Research Statement

Sachin Katti

My research interests are in the area of networked systems, with an emphasis on wireless systems. Cheap, fast and portable computing devices with ubiquitous wireless connectivity have the potential to revolutionize the personal computing landscape, creating an opportunity to design an unprecedented array of new applications and services. Standing in the way of this vision however, are several challenges across the entire spectrum of the system architecture. These range from the unpredictable wireless medium and interference at the lower layers, to limited throughput, unreliability and mobility at the network layer, and to complex, unsecure and closed software substrates at the application layer. Typical approaches to these problems have been piecemeal and isolated, leading to partial solutions at different layers which when put together introduce more complexity, inhibit extensibility, and often limit performance. My approach looks at the entire system architecture in a holistic fashion. I develop interdisciplinary solutions that encompass signal processing schemes, coding and decoding techniques, network protocols and software substrates.

In particular, I have designed a *network coded wireless architecture*, where I demonstrate through analyses, implementations and evaluations, how, if we use network coding, different layers can co-operate to exploit the physical characteristics of the wireless medium and significantly improve performance. The systems in this architecture provided the first implementations of network coding and integrated it into the network stack, showing that network coding is a viable practical technique. Examples include COPE, a packet-level network coding technique to take advantage of wireless broadcast, ANC, an analog network coding technique to take advantage of strategic interference, and MIXIT, a symbol-level network coding technique to take advantage of wireless spatial diversity. I have also built software systems which address security and privacy problems at the application layer; such as an overlay anonymity system that does not require a PKI to protect user identities, and a low-level security hypervisor to track the flow of private data and prevent it from falling into malicious hands.

A common thread in my research is the harnessing of theoretical techniques from varied disciplines to solve practical problems. My approach is to identify a systems problem, reduce it to its core by stripping away unnecessary details, and look for clean conceptual solutions. Such an approach often clarifies the problem and shows possible connections to other areas from where we can borrow ideas and develop innovative solutions. For example, in the COPE system, I observed that wireless broadcast creates shared redundancy, which I exploited by designing a network coding compression technique to improve throughput. In my other systems, I have utilized ideas from information theory to design anonymous routing protocols, control and optimization theory to design wireless routing protocols, and hardware virtualization techniques to build secure systems.

Real world systems, however, are complex, and cannot usually be captured by clean conceptual models. Hence, answering whether a technique really works requires careful system design, experimentation, and rigorous experimental evaluation. In my research, apart from theoretical analyses, I implement and deploy my ideas in working systems. By analyzing and evaluating ideas in practical settings, we gain a fundamental understanding of why certain techniques work and others do not. Such evaluations validate theory, and often suggest fresh avenues for theoretical investigation. For example, the experimental evaluation of COPE suggested a subtle connection between network coding and MAC scheduling, which was subsequently investigated by other researchers, while MIXIT's evaluation shed fresh light on the connection between wireless spatial diversity and end-to-end congestion control protocols.

My approach naturally leads to collaborations with researchers across a wide spectrum. I have worked with information and coding theorists, network architects, security experts and OS engineers. These diverse collaborations have exposed me to a wide variety of approaches and philosophies, and provided me with a unique perspective on systems research. Networked systems, especially the wireless mobile Internet, encompass all these disciplines. However, due to the original layering contract, these disciplines have developed separately, and over time this has created a language barrier between them. But as networking becomes a more complex field, problems become harder to solve in isolation. My research has strived to bridge this communication gap, allowing me to approach basic networking problems from an alternative perspective.

RESEARCH CONTRIBUTIONS

My inter-disciplinary approach has produced several pieces of research. Among them, I believe the most important are my contributions to wireless mesh networks and network coding, and to designing systems which preserve user privacy and security. Below I describe each of these in more detail.

Network Coded Wireless Architecture

My thesis research proposes a network coded wireless architecture to improve the performance and reliability of wireless mesh networks. I advocate a simple yet fundamental shift: instead of routers storing and forwarding packets, they are allowed to "mix" the contents of packets before forwarding them, i.e., perform network coding. I show through practical systems how routers can exploit this new functionality to harness the intrinsic characteristics of the wireless medium to improve performance. My research thus builds a strong connection between the theory of network coding and wireless system design. Prior work on network coding is mainly theoretical and focused on multicast traffic. The systems in my dissertation bridge network coding theory with practice, by addressing the common case of bursty and dynamic unicast traffic. They were also the first to show that network coding can be cleanly integrated into the network stack to deliver practical and measurable gains. Further, the work developed novel algorithms that enriched the theory of network coding, extending it to operate over multiple unicast flows, analog signals, and soft-information. Below I describe the three main components of this architecture.

COPE (SIGCOMM 2006) is a packet-level network coding technique that exploits wireless broadcast to improve throughput in congested networks. The key observation is that due to wireless broadcast, multiple nodes overhear every transmitted packet, creating a pool of shared redundancy around each router. We exploit this redundancy by designing a new network coding technique to compress multiple independent packets into a single transmission, and thus increase throughput. COPE was implemented as a coding shim which cleanly integrates between the network and link layers, and works with existing applications. The experimental evaluation of COPE in a 25 node deployment demonstrated the practical benefits, and also revealed a surprising connection between MAC scheduling and coding, which was subsequently investigated by other researchers. COPE also provided a fresh angle to attack the open problem of network coding in the presence of multiple unicast sessions, and has inspired significant follow-on work in the network coding community.

My second system, **ANC** (SIGCOMM 2007) sits in the physical layer. Traditionally interference has been considered harmful. ANC adopts the opposite approach, it encourages strategically picked senders to interfere and shows how we can exploit such strategic interference to increase concurrency and improve throughput. The key observation behind ANC is that in many cases a node has side information about one of the signals in an interfered signal, in the form of a prior transmitted/overheard packet. ANC exploits this observation by designing an analog network coding technique to disentangle interfered signals. Practical evaluations in a testbed of software radios show that ANC delivers large gains in practice, but the importance of ANC goes beyond the throughput improvements it provided. ANC advocated a shift in how interference was handled, exploiting and decoding interference instead of trying to avoid it using conservative techniques that limit performance. This approach was subsequently adopted by significant follow-on work that dealt with collisions in access point networks.

MIXIT (SIGCOMM 2008), the third component of the network coded architecture, exploits symbol-level wireless spatial diversity to provide large throughput and reliability increases. The key insight behind MIXIT is that due to spatial diversity in a dense wireless mesh network, even if no node receives a packet correctly, every bit is likely to be received correctly by some node. MIXIT leverages this observation by building a symbol-level network where nodes funnel their correctly received bits to the destination such that the entire packet is delivered reliably. Designing and implementing such a network is made possible by MIXIT's symbol-level network coding technique, which takes away the otherwise onerous co-ordination and scheduling overhead one would have incurred in building a symbol-level network. MIXIT also makes connections with theoretical work on co-operative diversity and backpressure routing for congestion control, demonstrating how a network coding substrate allows us to exploit these theoretical ideas in practice.

My research makes a clear departure from conventional network design. Decades ago, electrical engineers and computer scientists agreed on a digital contract, with the electrical engineers focusing on the physical and lower layers, while the computer scientists worked up from the network layer, with the packet being the only interface. This approach has worked well for the wired Internet, but the wireless medium is fundamentally different, and demands an inter-disciplinary approach. My research builds such a new contract, disposing of artificial abstractions such as indivisible packets and point-to-point links in favor of more natural abstractions that allow the network and lower layers to collaborate on the common objectives of improving throughput and reliability, using network coding as the building block. At the same time, the design maintains desirable properties such as being distributed, modular, implementable, and integrable with the rest of the network stack.

Secure Systems

My second main thread of research lies in improving the security and privacy of networked systems. I have primarily focused on protecting user privacy in two aspects: their identities and their private data. Anonymous routing is a useful service, since it protects user identity in an increasingly monitored Internet. However, adoption of the technique has been slow due to the requirement of a centralized public key infrastructure (PKI). I developed a technique **Information Slicing** (NSDI 2007), that gets rid of this requirement. The key insight is that the Internet provides a large number of node disjoint paths to any destination, thus if one were to code and divide information across these independent paths, no node except the destination can recover the original data. Information Slicing uses this basic idea to build an onion routing scheme without keys, where each node only knows its previous and next hops. Thus the identity of the source and destination are kept secret, and even the destination does not know the identity of the source. The technique borrowed ideas from coding theory (random linear codes to divide information) and information theory (to analyze secrecy) to solve a practical problem. I also implemented and evaluated the technique over the Planetlab overlay network. I am currently working on techniques to protect private user data from falling into malicious hands, which I will describe in detail in the next section.

CURRENT & FUTURE RESEARCH

My research goal is to design and build the future wireless mobile Internet. Challenges abound in every layer of this architecture, and have cascading effects on each other. I highlight what I believe are the salient ones below.

The Radio: At the lowest layer, the fundamental question is designing a radio for the next generation wireless Internet. Wireless spectrum is finite, but demand is expected to increase exponentially. As more devices jostle for high speed access, interference becomes the fundamental bottleneck. Historically, interference has been managed either centrally (like in cellular networks), or with inefficient MAC protocols (such as 802.11 carrier sense). But as demand increases, existing protocols are either infeasible or inefficient in managing interference. Technologies such as OFDM and smart antennas offer a potential solution, by giving us fine-grained control on independent axes apart from time such as space and frequency on which we can partition network resources. Further, wireless networks offer rich diversity, both spatial (multiple nodes overhear the same transmission) and technological (different wireless technologies such as Wi-Fi, cellular, Bluetooth etc in the same neighborhood). Designing an agile radio which can take advantage of all these degrees of freedom remains a fundamental challenge.

The Network: Network design is a broad question and encompasses infrastructure design (mesh vs. cellular), switch design (for handling policy, handoff and mobility) and security (authentication, DOS etc) among other things. The mesh architecture is attractive because of its low deployment cost and significant coverage benefits, but comes at the cost of lower spectrum efficiency. Recent research (including my own thesis work) has produced significant improvements in spectrum efficiency for mesh networks. I believe advances in radio design (like the one I discussed above) can be exploited at the network level to produce significant further gains, likely making the mesh competitive with cellular systems in spectral efficiency, while retaining the advantages of easy deployment and large coverage. Determining, designing and building the right "mix" between cellular and mesh for the next-generation wireless infrastructure remains an exciting research challenge.

The Software Substrate: A new generation of distributed applications (e.g. the IPhone App Store) is emerging to take advantage of the ubiquitous connectivity and the presence of rich contextual personal data (e.g., user preferences, location etc) in mobile wireless devices to deliver exciting new services. Yet the landscape is fraught with difficulties. Wireless connectivity often has variable bandwidth and delay. Further, wireless devices have relatively limited computing power (Moore's law works mainly to reduce the size and energy footprint of these devices). And finally, allowing untrusted third party application developers access to private user data is necessary to build useful applications, but raises large security concerns. Designing a distributed software substrate which can overcome these difficulties to expose a predictable, efficient and secure API to applications is an exciting long-term research challenge.

The system I am building currently in my post-doctoral research, called **S3**, is a step in this direction. S3's goal is to protect private data from leaking to malicious hands. The key observation is that emerging hardware trends such as hardware assisted virtualization and multi-core computing provides us the ability to build low-level security hypervisors which can track the flow of sensitive information at the instruction level. Hardware assisted virtualization provides efficient memory access trapping capabilities, while multiple cores provide the ability to let the operating system speculatively execute while the hypervisor keeps a watchful eye on sensitive information in a separate core. When information is about to go on to an unprotected network, the hypervisor provides users and application developers with provenance information for the data in the packets, i.e., what private data if any were the contents of the packet derived from. Users and application writers can use this provenance information to write informed high-level security policies specifying which data is allowed to leave the system. S3 thus makes two important contributions; first it does not require us to trust complex operating system code and second, it can enforce simple security policies on sensitive data (such as exfiltration policies) without requiring any modifications to existing OS or applications.

My philosophy in research is to look for connections. I identify practical real-world problems, strip them to their core, and look for possible connections with otherwise unrelated areas. I exploit such conceptual connections to develop concrete techniques, and evaluate and verify them through analyses and practical implementations. This approach will continue to underpin my future research independent of the areas I work in.