


Collect Production Network Traffic

As examples, we show below two very different network traces collected at two diverse locations on the Internet used in some of our studies. The first one from UNC was taken on the border link connecting the campus to the Internet service provider network. The second trace was taken at an aggregation switch for four internal networks, connecting one of IBM Corporation's largest development sites to the Internet. Both were 1-hour traces taken at peak times.


Traffic Data for Experiments



4.7 million TCP connections carrying 290 GB of application data

2.8 million connections carrying 320 GB of application data


Round Trip Time of Connections



- UNC
 - 80 ms - average RTT
 - 36 ms - median RTT
- IBM
 - 92 ms - average RTT
 - 68 ms - median RTT

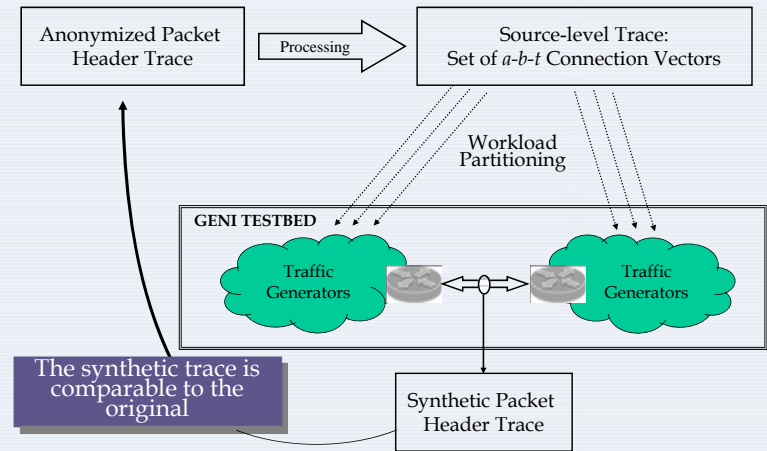
Convert TCP Traces to Connection Vectors

We analyze every connection found in a trace of TCP/IP headers to produce a "connection vector" for each connection. A connection vector includes the connection's start time relative to the beginning of the trace and a descriptor of each request-response exchange in the connection. A request-response exchange, called an "epoch," is described by a 4-tuple consisting of the request size (called the "a" unit size), the response size (called the "b" unit size) and two endpoint latency values (called the "l" values); one for the server-side "think time" between a request and its response and one for the user/client "think time" between successive requests. Unidirectional transfers have only an a or b value depending on the direction of transfer.




- | | | |
|------------|------------|------------|
| Epoch1: | Epoch2: | Epoch3: |
| > a1 bytes | > a2 bytes | > a3 bytes |
| < b1 bytes | < b2 bytes | < b3 bytes |
| tb1 ms | tb2 ms | tb3 ms |

Run Experiments on GENI



Why does this matter?

Experiment to test router queuing policy - keep state of the active connections



- Same traffic input
- Same path characteristics
- Different application models
- Drastically different outcomes!

Note: The block model emulates the same number of bytes in each direction, but omits the epoch structure and the endpoint latencies