Internet Traffic Measurement

Carey Williamson
Department of Computer Science
University of Calgary

November 24, 2001

Abstract

This tutorial article discusses the role of network traffic measurement in the design, testing, and evaluation of Internet protocols and applications. The article begins with some background information on Internet traffic measurement, and then proceeds to discuss the "tools of the trade", including examples of both hardware-based and software-based approaches to network traffic measurement. The article concludes with a summary of the main observations from the past fifteen years of network measurement research, along with pointers to the relevant literature for more information.

Keywords: Network traffic measurement, workload characterization, Internet protocols, TCP/IP, performance

1 Introduction

The evolution of the Internet over the last thirty years has been accompanied by the development, growth, and use of a wide variety of network applications. These applications range from text-based utilities such as file transfer, remote login, electronic mail, and network news from the early days of the Internet, to the advent of desktop videoconferencing, multimedia streaming, the World-Wide Web, and electronic commerce on today's Internet.

For the vast majority of Internet users, the Internet is simply a medium to provide the connectivity for these end-user applications. End users are shielded from having to know the details of how the Internet works, through the use of an "Internet protocol stack" (see Figure 1) that follows the principle of *information hiding*. This protocol stack, which is implemented in full on each Internet host, and in part on each Internet router, defines the rules for how user-level data is transformed into network packets for transport across the network, and put back together for delivery at the receiving application. For example, a user who clicks on a hyperlink on a Web page need not know how many network packets are generated, nor how these packets are routed through the network; the lower-layer protocols such as IP (Internet Protocol, which provides global addressing and routing on the Internet) and TCP (Transmission Control Protocol, which provides reliable end-to-end data delivery) look after this task. The user cares only about how long it takes for the chosen Web page to appear on the screen.

For many networking researchers, the Internet itself is the *raison d'être*, and the Internet protocols are themselves the subject of study. These researchers often rely on *network traffic measurement* as a methodology to collect and analyze data regarding the operation (and performance) of network protocols.

Network traffic measurement provides a means to go "under the hood", much like an Internet mechanic, to understand what is or is not working properly on a local-area or wide-area network. Using specialized network measurement hardware or software, a networking researcher can collect detailed information about the transmission of packets on the network, including their timing structure and contents. With detailed packet-level measurements, and some knowledge of the Internet protocol stack, it is possible to "reverse engineer" significant information about the structure of an Internet application, or the behaviour of an Internet user.

There are four main reasons why network traffic measurement is a useful methodology:

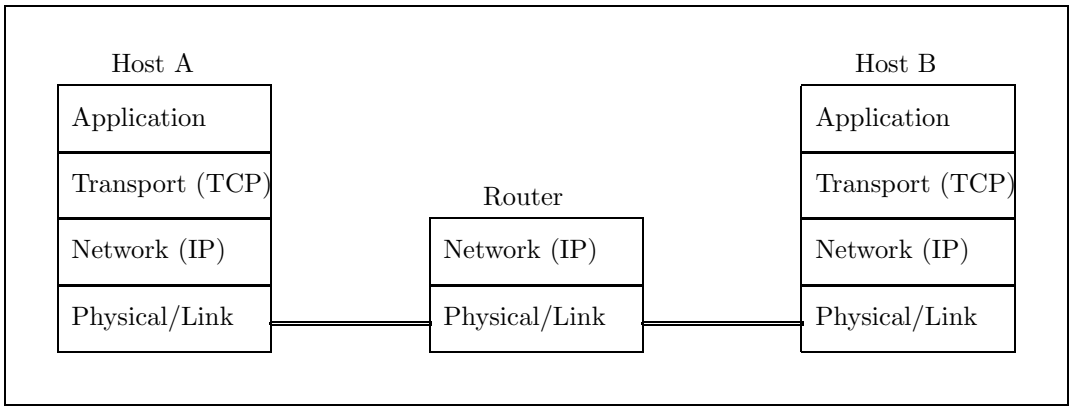

Figure 1: Illustration of the Internet TCP/IP Protocol Stack

- **Network Troubleshooting**. Computer networks are not infallible. Often, a single malfunctioning piece of equipment can disrupt the operation of an entire network, or at least degrade performance significantly. Examples of such scenarios include "broadcast storms", illegal packet sizes, incorrect addresses, and security attacks. In such scenarios, detailed measurements from the operational network can often provide a network administrator with the information required to pinpoint and solve the problem.
- **Protocol Debugging**. Developers often want to test out "new, improved" versions of network applications and protocols. Network traffic measurement provides a means to ensure the correct operation of the new protocol or application, its conformance to required standards, and (if necessary) its backward-compatibility with previous versions, prior to unleashing it on a production network.
- **Workload Characterization**. Network traffic measurements can be used as input to the *workload characterization* process, which analyzes empirical data (often using statistical techniques) to extract salient and representative properties describing a network application or protocol. Knowledge of the workload characteristics can then lead to the design of better protocols and networks for supporting the application.
- **Performance Evaluation**. Finally, network traffic measurements can be used to determine how well a given protocol or application is performing in the Internet. Detailed analysis of network measurements can help identify performance bottlenecks. Once these performance problems are addressed, new versions of the protocols can provide better (i.e., faster) performance for the end users of Internet applications.

Given these motivations for the use of network traffic measurements, the remainder of this article proceeds to describe the research methodology of network traffic measurement, and some of the highlights from the network traffic measurement literature in the past fifteen years.

2 Network Traffic Measurement Methodology

The "tools of the trade" for network traffic measurement research can be classified in several different ways:

- **Hardware-based vs. Software-based Measurement Tools**

  The primary categorization among network measurement tools is hardware-based versus software-based measurement tools. Hardware-based tools are often referred to as network traffic analyzers: special-purpose equipment designed expressly for the collection and analysis of network data. This equipment is often expensive, depending on the number of network interfaces, the types of network cards, the storage capacity, and the protocol analysis capabilities. Many vendors offer these types of products.

Software-based measurement tools typically rely on kernel-level modifications to network interfaces of commodity workstations to convert them into machines with packet capture capability. One widely used utility is *tcpdump*, a user-level tool for TCP/IP packet capture, made possible with the Berkeley Packet Filter (BPF). In general, the software-based approach is much less expensive than the hardware-based approach, but may not offer the same functionality and performance as a dedicated network traffic analyzer.

Another software-based approach to workload analysis relies on the access logs that are recorded by Web servers and Web proxies on the Internet. These logs record each client request for Web site content, including the time of day, client IP address, URL requested, and document size. Post-processing of such access logs provides useful insight into Web server workloads [1], without the need to collect detailed network-level packet traces.

- **Passive vs. Active Measurement Approaches**

  A *passive* network monitor is used to observe and record the packet traffic on an operational network, without injecting any traffic of its own onto the network. That is, the measurement device is non-intrusive. Most network measurement tools fall into this category.

  An *active* network measurement approach uses packets generated by a measurement device to probe the Internet and measure its characteristics. Examples of this approach include the *ping* utility for estimating network latency to a particular destination on the Internet, the *traceroute* utility for determining Internet routing paths, and the *pathchar* tool for estimating link capacities and latencies along an Internet path.

- **On-line vs. Off-line Traffic Analysis**

  Some network traffic analyzers support real-time collection and analysis of network data, often with graphical displays for on-line visualization of live traffic data. Most hardware-based network analyzers support this feature.

  Other network measurement devices are intended only for real-time collection and storage of traffic data; analysis is postponed to an off-line stage. The *tcpdump* utility falls into this category. Once the traffic data is collected and stored, a researcher can perform as many analyses as desired in the post-processing phase.

- **LAN vs. WAN Measurement**

  The early network traffic measurement research in the literature was undertaken in Local Area Network (LAN) environments, such as Ethernet LANs. LANs are easier to measure, for several reasons. First, a LAN is typically administered by a single well-known organization, meaning that obtaining security clearance for traffic analysis is a relatively straightforward process. Second, the broadcast nature of an Ethernet LAN means that all packets transmitted on the LAN are seen by all hosts. Network measurement can be done in this context by simply configuring a network interface into *promiscuous mode*, which means that the interface will receive and record (rather than ignore) the packets destined to other hosts on the network.

  Later measurement work extended traffic collection and analysis to Wide Area Network (WAN) environments [4, 11, 12]. The greater challenges here include administrative control of the network, and security and privacy issues. For organizations with a single access point to the external Internet, measurement devices can be put in-line on an Internet link near the default router for the organization.

  More recently, Barford and Crovella [2] discussed the deployment of a Wide-Area Web Measurement (WAWM) infrastructure that can collect simultaneous measurements of client, server, and network behaviours. With time coordination between these measurements, a more complete picture of end-to-end network performance is possible.

- **Protocol Level**

  Measurement tools differ in the protocol level at which data is collected and traffic analysis is performed. Many network traffic analyzers support multi-layer protocol analysis, but a specialized network card must be used for each type of network on which traffic data is to be collected. For example, specialized network

TIME STAMP	PROTOCOL	SOURCE IP_ADDRESS	SRC PORT	DESTINATION IP_ADDRESS	DST PORT	IP_PKT SIZE	TCP SEQ	TCP ACK
0	IP TCP	307.246.129.64	1060	427.86.12.704	80	40	920641	412791
14966	IP TCP	561.877.104.57	7410	427.86.12.704	80	508	410104	32779
15015	IP TCP	391.82.374.90	1105	891.82.59.75	80	40	2816846	7726
22090	IP TCP	719.327.502.359	1140	526.837.913.44	80	40	1010185	14762
22126	IP TCP	582.127.755.91	1291	419.74.87.6	80	40	9557082	50482
29960	IP TCP	561.877.104.57	3741	427.86.12.704	80	40	985526	58006
29960	IP TCP	419.74.87.6	80	582.127.755.91	1291	1500	653402	57082
31724	IP TCP	419.74.87.6	80	582.127.755.91	1291	1500	654862	57082
36055	IP TCP	512.84.9.317	1125	419.74.87.628	80	311	857517	89873
36279	IP TCP	512.84.9.317	1126	419.74.87.628	80	271	857661	3293
37181	IP TCP	407.84.92.183	1207	398.54.73.39	5190	40	64202	9407
41731	IP TCP	399.81.77.33	80	342.406.374.91	1116	40	1062629	68778

Figure 2: An Example of a TCP/IP Packet Trace File from ISP Network Measurements

cards exist for Ethernet, Frame Relay, Asynchronous Transfer Mode (ATM), and wireless networks, but the back-end protocol analysis engines for IP and higher-layer protocols may be the same.

2.1 Example 1: Hardware-based Measurements

This section presents a small example of network measurements collected using a NavTel IW95000 ATM network traffic analyzer. These measurements were collected from an Internet Service Provider (ISP) running an IP over ATM backbone network in 1998 [7].

The ATM network analyzer provides non-intrusive capture of cell-level ATM traffic streams, including packet headers and payloads. The analyzer timestamps each ATM cell with a one microsecond timestamp resolution, and records the captured traffic into memory in a compressed proprietary binary data format. The size of the memory capture buffer and the volume of the network traffic determine the maximum interval of time for which traces can be collected (e.g., typically several seconds at 155 Mbps OC-3 rates, and several minutes at 1.5 Mbps T1 rates).

Once the capture buffer is full, traces can be saved to disk or copied to another machine for off-line trace analysis. The analyses in [7] used a custom C program to decode the data recorded by the NavTel IW95000. This program converts the binary data file into an ASCII format with TCP/IP protocol information.

An example of the human-readable trace format is shown[1] in Figure 2. This format includes a timestamp (in microseconds, relative to the start of the trace), the protocol types recognized, and then selected fields from the IP and TCP packet headers, such as IP source and destination address, IP packet size (including TCP and IP headers), TCP source and destination port numbers, and TCP sequence number information, both for data packets and for acknowledgment packets.

Once available in this latter form, it is straightforward to construct customized scripts to process a trace file and extract the desired information, such as timestamp, packet size, as well as IP and TCP protocol information. Off-line analyses of traces such as these were used in a study of ISP Web traffic characteristics [7].

2.2 Example 2: Software-based Measurements

One of the most popular software-based measurement tools is *tcpdump*. The *tcpdump* tool enables IP packet-level capture from an operational network, along with filtering of captured traffic streams based on specific host

[1]Note that the IP addresses throughout the paper have been "sanitized" to conceal their true identities.

```
  TIME         SOURCE    PORT   DESTINATION  PORT  FLAG  SEQNUM    ACKNUM
19:52.731470 406.17.8.12.64826 > 723.65.19.6.www: S 4256930:4256930(0)
19:52.731889 723.65.19.6.www > 406.17.8.12.64826: S 768500:768500(0) ack 4256931
19:52.732200 406.17.8.12.64826 > 723.65.19.6.www: . ack 768501 win 17520
19:52.738205 406.17.8.12.64826 > 723.65.19.6.www: P 4256931:4257101(170) ack 768501
19:52.743248 723.65.19.6.www > 406.17.8.12.64826: P 768501:5769840(1339) ack 4257101
19:52.758535 406.17.8.12.64826 > 723.65.19.6.www: F 4257101:4257101(0) ack 5769840
19:52.758862 723.65.19.6.www > 406.17.8.12.64826: . ack 4257102
19:52.759700 723.65.19.6.www > 406.17.8.12.64826: F 5769840:5769840(0) ack 4257102
19:52.759935 406.17.8.12.64826 > 723.65.19.6.www: . ack 5769841
```

Figure 3: An Example of a tcpdump Packet Trace File

Table 1: Main Observations from Recent Network Traffic Measurement Research

1	Internet traffic continues to change.
2	Aggregate network traffic is multi-fractal in nature.
3	Network traffic exhibits "locality" properties.
4	Packet traffic is non-uniformly distributed.
5	Packet sizes are bimodally distributed.
6	The session arrival process is Poisson.
7	The packet arrival process is not Poisson.
8	Most TCP conversations are short-lived.
9	Traffic flows are bidirectional, but often asymmetric.
10	TCP accounts for most of the packet traffic on the Internet.

addresses, port numbers, or protocol types. Tcpdump has been widely used by networking researchers for the study of Internet applications, and for studying growth trends in Internet traffic over time [12].

An example of a *tcpdump* trace appears in Figure 3. This trace format shows a timestamp for each packet, and the IP and TCP headers, which carry address and control information. Post-processing of the trace in Figure 3 can extract application-level behaviours, such as the Web document transfer represented in this example.

3 Highlights

This section summarizes the "Top 10" observations from the network traffic measurement literature in the past fifteen years. These observations are summarized in Table 1, and discussed briefly here.

The main observations (in reverse order) are the following:

10. **TCP accounts for most of the packet traffic on the Internet.**

    Early network traffic measurement research [4] showed that TCP was the dominant protocol on the Internet in the early 1990's. Several recent (and popular) Internet applications, such as video streaming, napster, IP telephony, and multicast, rely predominantly on the UDP protocol rather than TCP, and may gradually shift the balance of traffic away from TCP. However, the measurements presented in [14] suggest that TCP is still the dominant traffic force on the Internet, and is likely to remain so for the foreseeable future. The primary reason is the advent of the World Wide Web: the growing number of Internet users, the widespread availability of easy-to-use Web browsers, and the proliferation of Web sites with rich multimedia content combine to contribute to the exponential growth of Internet TCP traffic. Web caching and content distribution networks help to soften this impact [3, 15], but the overall growth is still dramatic.

9. **Traffic flows are bidirectional, but often asymmetric.**

   Many of the network applications used on the Internet generate a bidirectional exchange of data, though the volume of data sent in each direction often differ greatly. This observation was true for the Internet applications studied in the early 1990's [4, 11], and is even truer today because of the download-intensive nature of the Web. It remains to be seen what effect large-scale peer-to-peer networking paradigms, such as napster and grids, will have on traffic asymmetry in the Internet.

8. **Most TCP conversations are short-lived.**

   Over 90% of the TCP conversations studied by Caceres *et al.* [4] exchanged fewer than 10 kilobytes of data, and lasted less than a few seconds. This prevalence of short-lived connections was somewhat surprising at the time, particularly for file transfer and remote login applications. However, this conversation paradigm has been significantly reinforced with the advent of the Web. The literature suggests that approximately 80% of Web document transfers are less than 10 kilobytes in size [1], though there is a significant heavy tail to the distribution [1, 6].

7. **The packet arrival process is not Poisson.**

   Much of the classical work in queueing theory and communication network design is based on the assumption that the packet arrival process is Poisson. In simple terms, the Poisson arrival process means that events (e.g., earthquakes, traffic accidents, telephone calls, customer arrivals, packet arrivals) occur independently at random times, with a well-defined average rate. More formally, the inter-arrival times between events are exponentially distributed and independent, and no two events happen at exactly the same time.

   Poisson models are attractive mathematically because of the "memoryless" property of the exponential distribution: even knowing the time that has elapsed since the last event provides no hint as to when the next event will occur. These types of models are often amenable to elegant mathematical analysis, leading (for example) to closed-form expressions for the mean waiting time (and variance) in queueing network models.

   Detailed studies of Internet network traffic show that the packet arrival process is not Poisson. That is, the inter-arrival times between packets are not exponentially distributed, nor are they independent. Rather, the packet arrival process is bursty: packets arrive in "clumps" that make the traffic far more bursty than predicted by a Poisson process [13]. As a result, the queueing behaviour can be much more variable than predicted by a Poisson model.

   This non-Poisson structure is due in part to the protocols used for data transmission. This observation casts doubt on the value of simple (Poisson) network traffic models used in network performance studies, and has been the impetus behind recent research on network traffic modeling [11].

6. **The session arrival process is Poisson.**

   Although the packet arrival process is not Poisson, there is strong evidence that the *session* arrival process is Poisson. That is, human Internet users seem to operate independently at random when initiating access to certain Internet resources. This observation has been noted for several network applications. For example, Paxson and Floyd [13] study telnet traffic and find that the session arrival process is well-modeled with a Poisson process, though with a time-varying rate (e.g., hourly). Similarly, Arlitt and Williamson [1] find that the user requests for individual Web pages on a Web server are often well-modelled by a Poisson process.

5. **Packet sizes are bimodally distributed.**

   The size (in bytes) of the network packets traversing the Internet have a "spiky" distribution [4]. Many of the packets (about 50%) are "as large as possible", meaning that they carry the maximum number of data bytes permitted based on the Maximum Transmission Unit (MTU) parameter defined for a network interface. Many packets (about 40%) are small in size (40 bytes), because of the prevalence of (header-only)

TCP acknowledgement packets for data received. The remaining 10% of the packets are somewhat randomly scattered in between these two extremes, based on how much user data remains in the "last" data packet of a multi-packet transfer. Occasionally, secondary spikes occur in the distribution due to IP fragmentation between networks with different MTU sizes.

4. **Packet traffic is non-uniformly distributed.**

   Analyzing the source and destination addresses carried in TCP/IP packets shows that packet traffic is often highly non-uniformly distributed amongst the hosts on the Internet. A common observation is that 10% of the hosts account for 90% of the traffic. In some sense, this observation is not that surprising, given the client-server paradigm for many network applications. However, the prevalence of this property in many network measurement studies suggests a fundamental power-law structure in many aspects of Internet traffic [1, 3, 6], and even in certain aspects of Internet topology [8].

3. **Network traffic exhibits "locality" properties.**

   The structure of network traffic is far from random. Traffic structure is imposed implicitly by the tasks initiated by Internet users at the application layer (e.g., a file transfer or a Web page download), and reinforced by the TCP/IP protocols used for data transer within the Internet. Packets are not independent and isolated entities; rather they are part of a higher-layer logical flow of information. This flow manifests itself at the network layer in recognizable (though not necessarily predictable) patterns of packet timing and source and destination addresses. In the literature, this structure is often referred to as "temporal locality" (i.e., time-based correlation of information) or "spatial locality" (i.e., geography-based correlation).

2. **Aggregate network traffic is multi-fractal in nature.**

   Characterizing aggregate network traffic is difficult, for many reasons: the heterogenous nature of the Internet, the diverse mix of network applications, the wide variations in link speeds and network access technologies, the time-varying nature of Internet evolution, and changing user behaviours.

   Nevertheless, networking researchers have identified a significant degree of long-range dependence (LRD) in network traffic, which has been referred to as "self-similar", "fractal", and "multi-fractal" behaviour [9, 10]. This LRD property appears to be ubiquitous: it is present in LAN, WAN, video, data, Web, ATM, FrameRelay, and SS7 signalling traffic. The LRD property is attributed in part to heavy-tailed on-off behaviours of Internet users, perhaps exacerbated by the TCP/IP protocols used on the Internet [6]. More recent research [5] addresses the non-stationarity of Internet traffic, and suggests that the multi-fractal traffic structure evident at the edges of the network diminishes within the core of a large network.

   Despite the complex multi-fractal structure of Internet traffic, surprisingly concise mathematical models have been developed and used for the characterization and analysis of Internet traffic. These models are used in performance studies that attempt to understand the implications of LRD traffic on the design of the future Internet infrastructure.

1. **Internet traffic continues to change.**

   Longitudinal studies, such as that by Paxson [12], have shown that Internet traffic continues to grow and change, over relatively short time scales. This change is not simply one of traffic volume, but also one of traffic mix, protocols, applications, and users. Despite the value of Internet traffic measurement as a research methodology, one must realize that any data set collected from an operational network represents but one snapshot at one point in time in the evolution of the Internet.

   Trying to identify the *invariants* in traffic structure is one means to cope with the never-ending battle of measuring and understanding Internet traffic.

```
Internet Traffic Archive (ITA). See http://ita.ee.lbl.gov
A public-domain repository of traces and data sets collected by networking researchers.

Internet Traffic Report (ITR). See http://www.InternetTrafficReport.com
Hourly statistics on global Internet traffic trends

National Laboratory for Applied Network Research (NLANR). See http://www.nlanr.net
A national (U.S.-based) initiative on high-performance networking.

NLANR Measurement and Operations Analysis Team (MOAT). See http://moat.nlanr.net
A subgroup of NLANR researchers specializing in Internet traffic measurement.
Provides on-line Internet traffic statistics, traces, and tools.

National Internet Measurement Infrastructure (NIMI). See http://www.ncne.nlanr.net/nimi/
An NLANR initiative to provide ubiquitous measurement capability
for Internet traffic, topology, routing, and quality of service.

tcpdump. See http://www.tcpdump.org/
Public-domain software for collecting network-level packet traces.
```

Figure 4: For More Information on Internet Traffic Measurement

4 Summary

Internet traffic measurement is an applied networking research methodology aimed at understanding packet traffic on the Internet. From its humble beginnings in LAN-based measurement of network applications and protocols, network measurement research has grown in scope and magnitude, and has helped provide insight into fundamental behavioural properties of the Internet, its protocols, and its users. Recent initiatives (see Figure 4) strive to provide a practical and scalable infrastructure for wide-scale operational measurement of today's global Internet.

References

[1] M. Arlitt and C. Williamson, "Internet Web Servers: Workload Characterization and Performance Implications", *IEEE/ACM Transactions on Networking*, Vol. 5, No. 5, pp. 815-826, October 1997.

[2] P. Barford and M. Crovella, "Measuring Web Performance in the Wide Area", *ACM Performance Evaluation Review*, Vol. 27, No. 2, pp. 37-48, September 1999.

[3] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web Caching and Zipf-like Distributions: Evidence and Implications", *Proceedings of the IEEE Infocom Conference*, New York, NY, pp. 126-134, March 1999.

[4] R. Caceres, P. Danzig, S. Jamin and D. Mitzel, "Characteristics of Wide-Area TCP/IP Conversations", *Proceedings of ACM SIGCOMM*, Zürich, Switzerland, pp. 101-112, September 1991.

[5] J. Cao, W. Cleveland, D. Lin, and D. Sun, "On the Nonstationarity of Internet Traffic", *Proceedings of ACM SIGMETRICS*, Cambridge, MA, pp. 102-112, June 2001.

[6] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", *IEEE/ACM Transactions on Networking*, Vol. 5, No. 6, pp. 835-846, December 1997.

[7] R. Epsilon, J. Ke, and C. Williamson, "Analysis of ISP IP/ATM Network Traffic Measurements", *ACM Performance Evaluation Review*, Vol. 27, No. 2, pp. 15-24, September 1999.

[8] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-Law Relationships of the Internet Topology", *Proceedings of ACM SIGCOMM*, pp. 251-262, Cambridge, MA, September 1999.

[9] A. Feldmann, A. Gilbert, and W. Willinger, "Data Networks as Cascades: Explaining the Multi-Fractal Nature of Internet Traffic", *Proceedings of ACM SIGCOMM*, pp. 42-55, Vancouver, BC, September 1998.

[10] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, pp. 1-15, February 1994.

[11] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections", *IEEE/ACM Transactions on Networking*, Vol. 2, No. 4, pp. 316-336, August 1994.

[12] V. Paxson, "Growth Trends in Wide Area TCP Connections", *IEEE Network*, Vol. 8, No. 4, pp. 8-17, July/August 1994.

[13] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", *IEEE/ACM Transactions on Networking*, Vol. 3, No. 3, pp. 226-244, June 1995.

[14] K. Thompson, G. Miller, and R. Wilder, "Wide-area Internet Traffic Patterns and Characteristics", *IEEE Network*, Vol. 11, No. 6, pp. 10-23, November/December 1997.

[15] C. Williamson, "On Filter Effects in Web Caching Hierarchies", *ACM Transactions on Internet Technology*, Vol. 2, No. 1, February 2002, in press.

Biographical Information

Carey Williamson is a Professor in the Department of Computer Science at the University of Calgary in Calgary, Alberta, Canada where he holds an iCORE Senior Research Fellowship in Broadband Wireless Networks, Applications, Protocols, and Performance. His research interests include Internet protocol performance, network traffic measurement, and network simulation.
Readers can contact the author at carey@cpsc.ucalgary.ca