properties under study—facilitates a more thorough understanding of a problem's fundamental causes. Design and implementation provide the opportunity to apply the intuition gained from measurement and modeling to make tangible improvements to real-world systems. Deployment demonstrates the feasibility and practicality of an implementation and provides the excitement of seeing research ideas applied in practice. In my future work, I intend to apply these techniques to specific problems in routing, network security, anonymous communication, anomaly detection, and monitoring in resource-limited environments.

Dissertation Work: Robust, Predictable Internet Routing

My dissertation solves some of the challenges raised above in the context of Internet routing, which requires that competing, autonomous networks ("autonomous systems", or ASes) cooperate to establish global connectivity. I have designed and implemented models and tools that make today's routing infrastructure more robust and easier to manage. I have examined fundamental tradeoffs between routing stability and expressiveness in generic policy-based routing protocols. Finally, I have proposed a new Internet routing architecture that solves many of the problems we discovered (*i.e.*, through measurement and analysis) with today's routing infrastructure.

Routing configuration is essentially a complex distributed program. Each AS independently configures local *policies* that control how routers select and re-advertise routes. These policies implicitly codify bilateral business relationships between ASes. Each AS may contain tens to

hundreds of routers, each of which is individually configured with hundreds to thousands of lines of code. The collection of configurations within an AS determines whether the routing protocol operates correctly. Faults in configuration can induce routing failures, such as forwarding loops, partitions, and instability, that can prevent packets from reaching their destinations.

My research has applied measurement, modeling, design and implementation, and deployment to help make today's Internet routing infrastructure less prone to failure, as well as more predictable. Network operators need assurances that today's routing protocols will operate correctly, and they need to know which route each router will select, given a set of configurations. My work on a *routing logic* defines a correctness specification for policy-based routing. Based on this specification, I developed *rcc* ("router configuration checker"), a tool that analyzes the set of router configurations from a single AS and detects configuration faults that could induce routing failures. *rcc* has been used by over sixty network operators and has successfully identified faults in the configurations of several Internet service providers with nationwide backbone networks. Experience with *rcc*'s deployment in real-world networks has provided a better understanding of the nature and extent of configuration faults that occur in practice. Additionally, my work on modeling route selection led to a tool that makes routing more predictable by helping network operators predict the effects of a configuration change before deploying it.

With collaborators, I have also applied the above techniques to design improvements to today's Internet routing protocols. Ideally, Internet routing should disseminate loop-free routes and converge to a stable routing topology, regardless of how each AS configures its local policies. We have applied an abstract model of today's Internet routing protocol to derive constraints on local policies that must be satisfied to guarantee that a policy-based routing protocol will not oscillate. To guarantee correct dissemination of loop-free routes, we proposed the Routing Control Platform (RCP), a system that computes routes on behalf of routers. By applying two design principles—(1) compute consistent routes with complete routing information and (2) control interactions between different routing protocols (*e.g.*, between the inter-AS routing protocol and an AS's internal routing protocol)—RCP explicitly prevents the forwarding loops and oscillations that plague today's Internet routing infrastructure.

Future Research Directions

I intend to continue working on improving the robustness, security, and diagnosis capabilities of large-scale systems in which potentially untrusted entities must cooperate to provide some service.

Robustness and Predictability

Wireless mesh networks. While Internet routing is perhaps the best studied example of a routing system that requires cooperation among multiple untrusted parties, other domains, such as public wireless "mesh" networks, present interesting issues. These networks are composed of nodes that are typically owned by different parties (*e.g.*, homes, businesses) that must cooperate to provide connectivity. Because each of these entities may have vastly different criteria for ranking preferred paths through the network and for carrying traffic over those paths (*e.g.*, minimizing loss rate, cost, etc.), these wireless networks may benefit from using policy-based routing protocols. I intend to explore routing problems in public wireless networks to see what design principles can tackle wireless-specific challenges (*e.g.*, contention for a shared channel, interference, mobility) to achieve robust and predictable routing.

Routing stability. Policy-based routing protocols provide each participant remarkable flexibility for implementing complex business arrangements in local policy; unfortunately, the interactions between these policies may conflict, resulting in instability. The tradeoff between routing policy flexibility and stability is poorly understood today. I would like to characterize the minimal set of constraints that must be imposed on each participant's policies to guarantee global stability. In the context of Internet routing, I would like to determine whether those constraints are expressive enough to implement important operational tasks (*e.g.*, load balancing traffic). My experience using game theory and mechanism design to study routing protocol oscillation, as well as my knowledge of network operations, should prove useful for solving these problems.

Secure Networked Systems

Data plane security. Previous work has studied routing protocol security, but little attention has been paid to security and policy enforcement in the *data plane* (*i.e.*, the path that data packets actually traverse). Today's Internet architecture provides scant support for a network to thwart unwanted packets (*e.g.*, spam, viruses, denial of service attacks) and essentially no control over the sequence of ASes that outgoing traffic traverses en route to a destination. I plan to design architectural modifications that could facilitate stronger security and policy enforcement capabilities in the data plane.

Anonymous and censorship-resistant communication. Governments of certain countries routinely implement firewalls to restrict communication to various destinations. To enable clients behind these firewalls to access restricted Web content, we designed and implemented Infranet, an anti-censorship system that embeds requests for blocked content in a covert channel that appears to the censor as innocuous traffic to permissible Web sites. I would like to design anti-censorship systems like Infranet that are robust to untrusted or malicious participants. More generally, I intend to examine how incorporating network-layer information can make anonymous communication systems more resistant to eavesdropping attacks.

Fault Diagnosis and Monitoring

Anomaly detection. Supporting both fault diagnosis and secure operation in large-scale networked systems typically requires the ability to collect, analyze, and audit large quantities of data. I intend to explore whether signal processing and clustering techniques can be useful for performing forensic analysis of spam (*e.g.*, determining groups of machines that are being controlled by a single sender). I have also begun investigating whether signal processing-based anomaly detection techniques such as principal component analysis can be useful for detecting routing anomalies.

Monitoring in resource-limited environments. I previously applied signal processing techniques to design and implement a video transcoder that allowed video streams that were originally encoded at very high bitrates to be transmitted over relatively low-bandwidth wireless links in real-time. I plan to investigate how signal processing can reduce communication costs in other bandwidth and resource-constrained environments. For example, sensor networks have strict bandwidth and energy budgets that often require data streams with high data rates to be processed in the network. I would like to design and implement distributed signal processing algorithms that help reduce computation and communication overhead in resource-constrained environments.

Teaching Statement

Nick Feamster

I look forward to the responsibility and privilege of teaching and mentoring students. I believe that a teacher has two primary responsibilities:

1. Exciting students about problems and helping them discover their interests;
2. Providing students with the necessary resources to succeed in pursuing those interests.

A good teacher uses a course not only to impart knowledge, but also to instill excitement about the material and help students discover new interests. As a teaching assistant for MIT's graduate networking course, I prepared lectures, recitations, and new questions for problem sets and quizzes. Lectures and recitations provide a unique opportunity to present new problems to students and challenge them to think in new ways. A good lecture or recitation should engage students by relating topics that they are not familiar with to those that they know well or some topic that they can be excited about (*e.g.*, motivating video transcoding with a wireless digital television application, rather than speaking exclusively in the abstract). It should also challenge students to think on the fly. Finally, I believe that a lecture should leave students with a clear intuition for the fundamental high-level problems and intellectual ideas in some area (which they could hopefully carry with them for years) but still allow a student who is interested in delving into the topic further with enough information to pursue his or her interests.

A good advisor helps students discover their interests by recognizing their specific strengths and guiding them towards interesting open problems that capitalize on those strengths. I have found that an effective way to excite students about a problem is to pique their interest with a concrete problem, a couple of (sometimes small) interesting results, and a handful of open questions to think about. To this end, I spend time thinking about interesting questions (*e.g.*, Do spammers steal addresses from other Internet service providers to untraceably send spam?) and performing some initial exploration of these problems so that I can pose problems to students in terms of concrete examples, rather than simply describing a problem in the abstract. I have found this approach to be successful in helping guide the research of Winston Wang, a Master's student who addressed some open problems I posed in the Infranet project and ultimately received an award for his thesis on the topic; and, more recently, of Mythili Vutukuru, a first-year Ph.D. student who is working on open issues related to the Routing Control Platform, based on the design that we previously proposed.

An advisor should also develop his or her students' research tastes, encouraging creativity while ensuring that students do not waste time on irrelevant or uninteresting problems (*e.g.*, problems that are too short-sighted, emphasize implementation without a research goal, etc.). One of the most valuable things that my advisor, Hari Balakrishnan, did for my research was to help steer it clear of these types of problems in a constructive way that encouraged me to refine my ideas; I intend to do the same for my students.

Teaching also requires providing students with the necessary resources to succeed once they have discovered their passions. Perhaps the most important resource an advisor can provide is the ability to draw connections between different research areas that do not initially appear to be

related. For example, my advisor recognized that information theory could help me design the covert channels in the Infranet project. This type of insight requires familiarity with many fields; I believe that my background in other fields, including game theory and signal processing, will help me provide the same support to my students.

One of the most important aspects of teaching undergraduates is conveying excitement about the material. I believe that one good way to do this is to pose a problem in terms of a concrete example or application that provides solid intuition. My discrete math professor explained merge sort with a Tower of Hanoi-style set of rings and combinatorics with card tricks. I believe that using concrete examples—whether they involve using cards to explain combinatorics or video streams to explain the effects of packet loss on streaming video quality—not only excites students about material, but also provides them with intuition that they will remember long after their memories of details have faded. Coursework should involve problems and projects that encourage students to both "get their hands dirty" and to develop creative problem solving skills based on their newfound knowledge.

Given my experience in Internet routing, Internet measurement and modeling, and network security, I would also be excited to hold graduate seminars on any of these topics. Finally, I am very interested in designing and teaching courses at both the undergraduate and graduate levels that could be applied to interdisciplinary research that I intend to pursue (*e.g.*, applying game theory and signal processing to networks).