AFRL-IF-RS-TR-2004-184
Final Technical Report
June 2004

ENFORCEABLE NETWORK PROTOCOLS

University of Washington

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K294

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-184 has been reviewed and is approved for publication

APPROVED: /s/
SCOTT S. SHYNE
Project Engineer

FOR THE DIRECTOR: /s/
WARREN H. DEBANY
Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 2004	**FINAL** **Jun 00 – Dec 03**

4. TITLE AND SUBTITLE

ENFORCEABLE NETWORK PROTOCOLS

5. FUNDING NUMBERS

G - F30602-00-2-0565
PE - 62301E
PR - K294
TA - 00
WU - 01

6. AUTHOR(S)

Tom E. Anderson
David J. Wetherall

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Washington
Department of Computer Science & Engineering
Box 352350, 185 Stevens Way
Seattle WA 98195-2350

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington VA 22203-1714

AFRL/IFGA
525 Brooks Road
Rome NY 13441-4505

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

AFRLIF-RS-TR-2004-184

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Scott S. Shyne/IFGA/(315) 330-4819 Scott.Shyne@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 Words)***

We propose to dramatically improve the reliability, fault tolerance, and survivability of wide area networks by systematically rethinking network design using the principle that the operation of network protocols should be enforceable – the correct operation of a protocol should not depend on trust. We propose to rethink four areas of network protocol design: the information exchange at the interface between nodes, the transient behavior of protocols under failure, the placement of enforcement functionality at the edges or inside the network, and the detailed control over resource allocation inside of routers. We will deliver a suite of network protocols – for routing, transport, multicast, and real –time – which have enforceable behavior without trust.

14. SUBJECT TERMS

security, networking, protocols, misbehavior, diagnostics

15. NUMBER OF PAGES 17

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

TABLE OF CONTENTS

SUMMARY:

This is the final technical report for "Enforceable Network Protocols" (Grant No. F30602-00-2-0565). Over the past three and a half years, work on this project at the University of Washington has laid the foundation for dramatic improvements in the security and reliability of networks. The overall thrust of the work has been to identify potential vulnerabilities in existing, widely used, network protocols, and to develop more robust and secure versions of those protocols. We have also developed practical, incremental approaches to improving robustness and security, where a redesign is impractical or only a long-term solution.

During the course of our investigation, we identified and pursued several commercial opportunities to bring our ideas into widespread practice, recognizing the immediate, highly visible problems being encountered by industry in developing highly secure and reliable network systems. Combined with subsequent cuts to the contract budget, this forced us to reduce the scope of effort somewhat, specifically the effort to examine the vulnerabilities of certain routing protocols to Byzantine transient behavior.

Nevertheless, the project was broadly successful in meeting its goals. Among its contributions was the first comprehensive study of past Internet failures, motivating a set of design guidelines for more robust and secure protocols. We applied these guidelines in several concrete case studies, including influencing IETF drafts on ECN and IP traceback to fix potential vulnerabilities before they could be exploited by attackers. We have developed practical tools for identifying the source of both BGP configuration errors and router/link level performance failures, key steps in reducing the mean time to repair. We have also taken a longer view, developing a concrete proposal based on receiver-driven capabilities to completely eliminate the potential for distributed denial of service attacks, as well as a novel protocol for identifying and containing misbehaving routers in an ad hoc network.

INTRODUCTION:

Our project work aims to dramatically improve the reliability, fault tolerance, and survivability of wide area networks by systematically rethinking network design using the principle that the operation of network protocols should be *enforceable* -- the correct operation of a protocol should not depend on trust that the other participating members are operating correctly. The Internet today, by contrast, has a bimodal security model -- trusted entities are authenticated and protected from untrusted attackers, but no further checks are placed on the information provided by an entity once it has been authenticated. The result is fragile systems -- once an attacker (or inadvertent error) slips in, the scope of error is potentially unlimited. Our approach is trust but verify -- to redesign network protocols to carefully limit the scope of information provided by a (putatively) trusted entity to be either (i) directly verifiable or (ii) to have no negative impact if the information is untrue. While this may seem difficult in general, in this contract we have demonstrated many examples where this design principle can apply.

We note that our work strongly complements traditional security approaches based on encrypted communication between authenticated participants. Encrypted communication is a necessary first step to securing a system – otherwise, an attacker can transparently manipulate packets to accomplish their ends. However, authentication is insufficient for a truly robust and secure system. Authentication controls only who can talk to you, while enforcement and verification protects you from what (authenticated but corrupt) peers might say.

Our efforts have examined vulnerabilities, and proposed solutions for, a broad spectrum of network protocols, to demonstrate that these ideas can be widely applied, specifically interdomain, intradomain, and ad hoc routing, endpoint and router-based congestion control, distributed denial-of-service, and diagnostic tools to quickly locate the source of problems.

METHODS, ASSUMPTIONS, AND PROCEDURES:

The research direction that we propose sits in-between two traditional kinds of system security. On one hand, the cryptography community has developed strong authentication and encryption technologies that can be applied to the design of distributed systems. These technologies aim to prevent unauthorized parties from gaining access to the system. IPSEC is an example of this approach. Their advantage is that, when used properly, they present a hard line of defense against outside attackers. Nonetheless, the resulting security is brittle in the sense that if one node of the system is compromised, then the system as a whole has been compromised. They also do not protect against damage that is caused when trusted nodes become faulty, or are simply misconfigured.

On the other hand, the intrusion detection community is developing techniques that may be used to determine when a system is being attacked. If it is possible to identify an attack at an early stage, then it should be possible to deter the attack or otherwise limit the damage that results. This task is immensely complicated by the design of existing protocols and the typically passive approach of observing activity from select vantage points. It is not possible to decide in general whether activity stems from a legitimate user or an attacker masquerading as a legitimate user.

Our strategy for this year is to build on our study of vulnerabilities in existing protocols and deliver a set of solutions -- proof of concept demonstrations and working systems to show that network protocols can be designed and implemented to be enforceable. This movement from analysis of existing protocols to the design and implementation of revised protocols is staggered for the different protocols that we are studying. That is, we will ramp up the study of new target protocols at the same time that we deliver revised solutions for our first target protocols, building on the experience that we have gained.

Our strategy has three phases. For each protocol domain that we study, we first study the most widely used protocols in significant depth, analyzing for potential attacks that can come from within the set of trusted participants. The next phase is then to design

mechanisms to eliminate, where possible, the attacks that we identify. Here, we convert the root cause of the attack into behavior that can be checked and thus safely enforced. This is not always possible, but our experience has shown that more often than not, hidden trust relationships embedded in protocols are unnecessary to the goals of the protocol. Finally, once we have designed new mechanisms we develop real prototype implementations for the Internet. This includes an assessment of backwards-compatible heuristics to defeat the new attacks and incremental deployment mechanisms that would allow the new solutions to be incorporated into the Internet. Throughout, a key part of our work was developing incentives for ISPs to deploy more robust and secure solutions, e.g., by designing more robust solutions that are also more efficient.

RESULTS AND DISCUSSION:

We discuss the major results of the project below, initially focusing on those results that identify potential problems, before moving to the more concrete proposals for solutions: new protocols and ways to fix old protocols in a backwardly-compatible fashion.

Design Guidelines for Robust Protocols. An early, highly successful part of our work was to do the equivalent of a FAA crash investigation for the Internet [1]. We examined the historical record to find as many actual failures as possible, dating from several decades ago to the better documented problems of the last few years. We then deconstructed each failure – not just how it was fixed at the time, but more importantly, how different protocol design guidelines, if followed, might have avoided the failure in the first place. Our intent was to demonstrate to the community the value of a more systematic approach to robust network protocol design.

Specifically, we recommended six design guidelines for robust network protocol design. Some are obvious (but still violated frequently in practice), such as to value conceptual simplicity – to avoid overloading mechanisms with multiple purposes, as this can cause other problems later. A second guideline is to minimize your dependencies – to carefully examine the information flow through the protocol, and eliminate any unnecessary dependence that may be a channel for errors to propagate. Again, this may seem obvious, but it has been violated repeatedly in the past. Where dependence is necessary, protocol designers should verify, where possible, to avoid blindly trusting the other participants in the protocol. Many denial of service attacks result from a failure to protect resources, the fourth guideline – as a recent example, Code Red caused router failures as an unintended side effect, because of the failure of the router development team to carefully analyze when received packets could cause resource exhaustion. Even if you believe a protocol is well-designed, practical problems can be avoided if you limit the scope of any vulnerability – many Internet problems start small, but propagate widely, the very definition of an asymmetric threat. Finally, if errors are to be fixed, they must be exposed. By contrast, most Internet protocols are designed according to the end to end principle. This tends to hide errors, allowing them to accumulate until they impact

robustness. To summarize, the goal of this work was to provide for us, in our work on new more robust protocol designs, and the community at large, with a list of best practices, so that protocol designers in the future will not repeat the errors of the past.

Interdomain Routing: Misconfiguration and Design. In a related study, we studied traces of BGP announcements, to determine how frequently BGP misconfigurations occurred in practice [2]. We focused on those announcements that were either short term or inconsistent with other BGP announcements. We found misconfigurations to be surprisingly frequent; for example, new prefixes announced via BGP are more likely than not to be due to a misconfiguration! Unfortunately, BGP misconfigurations are inherently difficult to prevent, because of the low level of the BGP specification. Instead, this work motivated us to target a higher level, more abstract protocol that uses BGP as an implementation mechanism, rather than as a specification language [3]. Users describe their policies to this higher level mechanism, resulting in more concise, easier to verify policies with fewer misconfigurations. As a first step in this direction, we developed a simple, yet surprisingly effective protocol for adjacent ISPs to use in choosing which traffic is to traverse each peering point [4].

Specifically, current interdomain routing policies are largely based on information local to each ISP, in part due to competitive concerns and the lack of effective information sharing mechanisms. This can lead to routes that are sub-optimal and even to routes that oscillate. This provides motivation for ISPs to use an automated protocol to replace what is currently a very manual, and error prone, process. We explored a setting in which two neighboring ISPs negotiate to determine the path of the traffic they exchange. We first asked the basic question: Is there an incentive to negotiate? The incentive exists only if both ISPs benefit relative to routing based on local information. Through simulation with over sixty measured ISP topologies, we found that negotiation is useful for both latency reduction and hotspot avoidance – that is, automated solutions have the potential to be both more efficient and more robust. Based on our results, we designed and evaluated a negotiation protocol which works within the real-world constraints of competing and independently managed ISPs. Specifically, our protocol reveals little information and works even when ISPs have different optimization criteria. We found that it achieves routing performance comparable to that of (the obviously infeasible) global optimization using complete information from both ISPs. Initial reactions from network architects at major ISPs are overwhelmingly positive.

Enforceable Congestion Control and ECN. Another major effort has been to identify and fix security vulnerabilities in popular congestion control mechanisms. It is well known that hosts sending TCP traffic can be corrupted to send traffic at arbitrary rates. Our work was the first to demonstrate that a malicious (or simply greedy) TCP receiver could spoof a TCP sender into ignoring congestion signals and sending at an arbitrary rate, in essence converting the sender into an unwitting partner in denying resources to competing traffic to the misbehaving receiver [5]. Further, we were able to show that existing proposals for ECN - congestion notification by routers – likewise assumes trust of the receiver to correctly forward congestion signals [6]. Such trust is unnecessary, however. We designed and implemented a set of small modifications to the TCP and

ECN specification that removed this dependence on the correct behavior of the receiver, by inserting random nonces in packets at the sender. Nonces are returned by the receiver to prove correct receipt of packets. Routers may force congestion to be signaled either by dropping a packet, or simply by erasing the nonce. This approach, very much in line with our design guideline of "trust but verify", is remarkably effective, and we have succeeded at having it added to the IETF ECN standard [7].

Distributed Denial of Service Traceback and Prevention. Vulnerability to distributed denial of service attacks has long been seen as a fundamental weakness of the Internet protocol architecture, but only recently have widespread attacks been mounted to exploit this vulnerability. The increasing severity and sophistication of attacks has meant that there can be no practical assurance of the ability for mission-critical traffic to be delivered over the public Internet. A consequence is that mission-critical communication is reserved for leased lines, at a huge increase in expense. Our initial work in this area was to design a practical protocol for embedding traceback information in packets, enabling victims to trace attacks back to their source [8]. While our traceback proposal was robust, a related proposal to send traceback information in separate ICMP packets was not. Our analysis identified a vulnerability in the ICMP-based traceback proposal, where the receiver could be confounded by an attacker sending spoofed traceback messages along with the attack packets. This vulnerability has since been fixed in later versions of the ICMP traceback draft.

However, traceback is only one step in the process of dealing with attacks, and to date router vendors have been slow to deploy the rich functionality needed to mount a practical defense based on traceback. Not only do we need to know the location of the source of attack packets, we also need rich, programmable filters in routers (since attack packets can be designed to be arbitrarily close in format to legitimate traffic), along with protocols to quickly push filters into the network in response to attacks, along with an authentication infrastructure for ensuring that only legitimate users install filters on their traffic. There is no evidence that such a complete solution is likely to be provided by router vendors anytime in the near future.

To address the fact that existing technologies to deal with denial of service attacks seems to be reaching a dead end, later in the contract we took a fresh look at the problem. From an architectural level, the reason that distributed denial-of-service attacks are possible is that any source is architecturally free to send any amount of traffic to any destination at any time. Instead, we argue that a better design would be to put recipients in control of which packets they want to receive. We developed a concrete protocol design to achieve this, based on short-term tokens, or capabilities, sent by receivers to senders [9]. Senders request permission to send, receive explicit tokens in reply, and include the tokens in packets as proof that the traffic was requested by the destination. Routers in the middle of the network can efficiently enforce the protocol by dropping packets that lack tokens. Although our proposal was speculative and time and budget constraints prevented us from fully implementing our ideas, several related proposals have recently appeared in the literature that build on our work, as an indication of the promise of our approach.

Robustness in Ad Hoc Wireless Networks. We recently completed the design and implementation of an enforcement-based protocol for deterring selfish behavior in multi-hop wireless networks [10]. There is little work in this area, despite the fact that most wireless MAC layers assume cooperation among all participants, and therefore are vulnerable to misconfigured and malicious participants. Our protocol detects selfish nodes that avoid forwarding packets for others and punishes them by refusing them multi-hop connectivity. The protocol works by having the neighbors of each node send it anonymous packets. If the node drops these packets, the neighbors can detect this, and refuse to send packets to the node. If the node forwards these packets, then that provides verification that the node can successfully receive and forward the packets, so that the neighbors can punish the node if it forwards the anonymous packets but not later data packets. Our simulation and live testbed experiments show that the protocol detects cheaters quickly, distinguishes between cheaters and non-cheaters, and handles sophisticated cheating strategies effectively. Our initial study in this area suggests that our approach can yield significant benefits and would strongly complement efforts to improve wireless encryption protocols.

Internet Diagnosis. Finally, we have also developed tools for diagnosing performance problems along Internet paths [11]. Currently, an attacker can effectively disrupt communication by causing problems in the middle of the network, and without diagnostic tools, these problems can be very difficult and time-consuming to track down. Our approach was two-fold. First, we show that annotating packets with diagnostic information is nearly as powerful as collecting a complete packet trace, for diagnosing performance faults along specific paths. Since collecting a complete packet trace is infeasible, this is a tremendously useful result, as it implies that a practical architecture for diagnostic support is possible. Second, we show that existing hooks in the Internet architecture provide much of the functionality required by our theoretical model. Based on these results, we built a practical tool for Internet diagnosis, called *Tulip*, that can isolate performance problems to the specific failing link. The tool is currently being used by Boeing to debug performance problems on cross-continental Internet paths, problems that would otherwise prevent teams of engineers from collaborating in real-time.

CONCLUSIONS:

In this report, we have described our work on Enforceable Network Protocols, our methods, and our major results. We have identified many potential vulnerabilities in existing, widely used, network protocols, and helped put other well-known vulnerabilities into context. More importantly, we have developed more robust and secure versions of those protocols that do not suffer from these vulnerabilities. Finally, we have developed several practical, incremental approaches to improving robustness and security. Our work on these topics is both important and timely -- addressing fundamental weaknesses in the Internet architecture and protocols in a manner that is not being tackled by industry and directly contributing to a strengthened and more robust cyber-infrastructure.

REFERENCES:

[1] Thomas Anderson, Scott Shenker, Ion Stoica, and David Wetherall. *"Design Considerations for More Robust Internet Protocols."* First Workshop on Hot Topics in Networks (HotNets-I), October 2002. Also appeared in Special Interest Group in Communications (SIGCOMM), Computer Communications Review, January 2003.

[2] Ratul Mahajan, David Wetherall, and Thomas Anderson. *"Understanding BGP Misconfiguration."* Special Interest Group in Communications (SIGCOMM), August 2002.

[3] Ratul Mahajan. *"Negotiation-Based Routing."* Workshop on Internet Routing Evolution and Design (WIRED), October 2003.

[4] Ratul Mahajan, David Wetherall, and Thomas Anderson. *"Interdomain Routing with Negotiation."* University of Washington Technical Report, May 2004.

[5] Stefan Savage. *"Protocol Design in an Uncooperative Internet."* Ph.D. Dissertation, University of Washington, January 2002.

[6] Neil Spring, David Ely, David Wetherall, Stefan Savage, and Thomas Anderson. *"Robust Congestion Signaling."* International Conference of Network Protocols (ICNP), November 2001.

[7] Neil Spring, David Ely, and David Wetherall. *"Robust ECN Signaling with Nonces."* Internet Draft, draft-ietf-tsvwg-tcp-nonce-02.txt, October 2001.

[8] Stefan Savage, David Wetherall, Anna Karlin, and Thomas Anderson. *"Practical IP Traceback."* IEEE/ACM Transactions on Networking, June 2001.

[9] Thomas Anderson, Timothy Roscoe, and David Wetherall. *"Preventing Internet Denial of Service with Capabilities."* Second Workshop on Hot Topic in Networks (HotNets-II), November 2003.

[10] Maya Rodrig, Ratul Mahajan, David Wetherall, and John Zahorjan. *"Deterring Selfish Behavior in Multi-Hop Wireless Networks."* UW Technical Report, February 2004.

[11] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. *"User-Level Internet Path Diagnosis."* Symposium on Operating Systems Principles (SOSP), October 2003.

Appendix A: Publications

Eric Anderson. *"A Multicommodity Flow Based Approach to Adaptive Internet Routing."* Ph.D. dissertation, University of Washington, June 2002.

Eric Anderson and Thomas Anderson. *"On the Stability of Adaptive Routing in the Presence of Congestion Control."* INFOCOM, June 2002.

Thomas Anderson, Scott Shenker, Ion Stoica, and David Wetherall. *"Design Considerations for More Robust Internet Protocols."* First Workshop on Hot Topics in Networks (HotNets-I), October 2002. Also appeared in Special Interest Group in Communications (SIGCOMM), Computer Communications Review, January 2003.

Thomas Anderson, Timothy Roscoe, and David Wetherall. *"Preventing Internet Denial of Service with Capabilities."* Second Workshop on Hot Topic in Networks (HotNets-II), November 2003.

Ratul Mahajan, David Wetherall, and Tom Anderson. *"Interdomain Routing with Negotiation."* Paper in progress, http://www.cs.washington.edu/research/networking/negotiation/bits/negotiation.pdf, May 2004.

Ratul Mahajan, Sally Floyd, and David Wetherall. *"Controlling High-Bandwidth Flows at Congested Routers."* International Conference of Network Protocols (ICNP), November 2001.

Ratul Mahajan, David Wetherall, and Thomas Anderson. *"Understanding BGP Misconfiguration."* Special Interest Group in Communications (SIGCOMM), August 2002.

Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. *"Inferring Link Weights Using End-to-End Measurements."* ACM Internet Measurement Workshop, November 2002.

Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. *"User-Level Internet Path Diagnosis."* Symposium on Operating Systems Principles (SOSP), October 2003.

Ratul Mahajan. *"Negotiation-Based Routing."* Workshop on Internet Routing Evolution and Design (WIRED), October 2003.

Ratul Mahajan, Miguel Castro, and Antony Rowstron. *"Controlling the Cost of Reliability in Peer-to-Peer Overlays."* International Workshop on Peer-to-Peer Workshop (IPTPS), November 2003.

Maya Rodrig, Ratul Mahajan, David Wetherall, and John Zahorjan. *"Deterring Selfish Behavior in Multi-Hop Wireless Networks."* UW Technical Report, February 2004.

Stefan Savage. *"Protocol Design in an Uncooperative Internet."* PhD Thesis, University of Washington, March 2002.

Stefan Savage, David Wetherall, Anna Karlin, and Thomas Anderson. *"Network Support for IP Traceback."* Special Interest Group in Communications (SIGCOMM), August 2000. Also appeared in IEEE/ACM Transactions on Networking, June 2001.

Neil Spring, David Ely, and David Wetherall. *"Robust ECN Signaling with Nonces."* Internet Draft, draft-ietf-tsvwg-tcp-nonce-00.txt , February 2001.

Neil Spring, David Ely, and David Wetherall. *"Robust ECN Signaling with Nonces."* Internet Draft, draft-ietf-tsvwg-tcp-nonce-02.txt, October 2001.

Neil Spring, David Ely, David Wetherall, Stefan Savage, and Thomas Anderson. *"Robust Congestion Signaling."* International Conference of Network Protocols (ICNP), November 2001.

Neil Spring, Ratul Mahajan, and David Wetherall. *"Mapping ISP Topologies with Rocketfuel."* Special Interest Group in Communications (SIGCOMM), August 2002. Received best student paper award.

Neil Spring, David Wetherall, and Thomas Anderson. *"Scriptroute: A Public Internet Measurement Facility."* Usenix Symposium on Internet Technologies and Systems (USITS), March 2003. Received best student paper award.

Neil Spring, Ratul Mahajan, and Thomas Anderson. *"Quantifying the Causes of Path Inflation."* Special Interest Group in Communications (SIGCOMM), August 2003.

Neil Spring, David Wetherall, and Thomas Anderson. *"Reverse-Engineering the Internet."* Second Workshop on Hot Topic in Networks HotNets-II, November 2003.

Neil Spring, Ratul Mahajan, David Wetherall, and Thomas Anderson. *"Measuring ISP Topologies with Rocketfuel."* ACM Transactions on Networking, February 2004

Andrew Whitaker, Parveen Patel, Jay Lepreau, and David Wetherall. *"TCP Meets Mobile Code."* Hot Topics in Operating Systems (HotOS), October 2003.

Andrew Whitaker, Parveen Patel, Jay Lepreau, and David Wetherall. *"Upgrading Transport Protocols Using Mobile Code."* Symposium on Operating Systems Principles (SOSP), October 2003.

Appendix B: Presentations at Meetings and Conferences

Thomas Anderson. "*Enforceable Network Protocols.*" DARPA Fault Tolerant Networks PI Meeting, August 2001.

Thomas Anderson. "*Design Considerations for Robust Internet Protocols.*" First Workshop on Hot Topics in Networks, October 2002.

Thomas Anderson. "*Design Considerations for Robust Internet Protocols.*" Distinguished Lecture, University of Illinois at Urbana-Champaign, February 2003.

Thomas Anderson. "*Design Considerations for Robust Internet Protocols.*" Invited talk, University of Wisconsin-Madison, February 2003.

Neil Spring. "*Robust Congestion Signaling.*" International Conference on Network Protocols (ICNP), November 2001.

Ratul Mahajan. "*De-Misconfiguring BGP.*" Presented at the UW Networking Workshop, Gig Harbor, WA, June 2001.

Ratul Mahajan. "*The Impact of BGP Misconfiguration on Connectivity.*" North American Network Operators Group (NANOG) 23, October 2001.

Ratul Mahajan. "*Controlling High-Bandwidth Flows at Congested Routers.*" International Conference of Network Protocols (ICNP), November 2001.

Ratul Mahajan. "*A Study of BGP Misconfiguration.*" Presented at UW-CSE Affiliates Meeting, February 2002.

Ratul Mahajan. "*Understanding BGP Misconfiguration.*" Presented at Special Interest Group in Communications (SIGCOMM), August 2002.

Ratul Mahajan. "*Towards Operator Fault-Tolerant Inter-Domain Routing.*" IEEE Computer Communication Workshop (CCW), October 2002.

Ratul Mahajan. "*Inferring Link Weights Using End-to-End Measurements.*" ACM Internet Measurement Workshop (IMW), November 2002.

Ratul Mahajan. "*Controlling the Cost of Reliability in Peer-to-Peer Overlays.*" International Workshop on Peer-to-Peer Systems (IPTPS), February 2003.

Ratul Mahajan. "*Pinpointing Performance Faults Along Internet Paths.*" UW-CSE Affiliates Program, February 2003.

Ratul Mahajan. *"User-level Internet Path Diagnosis."* Symposium on Operating Systems Principles (SOSP), October 2003.

Ratul Mahajan. *"Negotiation-Based Routing."* Workshop on Internet Routing Evolution and Design (WIRED), October 2003.

Stefan Savage. *"Protocol Design in an Uncooperative Internet."* Presented at University of Washington, Stanford University, University of California at Berkeley, Carnegie-Mellon University, MIT, University of California at San Diego, Winter 2002.

Stefan Savage. "*IP Traceback Techniques."* North American Network Operators Group (NANOG), June 2000.

Stefan Savage. "*Enforceable Network Protocols.*" Fault Tolerant Network PI Meeting, July 2000.

Stefan Savage. *"Practical IP Traceback."* Special Interest Group in Communications (SIGCOMM) 2000, Stockholm, Sweden, August 2000.

Neil Spring. "*Robust ECN Signaling with Nonces.*" Presented at the Transport Working Group, 50th IETF, Minneapolis, February 2001.

Neil Spring. *"Local Repair for OSPF."* Presented at the UW Networking Workshop, Gig Harbor, WA, June 2001

Neil Spring. *"Robust ECN Signaling with Nonces."* Presented at the Transport Working Group, 51st IETF, London, August 2001.

Neil Spring. *"Robust ECN Signaling with Nonces.,"* IEEE Computer Communications Workshop, October 2001.

Neil Spring. "*Robust Congestion Signaling."* International Conference of Network Protocols, November 2001.

Neil Spring. *"Mapping ISP Topologies with Rocketfuel."* Presented at the Systems/Networking Workshop, University of Washington, June 2002.

Neil Spring. *"Measuring ISP Topologies with Rocketfuel."* Symposium on Operating Systems Principles (SOSP), October 2003.

Neil Spring. *"Measuring ISP Topologies with Rocketfuel."* Special Interest Group in Communications (SIGCOMM), August 2002.

Neil Spring. *"Scriptroute: A Facility for Distributed Internet Measurement."* North American Network Operators Group (NANOG 26), Eugene, OR, October 2002.

Neil Spring. *"Measured Causes of Internet Path Inflation."* UW-CSE Affiliates Program, February 2003.

Neil Spring. *"Scriptroute: A Public Internet Measurement Facility."* USENIX Symposium on Internet Technologies and Systems (USITS), March 2003.

Neil Spring. *"Quantifying the Causes of Path Inflation."* Special Interest Group in Communications (SIGCOMM), August 2003.

Neil Spring. *"What Can We Do with Measurements?"* IRIS/PlanetLab Meeting, MIT, August 2003.

Neil Spring. *"Reverse-Engineering the Internet."* Hot Topics in Networks (HotNets), November 2003.

David Wetherall. "Enforceable Network Protocols." DARPA Fault Tolerant Networks PI Meeting, Florida, January 2001.

David Wetherall. *"BGP: A Surprisingly Robust Distributed System."* Work-In-Progress talk at Symposium on Operating Systems (SOSP), October 2001.

David Wetherall. *"A Measurement Study of BGP Misconfiguration."* Network Seminar, Stanford University, November 2001.

David Wetherall. *"Towards More Error-Tolerant Internet Protocols."* Presented at UW-CSE, January 2002.

David Wetherall. *"Enforceable Network Protocols: BGP Misconfiguration."* Presented at DARPA Fault Tolerant Networks PI Meeting, San Diego, January 2002.

SYMBOLS, ABBREVIATIONS, AND ACRONYMS

ACK ……. Acknowledgment

API …… Application Protocol Interface

ATM …….. Asynchronous Transfer Mode

BGP …….. Border Gateway Protocol

CHATS ……... Composable High Assurance Trusted Systems

ECN …….. Explicit Congestion Notification

FTN ……. Fault Tolerant Network

ICMP …….. Internet Control Message Protocol

IETF …….. Internet Engineering Task Force

IP …… Internet Protocol

IPSEC ……... IP Security Protocol

ISP ……. Internet Service Provider

MAC …….. Media Access Control

MTTF ……. Mean Time to Failure

MTTR …….. Mean Time to Repair

NAT ……... Network Address Translation

TCP …….. Transmission Control Protocol

VPN …… Virtual Private Network