CoMPACT-Monitor: Change-of-Measure based Passive/Active Monitoring
— Weighted Active Sampling Scheme to Infer QoS —

Masaki Aida, Keisuke Ishibashi, and Toshiyuki Kanazawa
NTT Information Sharing Platform Laboratories, NTT Corporation
3-9-11 Midori-cho, Musashino-shi 180-8585, Japan
`aida.masaki@lab.ntt.co.jp`

Abstract

This paper proposes a new performance measurement scheme which can infer performance and/or quality experienced by individual user, organization and application, via a scalable and lightweight measurement way. The proposed scheme is based on change-of-measure framework and is an active measurement transformed by using passively monitored data. We give the theoretical basis of the proposed scheme and show its typical implementations for inferring the delay distributions. The validation of the implementations of the proposed scheme is investigated through simulation with respect to both the accuracy of estimation and the amount of the extra traffic added by active measurement.

1. Introduction

The Internet has been growing rapidly with respect to the number of users and the amount of traffic and has been recognized as an important lifeline for information in social and business use. So, although initial and basic interest of the Internet has been its connectivity and transmission capacity, recent attention is also paid to its quality. The traffic conveyed by the Internet is generated by wide variety applications, which have different characteristics and have different quality requirements. Thus, Quality of Service (QoS) and performance measurements are crucially important in controlling and managing QoS and in provisioning networks.

In general, monitoring schemes to measure QoS and the performance of networks are, divided into two types, active and passive monitoring. Unfortunately, both types have drawbacks. They are briefly summarized as follows.

Active measurement monitors QoS and the performance of a network by sending probe packets and monitoring them. Many active monitoring tools have been developed to monitor network performance [1]. They monitor the performance of the probe packets to determine the performance of the network indirectly. This means that we implicitly assume the QoS/performance of a user/networks is identical with the values measured from active probe packets. There are following problems in these active monitoring schemes.

- If we use a probe-packet stream that simulates an actual user traffic,
  - The probe-packet stream incurs non-negligible extra traffic into the networks and it affects QoS/performance of users' traffic, and
  - The QoS/performance obtained from the probe packets is not equal to the unbiased one without influences of the probe-packet stream.
- If we use probe packets that have small length and are sent periodically, like ping,
  - The extra traffic may be negligible, but the QoS/performance obtained from the probe packets is not equal to the QoS/performance experienced by users, in general.

We add some explanation about the last case. Since the time for sending probe packets is independent of the users' behaviors, QoS/performance measured by the active monitoring scheme generally differs from the actual QoS/performance that users experience. If and only if we can assume that active monitoring measures the time average of network performance and that the user traffic is Poisson, then the performance experienced by the users and the actively-measured performance will be the same. This well-known property is called PASTA (which stands for "Poisson Arrivals See Time Average"). It is known, however, that current Internet traffic exhibits bursty properties and is not generally Poisson [2]. In that case, an average user experiences worse performance than the time-average performance measured by active monitoring.

Proceedings of the 2002 Symposium on Applications and the Internet (SAINT'02w)
0-7695-1450-2/02 $17.00 © 2002 **IEEE**

Incidentally, the relationship between the network performance experienced by a user and the measured performance at an arbitrary time can be obtained with the Palm measure [4]. However, to obtain the distribution of performance experienced by users, we require information about the distributions of interarrival time and service time to calculate the Palm measure. It is hard to measure these distributions, so this way is not feasible.

Passive measurement is mainly used to monitor traffic volume but can measure network performance as well. Passive monitoring is categorized into two types, two-point monitoring and one-point monitoring.

- Two-point monitoring requires two monitoring devices deployed at ingress and egress points in a network. The devices sequentially take packets' data, and network performance parameters such as delay and loss can be calculated by comparing the data of the corresponding packets taken at each point. If we apply the two-point monitoring to measure QoS/performance,
  - The all devices should have synchronized timing, and
  - The two-point monitoring requires identification of each packet at the two devices by its header or contents. Since this identification process is hard when the packet volume is huge as in a large-scale network, the two-point monitoring does not have scalablility.
- One-point monitoring uses the TCP acknowledgment mechanism. When a TCP-sink receives a packet from a TCP-source, the TCP-sink transmits the acknowledgement for the packet [3]. Thus, by monitoring the packet-ack pair at a point in the network, the round-trip delay between the point and the sink can be measured. The packet loss can also be detected in this way. If we apply the one-point monitoring, measurement is restricted for TCP flows.

In this paper, we propose a new performance measurement scheme to estimate the actual network performance experienced by users. Our scheme requires both active and passive monitoring using easy-to-measure methods. It is based on change-of-measure framework and is an active measurement transformed by using passively monitored data. Our scheme can estimate not only the mixed QoS/performance experienced by users but also the actual QoS/performance for individual users, organizations, and applications. In addition, the proposed scheme is scalable and lightweight.

The rest of the paper is organized as follows. In section 2, we give a mathematical formalization of the framework of our scheme. In section 3, as an application of the scheme, we propose a simple scheme for estimating the actual delay experienced by users, which is easy to implement. In addition, we show the validity of the scheme through simulation. In section 4, we extend the proposed scheme to estimate the performance experienced by an individual user. Finally, we conclude the paper in section 5.

2. Proposed Measurement Scheme

This section proposes a new measurement scheme based on a change-of-measure framework. The proposed scheme is scalable and lightweight and enables accurate estimation of detailed characteristics of performance for individual users, organizations, and applications. In this scheme, a combination of simple measurements of both the active and passive types enables the change-of-measure framework.

2.1. Change-of-Measure Based Measurement Scheme

We can recognize that almost all measurements fundamentally correspond to the integrals. This is because it is a cumulation of some quantities according to a certain rule. Our scheme enables us to obtain the measurement objective not from the integral describing a direct measurement of the objective, but from other integral that is easy-to-measure. These integrals are in different forms but their values are the same.

Let X be the measurement objective, *e.g.,* the delay for user packets, whose distribution function P is as follows:

$$\begin{aligned} \Pr(X > a) &= \int \mathbf{1}_{\{x>a\}} \, \mathrm{d}P(x) \\ &= \mathrm{E}_p \left[\mathbf{1}_{\{X>a\}} \right], \end{aligned} \tag{1}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. We consider how to estimate the distribution P via measurements of X. Suppose there are n arrivals in a measurement period, *e.g.,* n packets arrived. $X(i)$ denotes the i-th value of X. Then an estimator $Z_x(n,a)$ of the distribution P of X can be obtained by using $X(i)$ as follows:

$$Z_x(n,a) := \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\{X(i)>a\}}. \tag{2}$$

The estimator $Z_x(n,a)$ satisfies unbiasedness

$$\mathrm{E}_p \left[Z_x(n,a) \right] = \Pr(X > a). \tag{3}$$

In addition, if $X(i)$ is ergodic, then $\mathbf{1}_{\{X(i)>a\}}$ is also ergodic for arbitrary $a \in \mathbf{R}$, and the estimator $Z_x(n,a)$ strongly converges to

$$\lim_{n\to\infty} Z_x(n,a) = \Pr(X > a) \text{ a.s.} \tag{4}$$

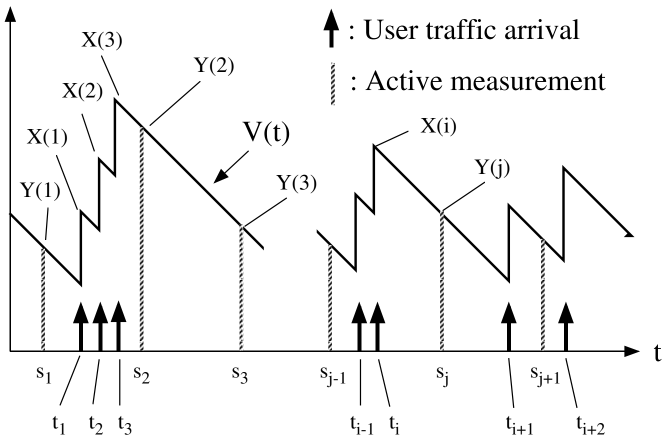

Figure 1. Relationship between $V(t)$ and $X(i)$, $Y(j)$.

Suppose we have a situation in which it is difficult to measure $X(i)$ directly, and an estimate of its distribution cannot be obtained with (2). Let $V(t)$ be the network performance at time t such that if the i-th arrival occurs at t_i, then $V(t_i) = X(i)$. In the example of delay measurements, $V(t)$ is the virtual weighting time. Also, let Y be the value of $V(t)$ measured independently of the arrivals $X(i)$ and let the distribution function of Y be Q. We assume that for $a, b \in \mathbf{R}$,

$$P(b) - P(a) > 0 \ \Rightarrow Q(b) - Q(a) > 0. \tag{5}$$

This assumption indicates that the network performance for arrivals can be measured with a positive probability. This is natural when the measurement lasts long enough. We can then define $\mathrm{d}P/\mathrm{d}Q$, and the distribution of X can be written using the distribution of Y as

$$\begin{aligned}\Pr(X > a) &= \int \mathbf{1}_{\{y>a\}} \frac{\mathrm{d}P(y)}{\mathrm{d}Q(y)} \,\mathrm{d}Q(y) \\ &= \mathrm{E}_q \left[\mathbf{1}_{\{Y>a\}} \frac{\mathrm{d}P(Y)}{\mathrm{d}Q(Y)} \right], \end{aligned} \tag{6}$$

by (1). Now, suppose $V(t)$ is measured m times, and let $Y(j)$ be the j-th measurement at s_j such that $Y(j) = V(s_j)$ $(j = 1, 2, \ldots, m)$ (Fig. 1). Then an estimator of $\Pr(X > a)$ can be derived with $Y(j)$ as follows:

$$Z_y(m, a) := \frac{1}{m} \sum_{j=1}^{m} \mathbf{1}_{\{Y(j)>a\}} \, L(j), \tag{7}$$

where

$$L(j) := \frac{\mathrm{d}P(Y(j))}{\mathrm{d}Q(Y(j))}, \tag{8}$$

which we call the likelihood ratio [5]. Equations (3) and (4) also hold for $Z_y(m, a)$ as

$$\mathrm{E}_q \left[Z_y(m, a) \right] = \Pr(X > a), \text{ and} \tag{9}$$

$$\lim_{m \to \infty} Z_y(m, a) = \Pr(X > a) \text{ a.s.} \tag{10}$$

If we can derive $L(j)$, then the estimator of the distribution P of X can be derived with the measurement value Y. The fundamental concept of our scheme is as follows: The estimation of the distribution P from direct measurements of X is difficult. However, since the values $Y(j)$ and $L(j)$ are easily obtained by respectively using active and passive monitoring, we can easily estimate the distribution P by using (7). The derivation of the likelihood ratio is described in the next subsection.

2.2. Likelihood Ratio

Let $\rho_x(t)$ be the traffic volume arriving (*e.g.,* the number of packets arriving) in an interval $[t, \, t + \delta(t))$. Let $\rho_y(t)$ be the number of measurements in $[t, \, t + \delta(t))$. We assume that the interval $\delta(t)$ is short enough compared to the time variance of $V(t)$ so that

$$V(s) \simeq V(s') \text{ for } \forall s, s' \in [t, \, t + \delta(t)). \tag{11}$$

The assumption indicates that one measurement of Y in the interval $[t, \, t + \delta(t))$ can be interpreted as $\rho_x(t)/\rho_y(t)$ measurements of X. Note that we can always define $\rho_x(s_j)/\rho_y(s_j)$ because at the time s_j of measurement $Y(j)$, $\rho_y(s_j) > 0$. Then the likelihood ratio can be obtained as

$$L(j) = \frac{\rho_x(s_j)}{\rho_y(s_j)} \frac{m}{n}, \tag{12}$$

because the total number of events of X is n and that of Y is m.

The likelihood ratio (12) can be obtained with passive measurements, and the distribution of X is estimated as

$$Z_y(m, a) = \frac{1}{n} \sum_{j=1}^{m} \mathbf{1}_{\{Y(j)>a\}} \frac{\rho_x(s_j)}{\rho_y(s_j)}. \tag{13}$$

2.3. Strength of Our Scheme

We can expect the proposed scheme has the following superior points:

- Since the extra traffic for active probe packets is negligible, users' traffic is little affected.
- We have a dependable estimation of QoS/performance measure.
- Since passive measurement is only required counting the amount of traffic (the number of packets), the two-point measurement is not necessary and the passive monitoring devices are simplified.
- We can apply our scheme to non-TCP protocols as well.

3. Simple Application for the Delay Measurement

As an application of the proposed scheme described in the previous section, we propose a simple scheme for estimating the actual delay experienced by user which is easy to implement.

3.1. Formalization

Let $Y(j)$ ($j = 1, 2, \ldots, m$) be the delay measured with probe packets, such as ping or other active monitoring, at time s_j. The probe packet interval $s_{j+1} - s_j$ is chosen to be a constant τ and $\delta(s_j)$ is also chosen to be the same interval, as a simple implementation.[1] Then suppose that the number of user packets arriving in $[s_j, s_{j+1})$ is $\rho(j)$ and the total number of packets arriving in the measurement period is $\sum_{j=1}^{m} \rho(j) = n$. As an example of a delay estimation case, $V(t)$ is considered as the virtual waiting time of the network, which is the delay for a packet arriving (virtually) at t. If we assume that τ is short enough compared to the fluctuation of $V(t)$, then we derive the estimator of the packet-delay distribution by applying (13) for $\rho_x(s_j) = \rho(j)$, $\rho_y(s_j) = 1$ and $\delta(s_j) = \tau$ as

$$Z_y(m, a) = \frac{1}{n} \sum_{j=1}^{m} \mathbf{1}_{\{Y(j)>a\}}\, \rho(j). \qquad (14)$$

As can be seen from (14), estimating the user delay requires measuring the network delay periodically with active probe packets and measuring the number of packets arriving between active measurements, which is far easier than measuring the delay for user packets directly with two probes deployed at the network edges.

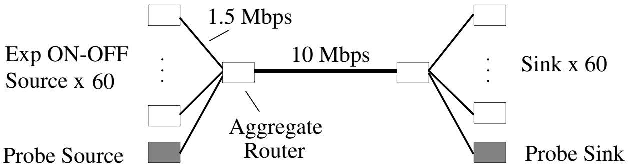

Figure 2. Network configuration for simulation.

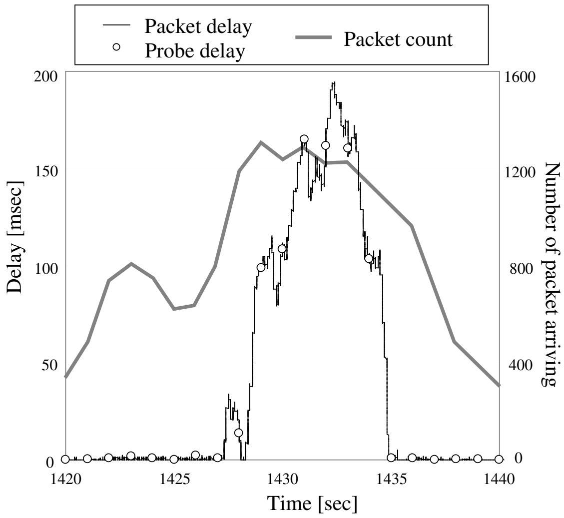

Figure 3. Sample paths for the number of packets arriving, user packet delay, and probe packet delay.

3.2. Evaluation

To demonstrate our simple application described above, we use the *ns2* [6] network simulator. Figure 2 shows the network topology for the simulation. Sixty sources are connected to a bottleneck router with 1.5-Mbps links, and two routers are connected with a 10-Mbps link.

We measure the queueing delay at the bottleneck router, which does not include the service time for the packets themselves. Other simulation conditions are as follows:

- The user packets are generated by ON-OFF sources. Both the ON and OFF durations are distributed as i.i.d. exponentials, where the mean ON duration is 1 second and the mean OFF duration is 14 seconds. The packet size is fixed at 1000 bytes.
- As the transport protocol for the user packets, both TCP and UDP are evaluated.

---

[1] It prefers that τ is exponentially distributed values because τ should not synchronize any traffic patterns in the networks.

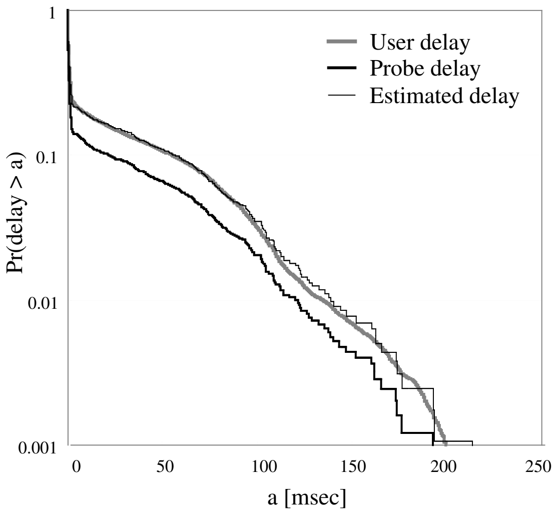

Figure 4. Distributions of queueing delay for packets generated by TCP On-Off sources, probe packets, and estimator.

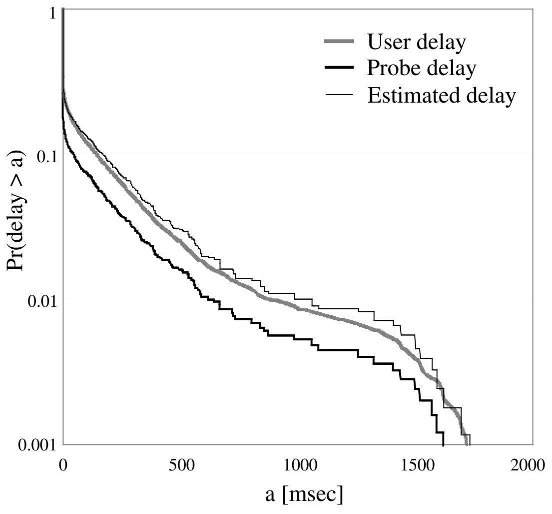

Figure 5. Distributions of queueing delay for packets generated by UDP On-Off sources, probe packets, and estimator.

- Probe packets to actively measure the queueing delay are generated every second. The size of each probe packet is fixed at 64 bytes.
- The bottleneck link utilization is 0.6.
- Simulation time is 2,500 sec, *i.e.*, 2,500 active probe packets are sent.

First, we show results for TCP. Figure 3 shows a sample path for the user packet delay, probe packet delay, and number of user packets arriving between probe packets. It can be seen that the delay measured with probe packets captures well the time variance of delay for the user packets. We can also, however, see fluctuation in the number of packets, which synchronizes with the delay fluctuation. This fluctuation causes the discrepancy between the distribution of delay for bursty user packets and periodical probe packets, because the number of packets with worse delay is larger for user packets than probe packets.

Figure 4 shows the delay distributions of user packets and probe packets and an estimation, for TCP. As expected from the sample path, we can observe the discrepancy between the distribution of user packet delay and that for active measurements. Using our proposed scheme, however, user delay can be estimated with high accuracy for active measurements.

Figure 5 shows results for UDP. Although the delay grows larger than for TCP because packets are sent without congestion control, it can be seen that our scheme estimates the user packet delay with high accuracy. In the rest of this paper, we use TCP as the transport protocol for user packets.

In the above examples, the extra traffic for active probe packets is 512 bps. This is about 0.005% of the bandwidth and is negligible.

4. Simple Application for the Individual User Delay Measurement

We describe here a proposed simple application using our change-of-measure-based measurement scheme, which can be extended to estimate the packet delay for individual users with one series of active measurements and passive traffic monitoring. user.

4.1. Formalization

Let X_k be the packet delay of user k $(k = 1, 2, \ldots, K)$ and $Y(j)$ $(j = 1, 2, \ldots, m)$ be the delay measured with active packets, such as ping, at time s_j. Assume the interval of probe packets $s_{j+1} - s_j$ is a constant τ, the number of packets for user k arriving in $[s_j, s_{j+1})$ is $\rho_k(j)$, and the number of total packets for user k is

$$n_k := \sum_{j=1}^{m} \rho_k(j). \tag{15}$$

Then the likelihood ratio for user k is

$$L_k(j) := \rho_k(j)\, \frac{m}{n_k}, \tag{16}$$

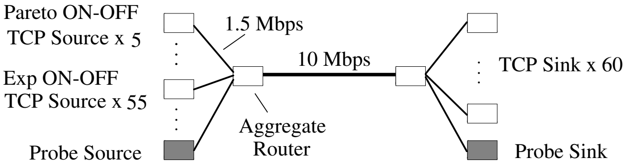

Figure 6. Network model.

and we can obtain the estimator as follows:

$$Z_y(k,m,a) = \frac{1}{n_k} \sum_{j=1}^{m} \mathbf{1}_{\{Y(j)>a\}}\, \rho_k(j). \qquad (17)$$

Thus, by counting the number of packets arriving for each user every τ seconds, we can estimate the delay experienced by individual users.

The classification of traffic is not limited to individual users or groups of users. An typical other example is the classification of traffic by by kind of applications. The traffic pattern generated by an application depends on the application, so the performance for packets may also depend on the application generating the packets. Using our scheme, we can monitor the performance for each application with one series of active measurements, on the condition that packets for every class are treated with the same priority in the network.

4.2. Evaluation

We next describe the proposed scheme to estimate the performance experienced separately by two groups of users. We separate the 60 sources from the simulation run in subsection 3.2 in 5 bursty sources and 55 non-bursty sources (Fig. 6). For the bursty sources, the ON-OFF durations are distributed in a Pareto distribution with a shape parameter of 1.5, where the mean ON duration is 1 second and the mean OFF duration is 14 seconds. For the non-bursty sources, the ON-OFF durations are distributed in a exponential, where the mean ON duration is 10 seconds and the mean OFF duration is 5 seconds.

Simulation time is 25,000 sec, *i.e.*, the number of generated probe packets is 25,000. The other parameters are the same as before.

Figure 7 shows the distribution of the user packet delay and the estimated distribution using (17). We see that the proposed scheme can estimate the distributions of both groups of users with high accuracy for one series of active measurements.

The extra traffic for active probe packets is 512 bps and this is about only 0.005% of the bandwidth.

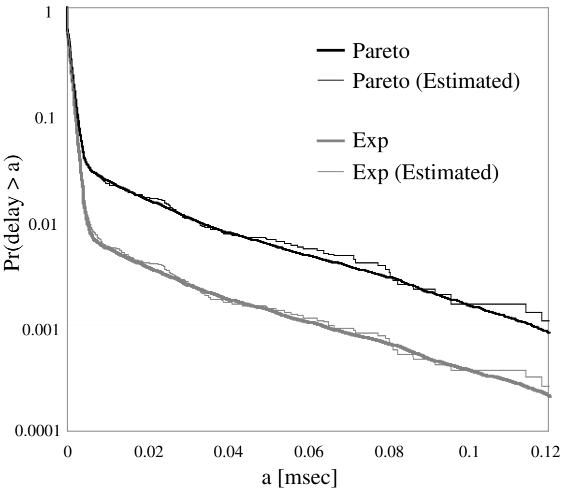

Figure 7. Estimations of delay distributions for two types of users.

5 Conclusion

In this paper, we have proposed a new performance measurement scheme, *the change-of-measure based measurement scheme*, which can estimate detailed QoS/performance via a scalable and lightweight method. Our scheme only requires counting the traffic volume as passive monitoring and simple measurement of network performance as active monitoring, which is feasible and tractable compared to a conventional scheme. We have applied our change-of-measure scheme as an example of a simple implementation. We used simulation to validate the proposed application. As a result, Our scheme has been shown to give a good estimation of the performance seen by a user. The extra traffic for active probe packet is negligible. We extended the scheme to estimate individual user performance, and confirmed the validity of this approach by simulation.

We should evaluate the relationship between measurement interval of active probe packets and the accuracy of the estimation. The result will address how to measure QoS/performance in various situation and will clarify the strength of our scheme in the practical implementations. These issues are for further study.

References

[1] CAIDA cooperative association for internet data analysis.
`http://www.caida.org/tools/.`

[2] V. Paxon and S. Floyd, "Wide-area traffic: The failure of Poisson modeling," IEEE/ACM Trans. on Networking, June 1995.

[3] W. R. Stevens, "TCP/IP illustrated, Vol. 1," Addison-Wesley, Reading MA, 1994.

[4] F. Baccelli and P. Brémaud, "Palm probabilities and stationary queues," Springer-Verlag, 1980.

[5] P. J. Smith, M. Shafi, and H. Gao, "Quick simulation: a review of importance sampling techniques in communications systems," IEEE J. Selected Areas in Commun., vol. 15, No. 4, May 1997.

[6] UCB/LBNL/VINT Network Simulator - ns (version 2). `http://www.isi.edu/nsnam/ns`.

[7] K. Thompson, G. J. Miller, and R. Wilder, "Wide-area Internet traffic patterns and characteristics," IEEE Network, November/December 1997.