Statement of Research Interests

Hitesh Ballani

I enjoy research. My research goal is to tackle problems afflicting the Internet. The explosive growth in the size of the Internet over the past couple of decades has meant that many of the assumptions that the Internet design is based on no longer hold true. This has led to a plethora of problems and has made it imperative that we rethink such assumptions and the concomitant design decisions. However, the tremendous success of the Internet has also been a bane for Internet research. It is difficult, if not impossible, to expect a wholesale change in Internet infrastructure. Throughout my graduate career, I have tried to stay cognizant of this ground truth and have strived to strike a balance between two competing urges regarding my research.

On one hand is the freedom of doing blue sky research that openly questions the fundamentals underlying the Internet architecture in the face of new needs and challenges. For instance, as part of my dissertation I argue that the fact that the Internet and its precursors started off as simple research networks meant that "manageability" was never a first-class design goal. Instead, humans operating the network were expected to delve into low-level network details in order to make it work. However, as the Internet becomes bigger and more complex, such an approach becomes intractable. Frustrated by the lamentable state-of-art in network management, I proposed and implemented *CONMan* [2,5,9], a cohesive architecture that leads to easy-to-manage and even largely self-managing networks. Similarly, my work on *DefaultOff* [10], a DoS-resilient Internet architecture recognises the mismatch between the security needs of the original and today's Internet and shows how this can be addressed by flipping default Internet reachability from "on" to "off".

On the other hand, I want my research to have practical impact through solutions that are immediately deployable. While it is well accepted that it is difficult to address the Internet's problems without changing the protocols involved, I have found that in many cases, simply by focussing on a subset of the given problem space, it is possible to devise a solution that does not require architectural change. In other words, it is possible to use existing protocols in novel ways to address some of the problems. I call this style of research *dirty-slate*. Such incremental solutions have a couple of important benefits. First, they offer a better alignment of cost vs benefits and hence, have a good chance of real-world adoption. Second, if the subset of problems solved happen to be the most pressing of the lot, such solutions buy the time needed for architectural proposals to mature and get deployed.

Over the past few years, I have developed a knack for recognising problems areas that can be alleviated in an incremental fashion and actually devising solutions that do not require change to deployed protocols and devices. For instance, I have invented *ViAggre* [1,4], a "configuration-only" approach to shrinking the routing table on Internet routers. While there are many other problems that afflict the Internet's routing system and *ViAggre* is not a cure-all, it does solve an important part of the routing problem area without requiring changes to both routing protocols and the routers themselves and without requiring global agreement. The same theme of tackling problems without protocol changes and global deployment pervades my work on mitigating DoS attacks on DNS nameservers [3,7] and making IP Anycast deployments practically feasible [8,11,12].

Overall, I feel that I have been successful at striking a good balance between clean-slate and incremental research. I have made a conscious effort to explore different facets of systems research in the context of my interest in networking. This appears in the fact that my work has included measurement studies, analytical studies, architectural proposals, implementation and even wide-area deployment.

1 Previous Work

Clean-slate research.

– **Network Management**. IP networks are hard to manage (install, configure, provision, monitor, test, debug). I have been able to boil down a large fraction of our management troubles to one specific shortcoming:

"Today, protocols and devices expose their internal details leading to a deluge of complexity that burdens the management plane." Consequently, a basic requirement for an Internet architecture conducive to management is that the management interface of protocols contain as little protocol-specific information as possible. This concept is at the core of my dissertation proposal called *CONMan* or Complexity Oblivious Network Management. Among other things, I have developed a generic abstraction that is used by CONMan-compliant protocols and devices to express their functionality to management applications. I have shown that this abstraction can serve as the narrow waist for the Internet's management plane, not much different from the way IP has served as the narrow waist for the Internet's data plane.

While the CONMan idea can be applied to all aspects of network management, my thesis focusses on the use of CONMan for configuration [5,9] and fault management [2]. To this effect, I implemented a suite of CONMan protocols and management applications. The resulting testbed can thus be configured by humans simply by specifying desired high-level goals. Further, any faults occurring in the network are automatically detected and localised. In effect, the CONMan testbed is a largely self-managing network. This work was one of the top-rated papers at ACM Sigcomm'07 and has evoked interest and received funding support from Cisco and NSF FIND.

– **Denial-of-Service attacks**. The original Internet architecture was designed to provide universal reachability; any host can send any amount of traffic to any destination. Unfortunately, today's less trustworthy Internet environment has revealed the downside of such openness—*every* host is vulnerable to attack by *any* other host(s). I argue that the simplest and the most direct approach to this problem is to flip Internet reachability on its head – end hosts should be "off by default" and should explicitly declare what traffic they want to be routed to them. This was the basis of the *DefaultOff* architecture [10] that, despite the seemingly intractable burden imposed on the Internet infrastructure, proved to be technologically feasible.

Dirty-slate research.

– **Routing Scalability**. The Internet routing table has been growing at a rapid rate for the past few years. Till now, it was almost universally accepted that maintaining a large routing table is just a fact of life in the existing architecture and there is no way to reduce the burden on Internet routers other than through a significant redesign. Consequently, most recent routing research has focussed on new architectures. While I agree that maintaining the entire routing table on routers is the de-facto mode of operation for Internet routing, I don't think that any part of the Internet architecture necessitates it. Instead, the load imposed by the routing table can be reduced simply by dividing the task of maintaining it amongst the routers.

This insight led me to invent *ViAggre* (Virtual Aggregation) [1,4], a scalability technique that shrinks the routing table on routers. ViAggre's most exciting feature is that it is a "configuration-only" approach. ViAggre does not require any changes to routers or routing protocols and can de deployed independently by any ISP on the Internet. It is mostly due to these reasons ViAggre has attracted a lot of positive attention in the IETF community and has been put on the standards track by IETF. As a matter of fact, a major router vendor (Huawei) is experimenting with implementing ViAggre natively into its routers.

Measurement results show that ViAggre can shrink routing tables by more than an order of magnitude. I also implemented a configuration tool that allows network operators to adopt ViAggre without manual router reconfiguration. Finally, I deployed ViAggre on a testbed of Cisco hardware routers. This exercise provided me with useful hands-on experience with routers used by commercial ISPs and am sure will help me in future routing research endeavors.

– **DNS Denial-of-Service attacks**. Recent years have seen many instances of Denial-of-Service (DoS) attacks on DNS, the Internet's naming system. I realised that such attacks are essentially targeting the skewed division of name resolution functionality between DNS servers and clients. Specifically, the attacks aim to make DNS servers unavailable and hence, can be tackled by doing away with the need for 100% server availability. Guided by this observation, I proposed a minor modification in DNS caching behavior that mitigates the impact of such attacks [3,7]. A long-term measurement study showed that even a single DNS resolver in the Internet can adopt this modification and acquire significant protection against DNS DoS

attacks. I am especially proud of this work; while certainly not the highlight of my research credentials, it does underscore my belief in dirty-slate research and shows how a very trivial hack can be used to address daunting problems in the existing setup.

– **Practical IP Anycast**. IP anycast, with its innate ability to find nearby resources in a robust and efficient fashion, has long been considered an important means of service discovery. However, it suffers from severe scalability problems. I argue that IP Anycast entails a tight coupling of the the anycast functionality to Internet routing mechanisms and this is the root-cause of its scalability problems. Hence, I designed, implemented and deployed *PIAS* (Proxy IP Anycast Service) [11,12], an anycast service that decouples anycast functionality from Internet routing.

As part of the PIAS deployment, I deployed a small testbed comprising of seven nodes spread across the Internet with each node advertising an address prefix into Internet routing. This, apart from the technical know-how involved, equipped me with an insight into the workings of commercial routing agreements. I have been maintaining this testbed for the past four years. Apart from the anycast service, the testbed provides a unique and very useful tool for researchers to perform active routing experiments. For instance, the testbed has been used by me [6,8] and other researchers [13,14] for wide-area BGP studies.

2 Future Work

From the beginning, I have focussed my research on problems in networking. This has allowed me to really delve deeply into one research area so as to have the most impact. This approach has worked well for me, and has led to a number of new ideas that I plan to pursue.

What to solve? A spate of routing research over the past few years has led to an agreement that severe problems afflict the routing system. For instance, scalability and manageability are commonly cited as the two biggest challenges facing Internet routing. However, there is still no agreement on what specific factors represent a scalability bottleneck and in what order: Is it memory or processing power or the ability to power the routers and dissipate the generated heat. Further, the answer depends not only on the characteristics of the routing system (such as the routing table size), but also on the characteristics of the routers themselves (such as the number of peers). My thesis is that only after understanding the key pain points can we come up with a practical alleviative for our routing pains. I am working on such benchmarking of the control plane on Internet routers. The ultimate goal of this exercise is to guide the design of the next generation routing system.

The Power of Tunnels. To a large extent, multihoming and traffic engineering have been the biggest contributors to the growth in the Internet routing table. This, in turn, is the result of the unholy marriage between Routing and Traffic Engineering. Specifically, Internet operators lack the appropriate primitives to control traffic flowing in and out of their networks. Instead, they try to achieve goals regarding network traffic through archaic knobs of routing protocols. The key problem with this marriage is the mismatch between the granularity provided by the routing system (prefix-level) and the granularity required for traffic engineering (flows or aggregates of flows).

On the other hand, recent years have seen an increasing use of tunnels in both enterprise networks and wide-area Internet. The adoption of MPLS by ISPs and the surge in the use of VPN technologies has promoted tunneling to a first-class Internet mechanism. In spite of this very public evolution, I am surprised at how little attention tunnels have received in the research community. My work on ViAggre exploits tunnels to shrink the routing table size on routers. Extending this, I believe that tunneling is a very powerful primitive and represents an opportunity to have a genuine impact on many Internet problems. Hence, my vision is to move the Internet towards an inter-domain tunneling system with the goal of tackling the complementary problems of routing scalability and traffic engineering.

Delay-tolerant Traffic. The rise of P2P and streaming applications has led to a substantial increase in traffic carried by ISPs and has caused them to take steps ranging from capacity upgrades to the much-maligned selective manipulation of traffic. However, a lot of such traffic is delay-tolerant. For instance, it is common for end-users to download movies only to watch them later. Such traffic can thus be "time-shifted" (delayed across time) to make it more amenable to the network. This, apart from reducing the peak burden on

the ISPs, can offer better performance to end-users. To this effect, I am collaborating with researchers at Telefonica Research towards the use of storage in the network to improve the network friendliness of delay-tolerant bulk traffic. Our current focus is on a CDN-like deployment model that has a symbiotic relationship with ISPs while benefiting the end-users. Looking ahead, I am also interested in an architectural solution to the bulk-traffic problem. For instance, is it feasible for the network to provide storage as a service or is it easier to deal with bulk traffic through a separate low-priority channel that does not interfere with existing (non delay-tolerant) traffic.

Beyond this, I believe that my past work has equipped me with the requisite tools to diversify my research and investigate problems beyond my core interests. For instance, my projects have included analysis, measurement, implementation and deployment and should hold me in good stead for research in networked systems, including enterprise networks, wireless networks, peer-to-peer systems and web services. In the near term, I want to build upon my work in security [3,6] to tackle challenges in securing networked systems. Further, I have a minor in statistics and am interested in game theory. Hence, I am exploring opportunities to apply the corresponding tools to solve network problems.

To summarise, my research will be geared towards building better networked systems with a focus on improving their scalability, manageability, security and reliability. Furthermore, my proposals will retain the dual theme of instant and delayed gratification. To this effect, I want my future research endeavors to track the following roadmap: Given a problem area, the first step will determine the largest subset that can be solved within the existing setup and develop such a solution. Next, the insights from the dirty-slate solution will guide a complete solution that will most likely require infrastructure change. While I do not claim that my proposals will be better at spurring architectural change, I do believe that when the cost of dealing with a problem becomes high enough to justify the costs associated with a change, my solutions will have a better alignment of costs and benefits and hence, a better chance of deployment.

References

[1] **Hitesh Ballani**, Paul Francis, Tuan Cao and Jia Wang, "Making Routers Last Longer with ViAggre," in *Proc. of USENIX Networked Systems Design and Implementation (NSDI)*, April 2009.

[2] **Hitesh Ballani** and Paul Francis, "Fault Management Using the CONMan Abstraction," in *Proc. of IEEE INFOCOM*, April 2009.

[3] **Hitesh Ballani** and Paul Francis, "Mitigating DNS DoS Attacks," in *Proc. of ACM Conference on Communications and Computer Security (CCS)*, October 2008.

[4] **Hitesh Ballani**, Paul Francis, Tuan Cao and Jia Wang, "ViAggre: Making Routers Last Longer!" in *Proc. of workshop on Hot Topics in Networks (Hotnets-VII)*, October 2008.

[5] **Hitesh Ballani** and Paul Francis, "CONMan: A Step towards Network Manageability," in *Proc. of ACM SIGCOMM*, August 2007.

[6] **Hitesh Ballani**, Paul Francis and Xinyang Zhang, "A Study of Prefix Hijacking and Interception in the Internet," in *Proc. of ACM SIGCOMM*, August 2007.

[7] **Hitesh Ballani** and Paul Francis, "A Simple Approach to DNS DoS Defense," in *Proc. of workshop on Hot Topics in Networks (HotNets-V)*, November 2006.

[8] **Hitesh Ballani**, Paul Francis and Sylvia Ratnasamy, "A Measurement-based Deployment Proposal for IP Anycast," in *Proc. of Internet Measurement Conference (IMC)*, October 2006.

[9] **Hitesh Ballani** and Paul Francis, "CONMan - Taking the Complexity out of Network Management," in *Proc. of Sigcomm Workshop on Internet Network Management (INM)*, September 2006.

[10] **Hitesh Ballani**, Yatin Chawathe, Sylvia Ratnasamy, Timothy Roscoe and Scott Shenker, "Off by default!" in *Proc. of workshop on Hot Topics in Networks (Hotnets-IV)*, November 2005.

[11] **Hitesh Ballani** and Paul Francis, "Towards a global IP Anycast service," in *Proc. of ACM SIGCOMM*, August 2005.

[12] **Hitesh Ballani** and Paul Francis, "Towards a Deployable IP Anycast Service," in *Proc. of First Workshop on Real, Large Distributed Systems (WORLDS)*, December 2004.

[13] Tongqing Qiu, L. Ji, D. Pei, J. Wang, J. J. Xu, and **Hitesh Ballani**, "LOCK: Locating Countermeasure-Capable Prefix Hijackers," GeorgiaTech, Tech. Rep. GT-CS-08-04, 2008.

[14] Zheng Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, , and R. Bush, "iSPY: Detecting IP Prefix Hijacking on My Own," in *Proc. of ACM SIGCOMM*, August 2008.