Internet Engineering Task Force   S. Floyd
INTERNET-DRAFT   E. Kohler
draft-irtf-tmrg-tools-00.txt   Editors
Expires: March 2006   13 September 2005

          Tools for the Evaluation of Simulation and Testbed Scenarios

Status of this Memo

Copyright Notice

Abstract

   This document describes tools for the evaluation of simulation and
   testbed scenarios used in research on Internet congestion control
   mechanisms. We believe that research in congestion control

mechanisms has been seriously hampered by the lack of good models
underpinning analysis, simulation, and testbed experiments, and that
tools for the evaluation of simulation and testbed scenarios can
help in the construction of better scenarios, based on better
underlying models. One use of the tools described in this document
is in comparing key characteristics of test scenarios with known
characteristics from the diverse and ever-changing real world.
Tools characterizing the aggregate traffic on a link include the
distribution of per-packet round-trip times, the distribution of
packet sequence numbers, and the like. Tools characterizing end-to-
end paths include drop rates as a function of packet size and of
burst size, the synchronization ratio between two end-to-end TCP
flows, and the like. For each characteristic, we describe what
aspects of the scenario determine this characteristic, how the
characteristic can affect the results of simulations and experiments
for the evaluation of congestion control mechanisms, and what is
known about this characteristic in the real world. We also explain
why the use of such tools can add considerable power to our
understanding and evaluation of simulation and testbed scenarios.

1. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in [RFC 2119].

2. Introduction

This document discusses tools for the evaluation of simulation and
testbed scenarios used in research on Internet congestion control
mechanisms. These tools include but are not limited to measurement
tools; the tools discussed in this document are largely ways of
characterizing aggregate traffic on a link, or characterizing the
end-to-end path. One use of these tools is for understanding key
characteristics of test scenarios; many characteristics, such as the
distribution of per-packet round-trip times on the link, don't come
from a single input parameter but are determined by a range of
inputs. A second use of the tools is to compare key characteristics
of test scenarios with what is known of the same characteristics of
the past and current Internet, and with what can be conjectured
about these characteristics of future networks. This paper follows
the general approach from "Internet Research Needs Better Models"
[FK02].

As an example of the power of tools for characterizing scenarios, a
great deal is known about the distribution of connection sizes on a
link, or equivalently, the distribution of per-packet sequence
numbers. It has been conjectured that a heavy-tailed distribution

of connection sizes is an invariant feature of Internet traffic. A
test scenario with mostly long-lived traffic, or with a mix with
only long-lived and very short flows, does not have a realistic
distribution of connection sizes, and can give unrealistic results
in simulations or experiments evaluating congestion control
mechanisms. For instance, the distribution of packet sequence
numbers makes clear the fraction of traffic on a link from medium-
sized connections, e.g., with packet sequence numbers from 100 to
1000. These medium-sized connections can slow-start up to a large
congestion window, possibly coming to an abrupt stop soon
afterwards, contributing significantly to the burstiness of the
aggregate traffic, and to the problems facing congestion control.

In the sections below we will discuss a number of tools for
describing and evaluating scenarios, show how these characteristics
can affect the results of research on congestion control mechanisms,
and summarize what is known about these characteristics in real-
world networks.

3. Tools

The tools or characteristics that we discuss are the following.

3.1. Characterizing Aggregate Traffic on a Link

o Distribution of per-packet round-trip times.

o Distribution of packet sequence numbers.

o Distribution of packet sizes.

o Distribution of peak flow rates.

o Distribution of transport protocols.

3.2. Characterizing an End-to-End Path

o Synchronization ratio.

o Drop rates as a function of packet size.

o Drop rates as a function of burst size.

o Degree of packet drops.

o Range of queueing delay.

3.3. Other Characteristics

o Congestion control mechanisms for background or competing
   traffic.

o Characterization of congested links in terms of bandwidth and
   typical levels of congestion (in terms of packet drop rates).

o Characterization of challenging lower layers in terms of
   reordering, delay variation, packet corruption, and the like.

o Characterization of network changes affecting congestion, such as
   routing changes or link outages.

Below we will discuss each characteristic in turn, giving the
definition, the factors determining that characteristic, the effect
on congestion control metrics, and what is known so far from
measurement studies in the Internet.

4. Distribution of per-packet round-trip times

Definition: The distribution of per-packet round-trip times on a
link is defined formally by assigning to each packet the most recent
round trip time measured for that end-to-end connection. In
practice, coarse-grained information is generally sufficient, even
though it has been shown that there is significant variability in
round-trip times within a TCP connection [AKSJ03], and it is
sufficient to assign to each packet the first round-trip time
measurement for that connection, or to assign the current round-trip
time estimate maintained by the TCP connection.

Determining factors: The distribution of per-packet round-trip times
on a link is determined by end-to-end propagation delays, by
queueing delays along end-to-end paths, and by the congestion
control mechanisms used by the traffic. For example, for a scenario
using TCP, TCP connections with smaller round-trip times will
receive a proportionally larger fraction of traffic than competing
TCP connections with larger round-trip times, all else being equal,
due to the dynamics of TCP favoring flows with smaller round-trip
times. This will generally shift the distribution of per-packet
RTTs lower relative to the distribution of per-connection RTTs,
since short-RTT connections will have more packets.

Effect on congestion control metrics: The distribution of per-packet
round-trip times on a link affects the burstiness of the aggregate
traffic, and therefore can affect congestion control performance in
a range of areas such as delay/throughput tradeoffs. The

distribution of per-packet round-trip times can also affect metrics
of fairness, degree of oscillations, and the like. For example,
long-term oscillations of queueing delay are more likely to occur in
scenarios with a narrow range of round-trip times [FK02].

Measurements: The distribution of per-packet round-trip times for
TCP traffic on a link can be measured from a packet trace with the
passive TCP round-trip time estimator from Jiang and Dovrolis
[JD02]. [Add pointers to other estimators, such as ones mentioned
in JD02. Add a pointer to Mark Allman's loss detection tool.]
Their paper shows the distribution of per-packet round-trip times
for TCP packets for a number of different links. For the links
measured, the percent of packets with round-trip times at most
100 ms ranged from 30% to 80%, and the percent of packets with
round-trip times at most 200 ms ranged from 55% to 90%, depending on
the link.

In the NS simulator, the distribution of per-packet round-trip times
for TCP packets on a link can be reported by the queue monitor,
using TCP's estimated round-trip time added to packet headers. This
is illustrated in the validation test "./test-all-simple stats3" in
the directory tcl/test.

Scenarios: [FK02] shows a relatively simple scenario, with a
dumbbell topology with four access links on each end, that gives a
fairly realistic range of round-trip times. [Look for the other
citations to add.]

5. Distribution of packet sequence numbers

Definition: The distribution of packet sequence numbers on a link is
defined by giving each packet a sequence number, where the first
packet in a connection has sequence number 1, the second packet has
sequence number 2, and so on. The distribution of packet sequence
numbers can be derived in a straightforward manner from the
distribution of connection sizes, and vice versa; however, the
distribution of connection sizes is more suited for traffic
generators, and the distribution of packet sequence numbers is more
suited for measuring and illustrating the packets actually seen on a
link over a fixed interval of time. There has been a considerably
body of research over the last ten years on the heavy-tailed
distribution of connection sizes for traffic on the Internet.
[CBC95] [Add citations.]

Determining factors: The distribution of connection sizes is largely
determined by the traffic generators used in a scenario. For
example, is there a single traffic generator characterized by a
distribution of connection sizes? A mix of long-lived and web

traffic, with the web traffic characterized by a distribution of
connection sizes? Or something else?

Effect on congestion control metrics: The distribution of packet
sequence numbers affects the burstiness of aggregate traffic on a
link, thereby affecting all congestion control metrics for which
this is a factor. As an example, [FK02] illustrates that the
traffic mix can affect the queue dynamics on a congested link.
[Find more to cite, about the effect of the distribution of packet
sequence numbers on congestion control metrics.]

[Add a paragraph about the impact of medium-size flows.]

[Add a paragraph about the impact of flows starting and stopping.]

[Add a warning about scenarios that use only long-lived flows, or a
mix of long-lived and very short flows.]

Measurements: [Cite some of the literature.]

Traffic generators: Some of the available traffic generators are
listed on the web site for "Traffic Generators for Internet Traffic"
[TG]. This includes pointers to traffic generators for peer-to-peer
traffic, traffic from online games, and traffic from Distributed
Denial of Service (DDoS) attacks.

In the NS simulator, the distribution of packet sequence numbers for
TCP packets on a link can be reported by the queue monitor at a
router. This is illustrated in the validation test "./test-all-
simple stats3" in the directory tcl/test.

6. **The Distribution of Packet Sizes**

Definition: The distribution of packet sizes is defined in a
straightforward way, using packet sizes in bytes.

Determining factors: The distribution of packet sizes is determined
by the traffic mix, the path MTUs, and by the packet sizes used by
the transport-level senders.

The distribution of packet sizes on a link is also determined by the
mix of forward-path TCP traffic and reverse-path TCP traffic in that
scenario, for a scenario characterized by a `forward path' (e.g.,
left to right on a particular link) and a `reverse path' (e.g.,
right to left on the same link). For such a scenario, the forward-
path TCP traffic contributes data packets to the forward link and
acknowledgment packets to the reverse link, while the reverse-path
TCP traffic contributes small acknowledgment packets to the forward

link. The ratio between TCP data and TCP ACK packets on a link can
be used as some indication of the ratio between forward-path and
reverse-path TCP traffic.

Effect on congestion control metrics: The distribution of packet
sizes on a link is an indicator of the ratio of forward-path and
reverse-path TCP traffic in that network. The amount of reverse-
path traffic determines the loss and queueing delay experienced by
acknowledgement packets on the reverse path, significantly affecting
the burstiness of the aggregate traffic on the forward path. [Add
citations.] [In what other ways does the distribution of packet
sizes affect congestion control metrics?]

Measurements: There has been a wealth of measurements over time on
the packet size distribution of traffic [A00], [HMTG01]. These
measurements are generally consistent with a model of roughly 10% of
the TCP connections using an MSS of roughly 500 bytes, and with the
other 90% of TCP connections using an MSS of 1460 bytes.

7. The Distribution of Per-Packet Peak Flow Rates

Definition: The distribution of peak flow rates is defined by
assigning to each packet the peak sending rate in bytes per second
of that connection, where the peak sending rate is defined over
0.1-second intervals. The distribution of peak flow rates gives
some indication of the ratio of "alpha" and "beta" traffic on a
link, where alpha traffic on a congested link is defined as traffic
with that link at the main bottleneck, while the beta traffic on the
link has a primary bottleneck elsewhere along its path [RSB01].

Determining factors: The distribution of peak flow rates is
determined by flows with bottlenecks elsewhere along their end-to-
end path, e.g., flows with low-bandwidth access links. The
distribution of peak flow rates is also affected by applications
with limited sending rates.

Effect on congestion control metrics: The distribution of peak flow
rates affects the burstiness of aggregate traffic, with low-peak-
rate traffic decreasing the aggregate burstiness, and adding to the
traffic's tractability.

Measurements: [RSB01]. The distribution of peak rates can be
expected to change over time, as there is an increasing number of
high-bandwidth access links to the home, and of high-bandwidth
Ethernet links at work and at other institutions.

Simulators: [For NS, add a pointer to the DelayBox,
"http://dirt.cs.unc.edu/delaybox/", for more easily simulating low-

bandwidth access links for flows.]

8. **The Distribution of Transport Protocols.**

    Definition: The distribution of transport protocols on a congested
    link is straightforward, with each packet given its associated
    transport protocol (e.g., TCP, UDP). The distribution is often
    given both in terms of packets and in terms of bytes.

    For UDP packets, it might be more helpful to classify them in terms
    of the port number, or the assumed application (e.g., DNS, RIP,
    games, Windows Media, RealAudio, RealVideo, etc.) [MAWI]). Other
    traffic includes ICMP, IPSEC, and the like. In the future there
    could be traffic from SCTP, DCCP, or from other transport protocols.

    Effect on congestion control metrics: The distribution of transport
    protocols affects metrics relating to the effectiveness of AQM
    mechanisms on a link.

    Measurements: In the past, TCP traffic has typically consisted of
    90% to 95% of the bytes on a link [UW02], [UA01]. [Get updated
    citations for this.] Measurement studies show that TCP traffic from
    web servers almost always uses conformant TCP congestion control
    procedures [MAF05].

9. **The Synchronization Ratio**

    Definition: The synchronization ratio is defined as the degree of
    synchronization of loss events between two TCP flows on the same
    path. Thus, the synchronization ratio is defined as a
    characteristic of an end-to-end path. When one TCP flow of a pair
    has a loss event, the synchronization ratio is given by the fraction
    of those loss events for which the second flow has a loss event
    within one round-trip time. Each connection in a flow pair has a
    separate synchronization ratio, and the overall synchronization
    ratio of the pair of flows is the higher of the two ratios. When
    measuring the synchronization ratio, it is preferable to start the
    two TCP flows at slightly different times, with large receive
    windows.

    Determining factors: The synchronization ratio is determined largely
    by the traffic mix on the congested link, and by the AQM mechanism
    (or lack of AQM mechanism).

    Different types of TCP flows are also likely to have different
    synchronization measures. E.g., Two HighSpeed TCP flows might have
    higher synchronization measures that two Standard TCP flows on the
    same path, because of their more aggressive window increase rates.

Effect on congestion control metrics: The synchronization ratio
affects convergence times for high-bandwidth TCPs. Convergence
times are known to be poor for some high-bandwidth protocols in
environments with high levels of synchronization. [Cite the papers
by Leith and Shorten.] However, it is not clear if these
environments with high levels of synchronization are realistic.

Measurements: Grenville Armitage and Qiang Fu have performed initial
experiments of synchronization in the Internet, using Standard TCP
flows, and have found very low levels of synchronization.

[Appenzeller et al, Sizing Router Buffers, reports that
synchronization is rare as the number of competing flows increases.
Kevin Jeffay has some results on synchronization also.]

Needed: We need measurements of the synchronization ratio for flows
that use high-bandwidth protocols over high-bandwidth paths, given
typical levels of competing traffic and with typical queuing
mechanisms at routers (whatever these are), to see if there are
higher levels of synchronization with high-bandwidth protocols such
as HighSpeed TCP, Fast TCP, and the like, which are more aggressive
than Standard TCP. The assumption would be that in many
environments, high-bandwidth protocols have higher levels of
synchronization than flows using Standard TCP.

10. **Drop or Mark Rates as a Function of Packet Size**

Definition: Drop rates as a function of packet size are defined by
the actual drop rates for different packets on an end-to-end path or
on a congested link over a particular time interval. In some cases,
e.g., Drop-Tail queues in units of packets, general statements can
be made; e.g., that large and small packets will experience the same
packet drop rates. However, in other cases, e.g., Drop-Tail queues
in units of bytes, no such general statement can be made, and the
drop rate as a function of packet size will be determined in part by
the traffic mix at the congested link at that point of time.

Determining factors: The drop rate as a function of packet size is
determined in part by the queue architecture. E.g., is the Drop-
Tail queue in units of packets, of bytes, of 60-byte buffers, or of
a mix of buffer sizes? Is the AQM mechanism in packet mode,
dropping each packet with the same probability, or in byte mode,
with the probability of dropping or marking a packet being
proportional to the packet size in bytes.

The effect of packet size on drop rate would also be affected by the
presence of preferential scheduling for small packets, or by
differential scheduling for packets from different flows (e.g., per-

flow scheduling, or differential scheduling for UDP and TCP
traffic).

In many environments, the drop rate as a function of packet size
will be heavily affected by the traffic mix at a particular time.
For example, is the traffic mix dominated by large packets, or by
smaller ones? In some cases, the overall packet drop rate could
also affect the relative drop rates for different packet sizes.

Effect on congestion control metrics: The drop rate as a function of
packet size has a significant effect on the performance of
congestion control for VoIP and other small-packet flows.
[Citation: "TFRC for Voice: the VoIP Variant", draft-ietf-dccp-tfrc-
voip-02.txt, and earlier papers.] The drop rate as a function of
packet size also has an effect on TCP performance, as it affects the
drop rates for TCP's SYN and ACK packets. [Citation: Jeffay and
others.]

Measurements: We need measurements of the drop rate as a function of
packet size over a wide range of paths, or for a wide range of
congested links. For tests of relative drop rates on end-to-end
packets, one possibility would be to run successive TCP connections
with 200-byte, 512-byte, and 1460-byte packets, and to compare the
packet drop rates. The ideal test would include running TCP
connections on the reverse path, to measure the drop rates for the
small ACK packets on the forward path. It would also be useful to
characterize the difference in drop rates for 200-byte TCP packets
and 200-byte UDP packets, even though some of this difference could
be due to the relative burstiness of the different connections.

Ping experiments could also be used to get measurements of drop
rates as a function size, but it would be necessary to make sure
that the ping sending rates were adjusted to be TCP-friendly.

[Cite the known literature on drop rates as a function of packet
size.]

Our conjecture is that there is a wide range of behaviors for this
characteristic in the real world. Routers include Drop-Tail queues
in packets, bytes, and buffer sizes in between; these will have
quite different drop rates as a function of packet size. Some
routers include RED in byte mode (the default for RED in Linux) and
some have RED in packet mode (Cisco, I believe). This also affects
drop rates as a function of packet size.

Some routers on congested access links use per-flow scheduling. In
this case, does the per-flow scheduling have the goal of fairness in
bytes per second or in *packets* per second? What effect does the

per-flow scheduling have on the drop rate as a function of packet
size, for packets in different flows (e.g., a small-packet VoIP flow
competing against a large-packet TCP flow) or for packets within the
same flow (small ACK packets and large data packets on a two-way TCP
connection).

11. Drop Rates as a Function of Burst Size.

Definition: Burst-tolerance, or drop rates as a function of burst
size, can be defined in terms of an end-to-end path, or in terms of
aggregate traffic on a congested link.

The burst-tolerance of an end-to-end path is defined in terms of
connections with different degrees of burstiness within a round-trip
time. When packets are sent in bursts of N packets, does the drop
rate vary as a function of N? For example, if the TCP sender sends
small bursts of K packets, for K less than the congestion window,
how does the size of K affect the loss rate? Similarly, for a ping
tool sending pings at a certain rate in packets per second, one
could see how the clustering of the ping packets in clusters of size
K affects the packet drop rate. As always with such ping
experiments, it would be important to adjust the sending rate to
maintain a longer-term sending rate that was TCP-friendly.

Determining factors: The burst-tolerance is determined largely by
the AQM mechanisms for the congested routers on a path, and by the
traffic mix. For a Drop-Tail queue with only a small number of
competing flows, the burst-tolerance is likely to be low, and for
AQM mechanisms where the packet drop rate is a function of the
average queue size rather than the instantaneous queue size, the
burst tolerance should be quite high.

Effect on conjestion control metrics: The burst-tolerance of the
path or congested link can affect fairness between competing flows
with different round-trip times; for example, Standard TCP flows
with longer round-trip times are likely to have a more bursty
arrival pattern at the congested link that that of Standard TCP
flows with shorter round-trip times. As a result, in environment
with low burst tolerance (e.g., scenarios with Drop-Tail queues),
longer-round-trip-time TCP connections can see higher packet drop
rates than other TCP connections, and receive an even smaller
fraction of the link bandwidth than they would otherwise.
[FJ92](Section 3.2). We note that some TCP traffic is inherently
bursty, e.g., Standard TCP without rate-based pacing, particularly
in the presence of dropped ACK packets or of ACK compression. The
burst-tolerance of a router can also affect the delay-throughput
tradeoffs and packet drop rates of the path or of the congested
link.

Measurements: One could measure the burst-tolerance of an end-to-end
path by running successive TCP connections, forcing bursts of size
at least K by dropping an appropriate fraction of the ACK packets to
the TCP receiver. Alternately, if one had control of the TCP
sender, one could modify the TCP sender to send bursts of K packets
when the congestion window was K or more segments.

[Look at Crovella's paper on loss pairs.]

[Look at: M. Allman and E. Blanton, "Notes on Burst Mitigation for
Transport Protocols", ACM Computer Communication Review, vol. 35(2),
(2005).]

Making inferences about the AQM mechanism for the congested router
on an end-to-end path: One potential use of measurement tools for
determining the burst-tolerance of an end-to-end path would be to
make inferences about the presence or absence of an AQM mechanism at
the congested link or links. As a simple test, one could run a TCP
connection until the connection comes out of slow-start. If the
receive window of the TCP connection was sufficiently high that the
connection exited slow-start with packet drops or marks instead of
because of the limitation of the receive window, one could record
the congestion window at the end of slow-start, and the number of
packets dropped from this window. A high packet drop rate might be
more typical of a Drop-Tail queue with small-scale statistical
multiplexing on the congested link, and a single packet drop coming
out of slow-start would suggest an AQM mechanism at the congested
link.

The synchronization measure could also add information about the
likely presence or absence of AQM on the congested link(s) of an
end-to-end path, with paths with higher levels of synchronization
being more likely to have Drop-Tail queues with small-scale
statistical multiplexing on the congested link(s).

[Cite the relevant literature about tools for determining the AQM
mechanism on an end-to-end path.]

12. Congestion Control Mechanisms for Background or Competing Traffic

Effect on congestion control metrics: Please don't evaluate AQM
mechanisms by using Reno TCP, or evaluate new transport protocols by
comparing them with the performance of Reno TCP! For an
explanation, see [FK02](Section 3.4).

Measurements: See [MAF05].

[13](#). **Characterization of Congested Links in Terms of Bandwidth and**
Typical Levels of Congestion

    [Pointers to the current state of our knowledge.]

[14](#). **Characterization of Challenging Lower Layers.**

    [This will just be a short set of pointers to the relevant
literature, which is quite extensive.]

[15](#). **Characterization of Network Changes Affecting Congestion**

    [Pointers to the current state of our knowledge.]

[16](#). **Using the Tools Presented in this Document**

    [To be done.]

[17](#). **Related Work**

    [Cite "On the Effective Evaluation of TCP" by Allman and Falk.]

[18](#). **Conclusions**

    [To be done.]

[19](#). **Security Considerations**

    There are no security considerations in this document.

[20](#). **IANA Considerations**

    There are no IANA considerations in this document.

[21](#). **Acknowledgements**

Informative References

    [RFC 2119] S. Bradner. Key Words For Use in RFCs to Indicate
        Requirement Levels. [RFC 2119](#).

    [MAWI] M.W. Group, Mawi working group traffic archive, URL
        "http://tracer.csl.sony.jp/mawi/".

    [AKSJ03] J. Aikat, J. Kaur, F.D. Smith, and K. Jeffay, Variability
        in TCP Roundtrip Times, ACM SIGCOMM Internet Measurement
        Conference, Maimi, FL, October 2003, pp. 279-284.

     [A00] M. Allman, A Web Server's View of the Transport Layer,
          Computer Communication Review, 30(5), October 200.

     [CBC95] C. Cunha, A. Bestavros, and M. Crovella, "Characteristics of
          WWW Client-based Traces", BU Technical Report BUCS-95-010, 1995.

     [FJ92] S. Floyd and V. Jacobson, On Traffic Phase Effects in Packet-
          Switched Gateways, Internetworking: Research and Experience, V.3
          N.3, September 1992, p.115-156.

     [FK02] S. Floyd and E. Kohler, Internet Research Needs Better
          Models, Hotnets-I, October 2002.

     [HMTG01] C. Hollot, V. Misra, D. Towsley, and W. Gong, On Designing
          Improved Controllers for AQM Routers Supporting TCP Flows, IEEE
          Infocom, 2001.

     [JD02] H. Jiang and C. Dovrolis, Passive Estimation of TCP Round-
          trip Times, Computer Communication Review, 32(3), July 2002.

     [MAF05] A. Medina, M. Allman, and A. Floyd. Measuring the Evolution
          of Transport Protocols in the Internet. Computer Communication
          Review, April 2005.

     [RSB01] R. Riedi, S. Sarvotham, and R. Varaniuk, Connection-level
          Analysis and Modeling of Network Traffic, SIGCOMM Internet
          Measurement Workshop, 2001.

     [TG] Traffic Generators for Internet Traffic Web Page, URL
          "http://www.icir.org/models/trafficgenerators.html".

     [UA01] U. of Auckland, Auckland-vi trace data, June 2001. URL
          "http://wans.cs.waikato.ac.nz/wand/wits/auck/6/".

     [UW02] UW-Madison, Network Performance Statistics, October 2002.
          URL "http://wwwstats.net.wisc.edu/".

Editors' Addresses

     Sally Floyd <floyd@icir.org>
     ICSI Center for Internet Research
     1947 Center Street, Suite 600
     Berkeley, CA 94704
     USA

     Eddie Kohler <kohler@cs.ucla.edu>
     4531C Boelter Hall
     UCLA Computer Science Department

      Los Angeles, CA 90095
      USA

Full Copyright Statement

Intellectual Property