Statement of Research Interests & Methods

An advisor once described me as "being interested in any interesting problem." I look forward with great anticipation to establishing a research group with students working on diverse problems in networking and computer systems that we mutually find intriguing. Supervising an M. Eng. student at MIT recently, I learned that working closely with students can be a wonderful source of creativity and insight. I hope to foster an environment where I can collaborate with students on related, but distinct projects that personally excite each of them. I enjoy tackling several projects at once. Inspiration strikes on its own schedule, and I often find it rewarding to pursue several promising ideas in parallel.

At MIT, I've worked with a number of different faculty members on research projects in a broad range of areas, including mobile networking, wireless network performance, robust content distribution, fault-tolerant Internet services, and distributed application management. In addition, I have worked part-time at BBN Technologies for the past three years, where I have focused on active networking and denial of service detection. While the focus of my projects has been varied, there are several unifying themes in my research. Chief among them is an emphasis on networking, particularly in the mobile and wide-area environments. Networks have become the life-blood of computer systems and provide a virtually inexhaustible source of new topics for exploration.

My projects endeavor to develop new functionality. Achieving good performance is critical, but the art of squeezing out the last few cycles is best used in moderation. I strive to develop systems that extend the state-of-the-art, and qualitatively, rather than only quantitatively, improve the fields of networking and computer systems. I derive great satisfaction from doing something today that was impossible yesterday.

There is also commonality in the way I approach problems. While all researchers enjoy charting new territory, it is often fruitful to re-examine classical problems with known (limited) solutions. I enjoy considering alternative approaches. Where the classical approach remains superior, examining unexplored alternatives may yield a deeper understanding of the limits or applicability of the current approach. However, if one seeks to identify domains in which circumstances have changed, or new technologies have been developed, one may find techniques previously discarded as foolish or infeasible are relevant once again. This is precisely the case in two of my recent research projects.

My dissertation research proposes, implements, and evaluates the Migrate architecture for Internet host mobility. A fundamental problem in mobile networking is preserving connectivity when communicating hosts change network locations and efficiently resuming communication after periods of disconnection. Current approaches to host mobility, such as Mobile IP and ad-hoc routing, assume Internet hosts are either rarely mobile or constantly mobile and always on and reachable. Neither model is appropriate for the intermittently connected laptops and PDAs that populate today's Internet, and both approaches lack the application support necessary for the cell phones and personal communication devices that will comprise tomorrow's. Migrate addresses these shortcomings by handling intermittent connectivity and a variety of mobility modes.

One interesting technical component of Migrate is a novel TCP connection migration scheme. It allows a TCP connection to be resumed from a different IP address, solving a problem that Vint Cerf described in a 1983 paper as "[having] plagued network designers since the design of the ARPANET in 1968." A TCP connection is identified by its end points; hence, a change in end point (IP address) invalidates the connection. TCP migration schemes have been proposed previously but failed in various ways to preserve the semantics, performance, or security of the connection. The Migrate option avoids such ill-effects by utilizing multiple, standard TCP connections. It negotiates a cryptographically secure identifier for each connection which can then be used to continue over an entirely new connection at another end point. The key observation is that only the connection end points need to be aware of the previous connection; all other entities in the network (routers, NATs, firewalls, etc.) should view the continuation as another, independent connection. So long as the end points ensure they can stitch together a TCP connection and its continuation, the resulting connection must perform identically to a standard TCP connection, as the network cannot tell the difference.

At BBN, I helped design the Source Path Isolation Engine (SPIE), which was developed to help combat denial of service and other IP-based attacks. SPIE identifies the true source of an IP packet and the path the packet followed while in the network, allowing victims to reliably identify their attacker(s). Paths are constructed by maintaining records of every packet forwarded at each router, and simply querying each router backwards along the path beginning at the victim. This general technique, known in the literature as packet logging, was previously regarded as intractable in the wide-area due to the massive packet histories required. The centerpiece of SPIE is a hash-based packet digesting procedure that reduces the storage requirement of an IP packet to only a few bits, bringing packet logging into the realm of practicality.

My experience with both the cryptographic aspects of Migrate and the hashing and data storage requirements of SPIE has reinforced my belief that systems research is often best conducted in a collaborative, cross-disciplinary fashion. It is critical that systems researchers spend time with experts in more formal disciplines, such as mathematicians and theoreticians. Recent systems projects in a wide range of areas, from distributed storage systems to overlay networks to event notification systems, have shown theoretical computer science can provide a powerful toolbox of techniques for those with sufficient understanding, desire, and patience to search for creative applications.

Do not mistake my predilection for clever theoretical and algorithmic approaches as an aversion to detailed implementation. I firmly believe that high-quality computer systems research requires building, deploying, and evaluating real systems. I have built every system I've proposed as part of my graduate research, and each implementation experience has presented additional complications, subtleties, and sometimes even principles that would otherwise have gone unnoticed. Not only is a full-scale implementation absolutely essential to fully understanding the implications of a proposed system, but it often uncovers fertile areas for future research as well—perhaps unrelated to the project at hand. In either case, system building keeps faculty in touch with today's trends, and produces students who are intimately familiar with the details of real systems, knowledge and skills they will find extremely valuable in both academic research and industry.

Moving forward, I plan to continue to focus on providing new functionality in computer systems. In particular, I believe the ever-increasing permeation of light-weight, mobile computing devices will challenge classical notions of computer systems and provide a fertile area for research. Just as recent interest in peer-to-peer systems has rejuvenated work in distributed file systems, I expect the significant computational power now available in a hand-held form factor will lead to a renaissance of distributed computing. I look forward to designing systems agile enough to function on such a substrate and networks capable of securely and robustly communicating between them. In the near term, I intend to explore topics exposed by my dissertation and recent work in robust content distribution.

The Migrate architecture proposed in my dissertation addresses host mobility, but stops short of handling user or service mobility, where an application or communication end point moves not only across network location, but across computational location (e.g. from laptop to cell phone) as well. As cell phones, PDAs, and other devices approach the computational power available in laptops, the ability to migrate sessions across devices becomes increasingly practical and desirable. This form of migration is complicated by the heterogeneity of system resources and methods of user interaction, but may yield to a continuation-based approach similar that advocated by Migrate.

In another thrust, I hope to extend my recent work on content distribution. Mesh-based, multi-path overlay routing has shown promise in delivering greater levels of reliability and performance than traditional distribution networks, but little is currently known about how to effectively construct distribution topologies in the wide area, or how their configuration affects overall network performance. By making a general-purpose, standards-based content distribution utility available to fellow researchers, I hope to build a distribution network large enough to study such effects.