Network Monitoring and Measurement at High Speeds—MAGNeTizing Traffic Through the Protocol Stack

RADIANT: Research And Development In Advanced Network Technology
http://www.lanl.gov/radiant

Network researchers often use traffic libraries such as `tcplib` or network traces such as those at http://ita.ee.lbl.gov/html/traces.html to drive their network experiments, particularly to test the performance of network-protocol enhancements. However, such libraries, traces, and models are based on measurements made by `tcpdump/libpcap` (or related tools), meaning that the traffic an application sends on the network is captured only after having passed through TCP (or more generally, any protocol stack) and into the network. That is, the tools capture traffic on the wire (or in the network) rather than at the application level. Thus, the above tools cannot provide any protocol-independent insight into the actual traffic patterns of an application.

So, researchers have not been testing the performance of network protocols using real, application-generated traffic traces; rather, they have used once-modulated (by the protocol stack) traffic traces as input, which are subsequently modulated a second time during the testing of the protocol. If the differences between application-generated traces and network-captured traces are negligible, such a simplification is acceptable. However, MAGNeT (monitor for application generated network traffic), a toolkit that we recently developed at Los Alamos National Laboratory,[1] demonstrates that the differences in the traces are substantial, indicating that the protocol stack adversely modulates the application-generated traffic patterns. This observation may invalidate the empirical results gathered from the performance evaluation of network protocols over the last decade because many researchers used network-captured traces

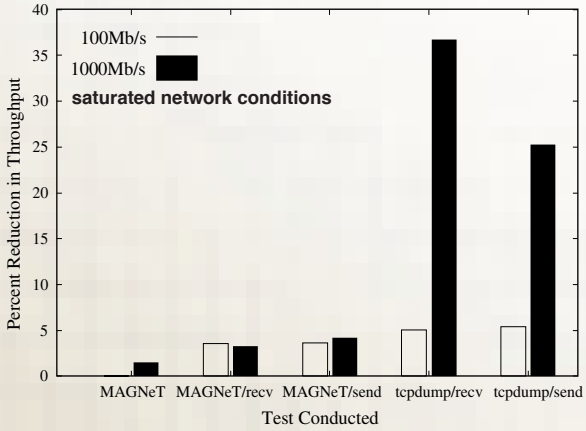

Fig. 1. Percent Reduction in Network Throughput.

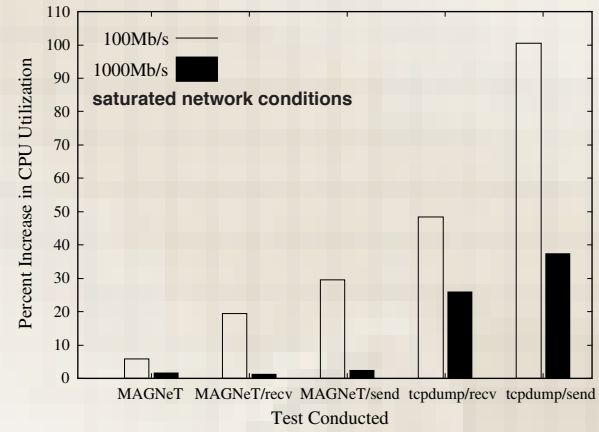

Fig. 2. Average Percent Increase in CPU Utilization.

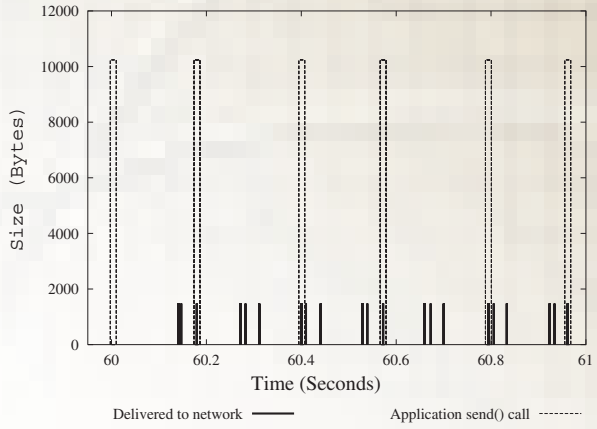

Fig. 3. MAGNeT FTP Trace.

Table 1. MAGNeT vs. `tcpdump`

Configuration	Throughput (Mb/s)	Send CPU (%)	Receive CPU (%)
Linux 2.4.3	94.1 ± 0.0	15.2 ± 0.1	33.5 ± 0.1
MAGNeT	94.1 ± 0.1	16.9 ± 0.2	33.5 ± 0.1
`magnet-read/rcv`	90.8 ± 0.8	20.7 ± 0.3	34.4 ± 1.0
`magnet-read/snd`	90.7 ± 0.9	23.7 ± 1.7	32.4 ± 0.4
`tcpdump/rcv`	89.4 ± 1.5	18.0 ± 0.4	59.8 ± 0.9
`tcpdump/snd`	89.4 ± 0.8	45.0 ± 0.6	31.9 ± 0.3

(a) 100Mbs (Fast Ethernet)

Configuration	Throughput (Mb/s)	Send CPU (%)	Receive CPU (%)
Linux 2.4.3	459.5 ± 1.6	61.0 ± 0.3	82.4 ± 0.2
MAGNeT	452.5 ± 1.8	63.0 ± 0.4	82.6 ± 0.3
`magnet-read/rcv`	444.3 ± 1.7	62.4 ± 0.3	82.0 ± 0.3
`magnet-read/snd`	440.2 ± 2.1	63.1 ± 0.5	81.1 ± 0.4
`tcpdump/rcv`	290.7 ± 15.6	36.1 ± 2.0	91.5 ± 0.5
`tcpdump/snd`	343.2 ± 18.7	93.2 ± 0.5	64.1 ± 3.3

(b) 1000Mbs (Gigabit Ethernet)

For more information on MAGNeT, visit http://www.lanl.gov/radiant or reference the following publications:

• W. Feng, J. Hay, and M. Gardner, "MAGNeT: Monitor for Application-Generated Traffic," Proc. of the 10th IEEE Int'l Conf. on Computer Communications and Networks (IC3N), October 2001.

• J. Hay, W. Feng, and M. Gardner, "Capturing Network Traffic with a MAGNeT," Proc. of the 5th Annual USENIX Linux Showcase & Conference, November 2001.

[1]*The MAGNeT software distribution is currently undergoing alpha testing at LANL. A beta prototype will soon be publicly available from* http://www.lanl.gov/radiant.

Contact Information

radiant-info@lanl.gov

RADIANT: Research And Development in Advanced Network Technology
(http://www.lanl.gov/radiant)
Los Alamos National Laboratory
Los Alamos, NM 87545

Sponsored in part by DOE LDRD and DOE Office of Science.

November 2001

LALP-01-246

A US DEPARTMENT OF ENERGY LABORATORY

rather than application-generated traces to drive their experiments. What network researchers need are traces of real, application-generated traffic—traces that can be provided by MAGNeT.

Performance

Here we demonstrate MAGNeT's ability to record application and network-stack events without adversely perturbing the traffic stream or application behavior. We also use MAGNeT-collected data to show significant differences between an application's network demands and the actual traffic delivered to the network.

Table 1 provides an indication of how much MAGNeT and `tcpdump` perturb application generated traffic. Along with the mean, the width of the 95% confidence interval is given. Figures 1 and 2 present this data graphically. To produce these results, we conducted our tests between two identical dual-400MHz Pentium IIs with a 100-Mbps or 1000-Mbps Ethernet card (connected via an Extreme Networks Summit 7i switch) and configured MAGNeT to record application `send()` and `recv()` socket calls as well as TCP and IP events. MAGNeT uses the default 256-KB kernel buffer to store event records. To generate the workload, we run netperf on the sender, transmitting data as fast as possible, i.e., saturating the network.

While no monitoring system can be completely transparent to the workload being monitored, Table 1 demonstrates that MAGNeT minimizes its impact as compared to `tcpdump`.

Traffic Pattern Analysis

We can use MAGNeT-collected data to investigate the differences between the traffic generated by an application and that same traffic as it appears on the network (i.e., after modulation by a protocol stack). As a simple example, we consider a trace of a FTP session from our facility in Los Alamos, NM to a location in Dallas, TX. Figure 3 shows one second of a MAGNeT trace taken one minute into the FTP transfer.

As can be seen from the graph, the FTP application attempts to send 10-KB segments of data every 20 ms, but the protocol stack (TCP and IP in this case) modulates the traffic into approximately 1500-byte packets at intervals of varying duration. Since the maximum data size on an Ethernet network is 1500 bytes, the protocol stack segments the data to this size. The variable spacing of the traffic intervals is mainly caused by TCP waiting for positive acknowledgements before sending more traffic.

Los Alamos
NATIONAL LABORATORY
Los Alamos NM 87545