Internet Scalability: Properties and Evolution

Matthew Roughan

Steve Uhlig

Walter Willinger

The modern Internet is a large-scale distributed system composed of many complex interoperating entities: traffic, topology, routing, addressing, and more. With its continued growth, scalability has become an increasingly important issue, as has the nature of the complex interactions between these entities. The current Internet architecture suffers from scalability problems in different respects: growth in traffic, topological size, multihoming, and complex interactions across the TCP/IP protocol stack are all stressing the current infrastructure. The evolution of multihoming impacts both the behavior of traffic and the routing plane. The deployment of IPv6 and node mobility might also fundamentally change how traffic and routing behaves in the future. Understanding any of these aspects requires a thorough study of the relationships between them. Obviously, we cannot expect to cover in a single special issue of *IEEE Network* all aspects of Internet scalability. In this issue we gather contributions that deal with routing, naming, attacks against hosts and routers, and, finally, its topology.

In recent years the exponential growth of the number of Internet users with increased bandwidth demands has led to the emergence of the next generation of IP routers. Distributed architectures are a promising trend, providing petabit routers with huge switching capacity and high-speed interfaces. Distributed routers are designed with an optical switch fabric interconnecting line and control cards. Computing and memory resources are available on both control and line cards in order to perform routing and forwarding tasks. These new hardware architectures are not efficiently utilized by the traditional software models where a single control card is responsible for all routing and management operations. The routing table manager (RTM) plays a highly critical role by managing routing information and, in particular, a forwarding information table (FIT). In “A Distributed and Scalable Routing Table Manager for Next Generation IP Router” Nguyen, Jaumard, and Agarwal present a distributed architecture set up around a distributed and scalable RTM that also provides improvements in robustness and resiliency.

Today’s Internet would be more efficient and robust if routers could flexibly divide traffic over multiple paths. Often, having just a few alternate paths is sufficient for customizing paths for different applications, improving security, reacting to failures, and balancing load. Yet alternate paths are often unavailable due to two barriers. First, in a network the size of the Internet, moving to flexible multipath routing can impose significant computational and storage overhead on routers. Second, the independent networks that make up the Internet will not relinquish control over the flow of traffic without appropriate incentives. In “Toward Internet-Wide Multipath Routing” He and Rexford survey flexible multipath routing techniques that are both scalable and incentive-compatible. Techniques covered include multihoming, tagging, tunneling, and extensions to existing Internet routing protocols.

A fundamental building block of the Internet is the naming system. The Domain Name System (DNS) is the global lookup service for network resources. To protect DNS information, the DNS Security Extensions (DNSSEC) have been developed and deployed on branches of the DNS to provide authentication and integrity protection using digital signatures. However, signed DNS nodes have been found to have an unfortunate side effect: an attacker can query them as reconnaissance before attacking hosts on a particular network. There are different ways a zone administrator can minimize information leakage while still taking advantage of DNSSEC for integrity and source authentication. In “Minimizing Information Leakage in the DNS” Rose and Nakassis describe the risk, examine the protocol and operational options, and look at their advantages and drawbacks.

While attacks against the DNS can be particularly disruptive to large numbers of nodes, traffic surges due to worm attacks are able to partially disrupt routing. Two of the articles in this issue are related to the spread of worms and detecting compromised routers. Through modeling, Debany considers in "Modeling the Spread of Internet Worms via Persistently Unpatched Hosts" the effects of Internet worms on persistently unpatched hosts and hosts for which vulnerabilities are refreshed. Current models have assumed that all hosts transition through the same set of states. Debany goes beyond this simplification of host behavior by allowing hosts to have inherently different characteristics. Equilibrium conditions are obtained under which an Internet worm will self-propagate indefinitely, which leads to thresholds below which worms will become extinct.

While compromising hosts is now very common for various purposes (e.g., bots, spyware, spam forwarding), attacks against routers are far less appreciated. The central role of routers in end-to-end communication makes it possible to leverage them for eavesdropping, man-in-the-middle, and denial-of-service attacks. In response, a range of specialized anomaly-detection protocols has been proposed to detect misbehaving packet forwarding between routers. Mizrak, Mazrullo, and Savage discuss in "Detecting Compromised Routers via Packet Forwarding Behavior" the detection of misbehaving packet forwarding between routers. The article describes a general framework for understanding the available design space and reviews the capabilities of various detection protocols.

Last but not least, in "Network Models with a 'Soft Hierarchy': A Random Graph Construction with Loglog Scalability" Norros and Reittu are concerned with the scalability of topologies in general and Internet-like topologies in particular. Many graph models have been proposed to describe the Internet's topology at the IP or autonomous system level, most of them having power-law node degrees. The article discusses how a special class of very simple network models, random graphs with infinite variance power law degree, have features that may or may not be desirable from an Internet perspective. The proposed random graphs exhibit loglog scalability and possess a fascinating architecture with a softly hierarchical core network.

Finally, we would like to thank the Liaison Editor, Professor Kazem Sohraby, for helping in the final steps of the reviewing process, as well as the external reviewers for their valuable time.

Biographies

MATTHEW ROUGHAN (matthew.roughan@adelaide.edu.au) joined the School of Applied Mathematics at the University of Adelaide in February 2004, where he is interested in the area of design and installation of Internet measurement equipment, and the analysis and modeling of Internet measurement data. He previously worked in this area for four years at AT&T, and at Ericsson in Australia (via the Universities of Melbourne, and the Royal Melbourne Institute for Technology) for another four. Prior to this, he worked at the Cooperative Research Centre for Sensor Signal and Information Processing (CSSIP), Adelaide, Australia, on diverse projects ranging from the analysis of ionograms to land-mine detection. He gained his Ph.D. from the University of Adelaide in 1994 in applied mathematics, and has now returned to the same department to teach.

STEVE UHLIG (S.P.W.G.Uhlig@tudelft.nl) received a Ph.D. degree in applied sciences from the University of Louvain-la-neuve, Belgium, in March 2004, on interdomain traffic engineering. He received the 2005 IBM/F.N.R.S. prize for his Ph.D. thesis. From 2004 to 2006 he was a postdoctoral fellow of the Belgian National Fund for Scientific Research. In October 2006 he joined the Network Architecture and Services group of the Technical University of Delft as an assistant professor. His research interests include interdomain routing, traffic engineering and characterization, and Internet measurements.

WALTER WILLINGER [F'05] (walter@research.att.com) received a Dipl.Math. from ETH Zurich, Switzerland, and M.S. and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, New York. He is currently a member of the Information and Software Systems Research Center at AT&T Labs -Research, Florham Park, New Jersey, and before that was a member of technical staff at Bellcore Applied Research (1986–1996). His research interests include studying the multiscale nature of Internet traffic and topology, and developing a theoretical foundation for dealing with large-scale communication networks such as the Internet. He is a Fellow of ACM (2005). For his work on the self-similar ("fractal") nature of Internet traffic, he received the 1996 IEEE W.R.G. Baker Prize Award, the 1994 W.R. Bennett Prize Paper Award, and the 2005 ACM/SIGCOMM "Test of Time" Paper Award.