Harpoon: A Flow-Level Traffic Generator for Router and Network Tests

Joel Sommers Hyungsuk Kim Paul Barford
University of Wisconsin—Madison

http://wail.cs.wisc.edu/

Motivation for Flow-Level Traffic Generation

- Emulation and simulation environments require tools for reproducible creation of a range of test conditions similar to those experienced in the live Internet.
- Existing tools (*e.g.*, SURGE, `iperf`, specialized hardware like Spirent AX/4000) for creating background traffic cannot recreate the rich variety of packet streams observed in the Internet.

Harpoon Features and Benefits

Harpoon is unique in that it:

- scalably generates statistically representative network traffic at the IP flow level;
- is application-independent;
- recreates temporal volume (byte, packet, flow) characteristics of live traces;
- recreates spatial characteristics (source and destination IP address frequencies);
- self-configures (IMC '04 submission) from Netflow logs or packet traces – there is no parametric estimation required.

Architecture of Harpoon

Flow records or packet traces collected from live environments are used in a self-configuration step to parameterize Harpoon's two-level model. As depicted in the figure to the right, empirical distributions of file sizes, inter-connection times, source and destination IP address frequencies, and the average number of sessions active over a series of consecutive time intervals are used to generate traffic that is statistically identical to the originally measured traffic.

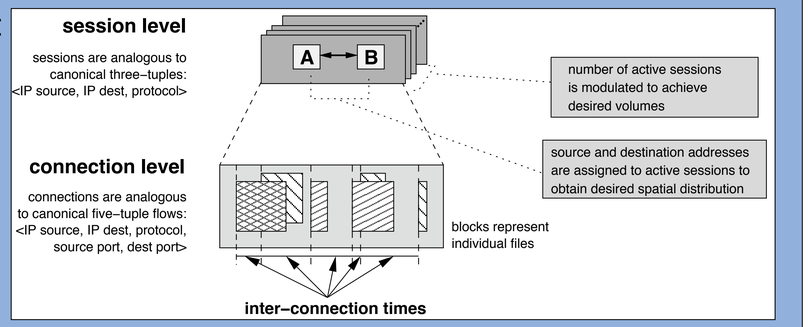

Validation of the Model and Tool

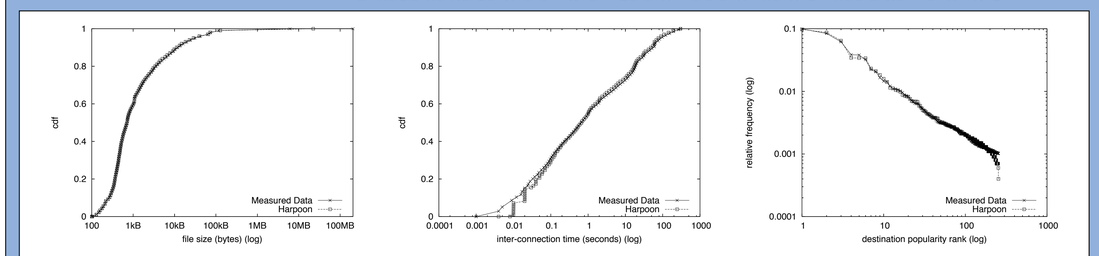

Figures show capability of Harpoon to recreate file size, inter-connection time, and spatial distributions given as input. Steps visible in test environment inter-connection times are caused by a coarse-grained operating system scheduler.

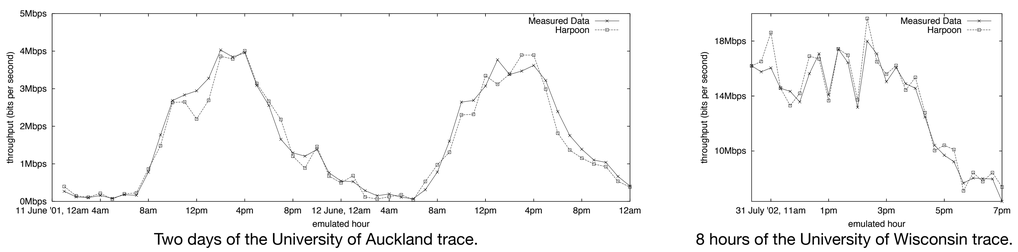

Two days of the University of Auckland trace.

8 hours of the University of Wisconsin trace.

Figures show capability of Harpoon to generate temporal volume characteristics from original data.

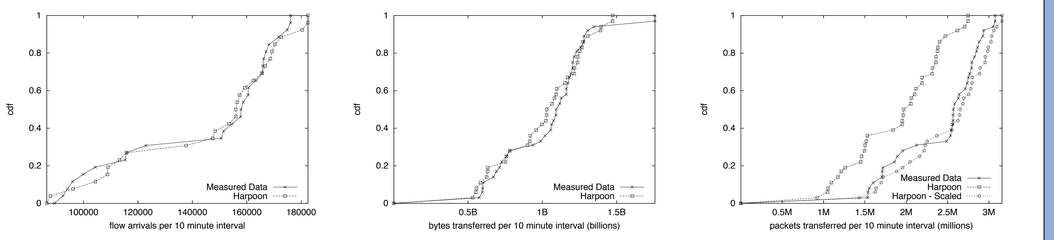

Detailed comparisons of flow, byte, and packet volumes from original data and modeled by Harpoon. For packet volumes, homogeneous maximum transmission unit sizes in the test environment lead to average packet sizes that are larger than the original data. Scaling the test environment volumes by the ratio of average packet sizes between test environment and original data yields a good match.

Using flow records captured at a border router of the University of Wisconsin and packet traces from the University of Auckland, we validated the capability of Harpoon to reproduce the distribution data supplied as input (top row) and to generate the original byte/packet/flow volumes (bottom two rows).

Comparison with Packet-Oriented Traffic Generators

We ran experiments based on RFCs 2544 and 2889 to compare loads placed upon a Cisco 6500 using Harpoon and a Spirent AX/4000, a high-performance, precise packet-oriented traffic generator.

- Two different load levels of 600Mbps and 900Mbps were used by first generating load with Harpoon, then matching average load using AX/4000.
- Configured Harpoon and the AX/4000 to generate traffic over a full class-B network, and used four different routing table sizes at the 6500.
- Used three packet size configurations and four burst size configurations at the AX/4000.
- Comparisons considered average forwarding rate at the 6500, switching fabric utilization, and packet loss rates.

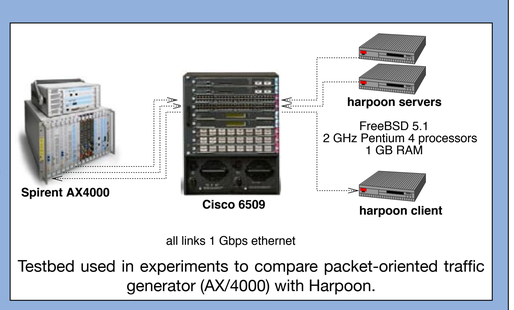

Testbed used in experiments to compare packet-oriented traffic generator (AX/4000) with Harpoon.

Major findings:

- Forwarding rates with Harpoon are much more variable than for AX/4000, especially at higher loads.
- Extreme conditions generated by AX/4000 (*e.g.*, only using 40 byte packets) do not generate the kinds of variability in forwarding rates and switching fabric utilization produced by Harpoon.

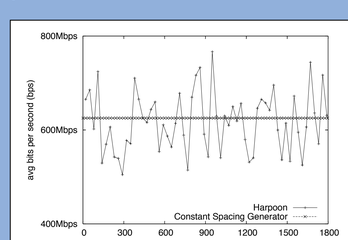

Packet forwarding rate time series for Harpoon and the AX/4000 at the low offered load level to demonstrate basic difference in load placed on a router or switch.

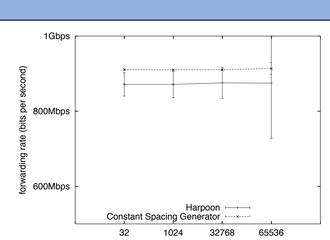

Average forwarding rates for Harpoon and the AX/4000 for different forwarding table sizes at the high offered load level. Vertical bars span one standard deviation above and below the mean. Lower average forwarding rate for Harpoon is due to packet loss at the 6500 (there is no packet loss for the AX/4000).

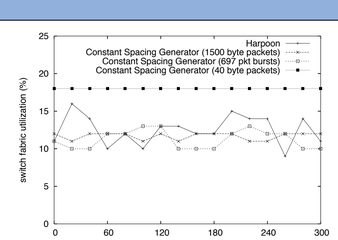

Switch fabric utilization using 2^{15} forwarding table entries and high offered load. Results shown for Harpoon, AX/4000 with uniformly spaced 1500 byte packets, bursts of 697 packets of 1500 bytes each, and uniformly spaced 40 byte packets.