

Stochastic Sampling for Internet Traffic Measurement

Tilman Wolf, Yan Cai, Patrick Kelly, and Weibo Gong
Department of Electrical and Computer Engineering
University of Massachusetts, Amherst, MA, USA
{wolf,ycai,kelly,gong}@ece.umass.edu

Abstract—The increasing complexity of the Internet demands continued improvements to measurement techniques and data analysis methods to aid our understanding of network operation. The availability of accurate measurement data is necessary in many areas ranging from attack detection, novel pricing schemes, buffer dimensioning and switch design to general network management. In this paper, we develop a theory for accurate and unbiased Internet traffic measurement using the tools of Poisson random sampling. We show how this approach helps in storing, managing, and aggregating data from different sources with independent clocks and sampling rates. We present results that show that stochastic sampling maintains important information about network measurements that would be lost when using conventional uniform sampling.

I. INTRODUCTION

Accurate Internet traffic measurements are critical for understanding and effectively managing network operations. Traffic data also is a vital source of information for planning, billing, anomaly and intrusion detection, and developing new network designs. Ever increasing data link rates to support traffic demands make it virtually impossible to record and store information on every packet that is transmitted. Instead, the common current practice for traffic measurement is to sample or aggregate information. For example, to estimate link utilization, the size of all packets is added and averaged over a fixed-length interval. Estimation over fixed intervals essentially amounts to *uniform sampling* of the integrated traffic.

The inherent problem with this approach is that the observed arrival process on a high data rate link inevitably contains extremely high frequencies. It is well known that uniform sampling of a stochastic process whose bandwidth exceeds half the sample rate (i.e., the “Nyquist rate” of the sampling process) introduces spectral folding known as *aliasing*. That is, the sampled process is both missing the true high-frequency components and containing false low-frequency components due to aliasing. No practical measurement and storage system can implement uniform sampling at frequencies that are high enough to avoid this problem.

A close look at factors influencing packet arrival rate reveals several that lead to high-frequency periodicity causing aliasing in uniform sampling of observed traffic. These include deterministic link transmission times, common speeds for numerous switches and routers, packet lengths that are mostly fixed due to long TCP flows, oscillations induced by TCP feedback loops, etc. As we show with examples from real data,

these periodicities are manifested as spikes in the frequency spectrum of the traffic process. The frequencies of these spikes (and parameters of the corresponding process components) provide exactly the information that is most valuable for understanding, modeling, planning and design, and it is exactly the information that uniform sampling misses. Misleading spectral information causes severe problems for the very purposes for which we are attempting traffic measurement in the first place: it impedes understanding of the traffic components and it leads to erroneous statistics (mean, variance, and correlations) used for network planning, design and operation. Given these considerations, it is perhaps not too strong a statement to say that the currently ubiquitous uniform sampling methods are fundamentally unsuitable for measuring Internet traffic.

It has long been recognized that one way to overcome aliasing in sampling is to sample at random intervals rather than at uniform intervals. In particular, *Poisson sampling* (that is, sampling with random intervals that are independent and exponentially distributed with a given average rate) has the property that the frequency domain error due to sampling appears as constant power noise rather than aliased components. While this noise could potentially drown out flat parts of the observed signal’s power spectrum, its constant-power nature also allows spectral peaks (at *any* frequency) to be correctly recognized in the sampled data (as long as the peak amplitudes exceed the noise floor).

In this paper, we develop a stochastic sampling framework for Internet traffic measurement. Starting from the theoretical foundations of Poisson sampling and Internet traffic models, we develop a modified sampling algorithm that is useful for practical, networking-specific measurement applications. We use recent advances in the theory of *Poisson counter-driven stochastic differential equations* (PCDSDE) for network traffic modeling to develop a technique for estimating important network characterization parameters from the sampled data. We illustrate the application of this measurement techniques with results from real network traffic.

In the remainder of this paper, Section II discusses related work. The theoretical foundations of stochastic sampling are introduced and refined for the network measurement domain in Section III. Section IV discusses practical implementation issues and applications. Section V presents results to illustrate the benefits of stochastic sampling over uniform sampling. Section VI summarizes and concludes this paper.

II. RELATED WORK

A large number of commercial and open-source network measurement tools use uniform sampling when obtaining and reporting such measurement results. Examples are:

- **SNMP Counts:** The Simple Network Management Protocol (SNMP) [1] provides a mechanism to access information on SNMP-enabled devices. Typically, information is retrieved by periodic polling the SNMP agent, which is equivalent to uniform sampling. Typical SNMP-based monitoring applications poll information at intervals of minutes.
- **Cisco Netflow:** Cisco Netflow is a tool that provides information about the data flows (defined by the five-tuple of source and destination IP addresses and port numbers and the protocol type) that are observed on a router [2]. The default interval for reporting flow information is five minutes. One Netflow feature is to consider only every n^{th} packet (“Sampled Netflow” [3]), which is equivalent to uniform subsampling within the sampling interval. An improved version of Netflow uses adaptive sampling [4], but the adaptation is relatively slow and subsampling is still done uniformly.
- **Active Probing:** Many active measurement tools, for example `ping` and `traceroute` in the simplest case, use uniform intervals to sample round-trip time delays. On a larger scale, for example in PlanetLab, several measurement services (e.g., all pairwise pings or path failure monitoring) use uniform sampling intervals.

The key observations from these existing systems are that (1) some sampling intervals are extremely long (e.g., SNMP and Netflow) and thus do not yield any information on any interactions that have periods below the minute range and (2) that practically all measurement applications use uniform sampling. However, many applications that *use* measurement data rely on accurate frequency information. For example, anomaly detection has been approached from a signal processing point of view [5]. For a comprehensive overview of measurement applications and general sampling techniques see [6].

The key properties of Poisson sampling, which is the basis for our work, were derived in the 1970s [7], [8], and are nicely summarized in [9]. We adapt these concepts to network measurement using Poisson counter-driven stochastic differential equations (PCDSDE) models. PCDSDE have recently been applied to a variety of modeling problems ranging from TCP window size [10] to admission control policies [11]. Poisson sampling is similar to packetized Bernoulli sampling, but only when packet streams exhibit deterministic inter-interval times. When packet streams are random – as is the case for real traffic – they are different.

We focus on the frequency spectrum of network traffic in the context of this work since it contains important information on traffic characteristics. Li et al. have observed that frequency characteristics of network traffic are preserved even when passing through finite-buffer systems [12]. However, these buffers can act as low-pass filters and truncate the high-frequency

spectrum. With networks moving towards small buffer router designs [13], more high-frequency information is preserved and aliasing caused by uniform sampling becomes a growing problem in network measurement.

III. POISSON SAMPLING FOR NETWORK MEASUREMENT

In order to put our approach to stochastic sampling into context, we briefly review the theoretical foundations of random sampling. It is important to note that even though our work is based on Poisson sampling, it does not imply that we assume a Poisson network traffic model (which has been shown to have significant shortcomings [14]). The issue of measurement, which is what we address, is independent of the issue of traffic modeling.

A. Theoretical Foundations of Random Sampling

Suppose that a signal (deterministic or random) is sampled uniformly at a rate of f_s samples/sec. As derivations of the sampling theorem [9], [15] show, the sampled signal *cannot* represent frequencies greater than $f_s/2$ Hz; instead, any components in the original signal at frequencies greater than $f_s/2$ Hz are “folded over” to be seen as components at frequencies less than $f_s/2$ Hz. (This is the effect known as aliasing.) Hence, statistics obtained from low-rate uniform samplings of high-bandwidth signals could be completely unreliable: high-frequency terms (that help to determine signal variance and correlations) are missing from the sampled data while low-frequency terms (that determine the mean and contribute to the variance) are distorted by aliasing.

To reduce the impact of aliasing, nonuniform adaptive sampling has been proposed [16], [17], where the sampling interval is adjusted according to estimates of the “instantaneous bandwidth” of the signal. The complexity of a bandwidth estimator can be avoided when using a stochastic sampling technique known as *Poisson random sampling*.

Poisson random sampling is a scheme for sampling with random intervals that are independent and exponentially distributed with a given average rate λ . In a sense, it can be viewed as a stochastic approximation to uniform sampling since it can be shown that Poisson samples over any finite interval have a uniform distribution [7], [8]. In particular, using the properties of Poisson point processes it is shown in [9] that if a deterministic signal $f(t)$ has Fourier transform $F(\omega)$, an *unbiased* estimate of $F(\omega)$ can be obtained in the form

$$\mathbf{P}(\omega) = \frac{1}{\lambda} \sum_i f(\mathbf{t}_i) e^{-j\omega \mathbf{t}_i} \quad (1)$$

where \mathbf{t}_i is a Poisson point process with rate λ . Furthermore, the variance of $\mathbf{P}(\omega)$ (for any ω) is equal to $\frac{E_f}{\lambda}$ where E_f is the energy in $f(t)$ (given by $E_f = \int_{-\infty}^{\infty} f^2(t) dt$). (Similar results are obtained if instead of a deterministic signal we are sampling a random process – the sampled signal’s power spectrum is the true power spectrum plus a constant that is proportional to the average power of the random process divided by λ [9], [18].)

These results have several key consequences for our application to Internet traffic measurement. In particular, they show that there is no aliasing with Poisson sampling – the expected value of the spectrum obtained from Poisson samples (no matter how low the rate) is equal to the true spectrum at *every* frequency. The sampling error is seen as constant power noise through the entire spectrum, with the noise power inversely proportional to the sample rate. As pointed out in [9], this latter fact implies that a satisfactory estimate of $F(\omega)$ is obtained from Poisson samples if $|F(\omega)| \gg \sqrt{\frac{E_f}{\lambda}}$. But again, our results presented in Section V show the presence of multiple periodic components in the Internet traffic spectrum (we present a model in Section III to account for these periodicities as resulting from clocks modulated by on-off signals). These periodicities appear as high-amplitude spectral peaks. So, as long as the peak amplitude is significantly larger than $\sqrt{\frac{E_f}{\lambda}}$, the corresponding component (at *any* frequency) are detectable in the Poisson samples. This is the crucial difference between our Poisson sampling scheme and current uniform sampling schemes, and it is the key to extracting useful network statistics from practical low-rate sampling.

B. Poisson Sampling for Network Measurement

Ideal Poisson sampling requires exponentially distributed sample intervals, so that a given interval has a non-zero probability of being either too short for any realistic sampling mechanism or too long for sampling over any finite interval. A variation on Poisson sampling called *minimum distance Poisson sampling* was proposed in [18] to enforce a minimum allowable distance between the sample points. The ideal Poisson sampling process has a power spectrum equal to a constant (the inverse of the rate) plus an impulse at zero frequency. The power spectrum of the modified sampling process is shown in [18] to have the form of a non-constant function of frequency added to a zero-frequency impulse. The non-constant power spectrum component increases the error in recovering high-frequency components from the samples. In addition to enforcing a non-zero lower bound on the distance between sample points, for practical sampling of Internet traffic we also need to impose a restriction on the maximum distance between samples (thus creating a *truncated* Poisson sampler).

A second practical consideration lies in averaging vs. sampling. Even the fastest current sampling techniques actually perform averaging over short intervals. The frequency domain effect of this is to multiply the sampled spectrum by a broad sinc envelope that diminishes high-frequency components. We need to account for this effect in our analysis. Another question is whether we should be just sampling (or short-term averaging), or as is the case in most current systems, counting the arrivals between sample times. The latter scheme is equivalent to sampling the *integrated* traffic. Integration has the effect of multiplying the original signal power spectrum by a factor of $1/\omega^2$. This would reduce the constant power noise in the spectrum of the Poisson sampled signal, at the

cost of diminished values for any high-frequency peaks.

Finally, in addition to recovering spectral peaks and the corresponding signal component parameters, it is of interest for some applications such as network forensics and post-mortem analysis to reconstruct the actual traffic signal from Poisson samples. One possibility would be simply to invert an FFT of the sampled data, or that part of the FFT that exceeds some threshold. A method that was shown in [19] to yield better results involved polygonal interpolation of the Poisson samples.

C. Traffic Modeling and Link Parameter Estimation

As we have noted above (and demonstrate with the real traffic data example in Section V), periodic components (at both low and high frequencies) predominate in Internet traffic. We have discussed how these components may be missed or erroneously identified from uniform traffic samples, with negative consequences for our ability to derive the network statistics needed for planning, design and pricing policy. We have also shown that in contrast, the measurements generated by our Poisson sampling approach are ideally suited for recovering periodic components in the traffic data. The questions that needs to be addressed is how these periodic components arise, how they relate to key link characteristic parameters, and how the link parameters may be estimated from the measurements.

We first note that the problem of identifying traffic characteristics from sampled data has received significant recent attention. For example, an algorithm for using counters to identify the top k categories (e.g., destination IP addresses) to which individual packets belong is presented in [20]. The work that is perhaps closest in spirit to our approach is described in [21]. That work, which attempts to recover the packet-level power spectral density and the distribution of the number of packets per flow, is based on flow sampling (thinning) rather than packet sampling. The authors demonstrate accurate estimation of the spectrum and flow statistics even after significant thinning. However, this approach requires that the observed traffic be modeled as the superposition of independent and identically distributed packet flows, and it requires a preprocessing step of classifying packets into flows before sampling.

In contrast, our approach is based on a hierarchical model for the observed packet stream on a link. This stream is a binary signal with relatively short and infrequent 1's (indicating packets) separated by strings of 0's (indicating no packets). To capture this characteristic, we assume that the fundamental packet generation process is a Markov on-off process with short on duration and long off duration. This process (say, $x(t)$) is generated by a *Poisson Counter Driven Stochastic Differential Equation* (PCDSDE) [22] of the form

$$dx(t) = -x(t)dN_1(t) + (1 - x(t))dN_2(t), \quad (2)$$

where $N_1(t)$ and $N_2(t)$ are independent Poisson counters having rates λ_1 and λ_2 , respectively. (Note that the relative sizes of the Poisson counter rates provide an indicator of

the “busy-ness” of the packet generation process.) Using the independence property of Poisson counter jumps, it is not difficult to show that the stationary process generated by Equation 2 has mean $\mu_x = (\lambda_1 + \lambda_2)^{-1} \lambda_2$ and autocovariance function $C_x(\tau) = C_x(0) \exp\{-(\lambda_1 + \lambda_2)|\tau|\}$. Hence, the power spectrum $S_x(\omega)$ of $x(t)$ is low-pass with an impulse at zero frequency of size μ_x^2 and a half-power frequency of $\lambda_1 + \lambda_2$ rads/sec.

However, what we observe when we measure link traffic is not $x(t)$; rather (because of factors such as those mentioned above), we see a process $y(t)$ that amounts to a high-frequency periodic signal $s(t)$ (which we call a clock) *modulated* by $x(t)$. That is, we obtain samples of $y(t) = x(t)s(t)$, which (assuming independence of $x(t)$ and $s(t)$) implies that the autocorrelation function of $y(t)$ is the product of the autocorrelations of $x(t)$ and $s(t)$ and hence that the power spectrum $S_y(\omega)$ of the observed process is the *convolution* of $S_x(\omega)$ and $S_s(\omega)$ (the clock power spectrum). The periodic nature of the clock signal autocorrelation function implies that its power spectrum consists of a series of impulses at the fundamental clock frequency and its harmonics. Hence, $S_y(\omega)$ consists of a series of shifted versions of $S_x(\omega)$ centered at harmonics of the clock frequency. So, in the observed signal’s power spectrum we see impulses (due to μ_x) and narrow pulses (due to $C_x(\tau)$) centered at the clock frequency harmonics. We believe that this basic PCDSDE-modulated clock model can both explain the presence of the high-frequency spikes seen in the traffic measurement spectrum and indicate how parameters of those spectral peaks (frequencies and amplitudes) relate to network characterization parameters such as clock frequencies and Poisson counter rates. Since our Poisson sampling algorithms can recover the power spectrum peaks, they provide us with a new opportunity to extract these fundamental network operation parameters from traffic measurements.

We can develop models that are extensions of the above basic model. In particular, it has been argued recently [23] that typical packet generation processes are more accurately modeled by a product of Markov on-off processes; that is, by $x_1(t)x_2(t) \cdots x_N(t)$ where each $x_k(t)$ is a Markov on-off process of the type described above. (In fact, it has been shown that such a model can account for the long-range dependency seen in some Internet traffic data [24].) This product process may in turn modulate several clocks having different frequencies, so in general we observe the sum of modulated signals $y(t) = x_1(t)x_2(t) \cdots x_N(t)\{s_1(t) + s_2(t) + \cdots + s_L(t)\}$. It is possible to validate these models by comparing the results in time and frequency domains to actual Internet traffic data sets.

IV. PRACTICAL CONSIDERATIONS

To illustrate how the theoretical foundations of stochastic sampling can be applied to practical network measurement scenarios, we discuss three general problems that appear in measurement: (1) how to merge multiple measurement data, (2) how to maintain information with limited storage, and (3) how to perform stochastic sampling online.

A. Merged Data Analysis

In many network measurement scenarios, data is collected in a distributed fashion to obtain a broader view of the state of the network than can be achieved with a single measurement point [25]–[27]. Utilizing multiple measurement systems raises a number of challenges that are related to the issue of sampling: (1) Data collected from different measurement nodes needs to be merged into a single coherent “view;” (2) Synchronization between system needs to be considered to ensure the results are consistent.

The merging process combines samples from multiple locations into a single set of samples. However, measurement systems typically do not have a globally synchronized clock with a level of accuracy better than millisecond range. NTP (Network Time Protocol) is limited due to variation in the end-to-end delay and GPS-based or CDMA-based devices are limited due to the typical use of the serial interface as communication port). When no synchronization can be achieved among distributed measurements, it is difficult to correlate information. A similar problem occurs when two measurement systems sample at different rates (e.g., due to different link rates).

In the case of stochastic sampling, this issue can be solved elegantly in two steps:

- Change of sampling rate: Data from stochastic sampling can be adapted to reflect any lower sampling rate by removing random samples. Uniform sampling only allows for downsampling by integer values (e.g., 1:2, 1:3, etc.), but stochastic sampling can be adapted down to any rate (for long sample traces) by randomly merging adjacent intervals.
- Superposition of samples: Two traces obtained from stochastic sampling (and typically adapted to be sampled at the same rate) can simply be merged to reflect a superposition of the information. The resulting trace is a stochastically sampled trace with a sampling rate equal to the sum of the rates of each trace.

To merge two traces, they both need to be adapted to the same sampling rate and then superimposed. Repeated merging of sampling data with different traces is possible.

B. Aging of Samples

The elegant property of being able to adjust the sampling rate of stochastically sampled trace is also useful when considering long-term storage. For reasons of Internet forensics, post-mortem analysis, etc., it is often desirable to store measurement information for days, months, and even years. With link speeds in the Gigabit per second range, it is practically impossible to store high-resolution measurement data for a long time. Instead, stochastic sampling can be applied to the data and the sampling rate can be adapted for data of different age. The sampling rate can be *continuously* varied from high (for recent data) to low (for oldest data) while keeping the total storage requirement constant.




Fig. 1. Comparison of Sampling Techniques. Packet counts are collected at 1 second intervals (shown in subfigure (a)). The FFT of this trace is shown in subfigure (b). When applying sampling techniques, uniform sampling (subfigure (c)) shows an accurate representation of frequency components below the Nyquist frequency (0.5×10^{-3} Hz in this example), but no information at higher frequencies. Stochastic sampling (subfigure (d)) shows a more complete representation of the original signal albeit with a slightly higher noise floor than the FFT.

C. Online Algorithms for Sampling

As is described in Section III, Poisson samples can generate an unbiased estimate of the complete signal spectrum through Equation 1. In practice, the Fourier Transforms would be calculated with FFTs (as in the results shown in Section V). There is a complication for doing this that is implied by Equation 1: even if we have only a very sparsely sampled subset of the original data, computation of an FFT-based spectrum estimate from Poisson samples requires the FFT length to be that of the *original* data, not the sampled data. (In other words, we have to zero-pad between the Poisson samples before we can take an FFT.) For very long data records obtained by measuring high-speed links over any significant time period, the resulting FFTs may be too computationally demanding for online algorithms. Alternate means of spectrum estimation from Poisson samples have been developed for uniformly sampled harmonic and pole-zero signal models [28], [29].

V. RESULTS

In this section, we present results that compare the information that can be extracted from uniformly sampled and stochastically sampled measurement data. We have performed several experiments to demonstrate the differences between these approaches.

The measurement data consists simply of packet arrival times (as it has been used for anomaly detection [5]). The arrival time traces were collected from the Internet access link at the University of Massachusetts Amherst, which is a 1 Gbps link that is paced to 350 Mbps. Data was collected with Endace DAG4.3E network measurement card [30], which provides packet traces with timestamps with a relative accuracy of tens of nanoseconds.

A. Uniform vs. Stochastic Sampling

In the first experiment, we collected data over a period of approximately 100 hours. Packet counts were taken for the uplink at 1-second intervals and are shown in Figure 1(a) (to obtain a clear frequency spectrum, the packet counts are shifted to a zero average). This figure clearly shows the diurnal cycles in

the traffic patterns (and two spikes due to artifacts in the data collection process). The complete power spectrum of the signal is shown in Figure 1(b), which is obtained by performing a Fast Fourier Transform (FFT) over the packet counts. When applying sampling to the packet count process, Figure 1(c) (uniform sampling) and Figure 1(d) (stochastic sampling) are obtained. In both cases, the signal is downsampled 1:1000. Clearly, Figure 1(c) shows a highly accurate representation of the low frequency components of the packet counts. However, due to uniform sampling, frequency components above the Nyquist frequency (0.5×10^{-3} Hz in this example) are cut off entirely. In comparison, stochastic sampling in Figure 1(d) shows that information over the entire frequency range is available. While the noise floor is somewhat increased over the original signal, there is still an accurate representation of low frequency components as in the uniform sampling case.

In the second experiment, we represented packet arrivals as a point process with a very high resolution of 50 ns. The goal is to illustrate the effect of aliasing in uniform sampling. The FFT of a 0.75 s long interval is shown in Figure 2(a). The uniform sampled version of the power spectrum is shown in Figure 2(b). Clearly, several of the high-frequency spikes above the Nyquist frequency have been folded into the low-frequency range. This aliasing leads to an erroneous interpretation of the measured data. In the stochastic sampling case shown in Figure 2(c), on the other hand, no such aliasing occurs despite sampling.

VI. SUMMARY AND CONCLUSIONS

We believe that stochastic sampling is an important step towards more accurate network measurements. As shown, commonly used uniform sampling misses high-frequency signal components and causes aliasing in low-frequency components. Poisson random sampling avoids these problems and simplifies distributed measurement due to its inherent asynchrony. We have discussed how stochastic sampling can be practically implemented in the context of network measurement. Our results show that stochastic sampling retains wide spectral information, creates no aliasing, and thus avoids the common problems of today's widely used uniform sampling.




Fig. 2. Comparison of Sampling Techniques. Packet arrivals are represented as a point process with 50ns resolution. The full FFT is shown in subfigure (a). Uniform sampling shows the expected cut-off at the Nyquist frequency (1 MHz), but also shows aliasing. The peaks in the original frequency range between 1 and 10 MHz are folded into the frequency range below 1 MHz. Stochastic sampling (subfigure (c)) does not show this aliasing.

REFERENCES

- [1] J. Case, M. Fedor, M. Schoffstall, and J. Davin, "A simple network management protocol (SNMP)," Network Working Group, RFC 1157, May 1990.
- [2] *Cisco IOS Netflow Data Sheet*, Cisco Systems, Inc., Sept. 2004, http://www.cisco.com/application/pdf/en/us/guest/tech/tk812/c1482/cdccont_0900aecd80173f71.pdf.
- [3] *Sampled Netflow*, Cisco Systems, Inc., 2003, http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s11/12s_sanf.pdf.
- [4] C. Estan, K. Keys, D. Moore, and G. Varghese, "Building a better NetFlow," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, Portland, OR, 2004.
- [5] P. Barford, J. Kline, D. Plonka, and A. Ron, "A signal analysis of network traffic anomalies," in *Proc. of the 2nd ACM SIGCOMM Workshop on Internet Measurement (IMW)*, Marseille, France, Nov. 2002, pp. 71–82.
- [6] N. Duffield, "Sampling for passive internet measurement: A review," *Statistical Science*, vol. 19, no. 3, pp. 472–498, May 2004.
- [7] F. Beutler, "Alias-free randomly timed sampling of stochastic processes," *IEEE Transactions on Information Theory*, vol. 16, no. 2, pp. 147–152, Mar. 1970.
- [8] E. Masry, "Poisson sampling and spectral estimation of continuous-time processes," *IEEE Transactions on Information Theory*, vol. 24, no. 2, pp. 173–183, Mar. 1978.
- [9] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*, 2nd ed. McGraw-Hill, 1984.
- [10] V. Misra, W. Gong, and D. F. Towsley, "Stochastic differential equation modeling and analysis of TCP-window size behavior," Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA, Tech. Rep. ECE-TR-CCS-99-10-01, 1999.
- [11] S. W. Tam, D. M. Chiu, C. S. Lui, and Y. C. Tay, "A case for TCP-friendly admission control," Dept. of Information Engineering, Chinese University of Hong Kong, Tech. Rep., 2005.
- [12] S.-Q. Li, S. Chong, and C.-L. Hwang, "Link capacity allocation and network control by filtered input rate in high-speed networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 1, pp. 10–25, Feb. 1995.
- [13] M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, and T. Roughgarden, "Routers with very small buffers," in *Proceedings of the Twenty-fifth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2006)*, Barcelona, Spain, Apr. 2006.
- [14] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE Transactions on Networking*, vol. 3, no. 3, pp. 226–244, June 1995.
- [15] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, *Discrete-Time Signal Processing*, 2nd ed. Prentice Hall, Feb. 1999.
- [16] J. Clark, M. Palmer, and P. Lawrence, "A transformation method for the reconstruction of functions from nonuniformly spaced samples," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 5, pp. 1151–1165, Oct. 1985.
- [17] M. R. Shankar and P. Sircar, "Nonuniform sampling and polynomial transformation method," in *Proc. of IEEE International Conference on Communications (ICC)*, vol. 3, New York, NY, Apr. 2002, pp. 1721–1725.
- [18] M. A. Z. Dippé and E. H. Wold, "Antialiasing through stochastic sampling," in *SIGGRAPH '85: Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, San Francisco, CA, July 1985, pp. 69–78.
- [19] A. Ouahabi, C. Depollier, L. Simon, and D. Koume, "Spectrum estimation from randomly sampled velocity data," *IEEE Transactions on Instrumentation and Measurement*, vol. 47, no. 4, pp. 1005–1012, Aug. 1998.
- [20] E. D. Demaine, A. López-Ortiz, and J. I. Munro, "Frequency estimation of internet packet streams with limited space," in *Proc. of the 10th Annual European Symposium on Algorithms (ESA 2002)*, ser. Lecture Notes in Computer Science, vol. 2461, Rome, Italy, Sept. 2002, pp. 348–360.
- [21] N. Hohn and D. Veitch, "Inverting sampled traffic," in *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, Miami Beach, FL, 2003, pp. 222–233.
- [22] R. W. Brockett, "Lecture notes: Stochastic control," Harvard University, Tech. Rep., 1983.
- [23] V. Misra and W. Gong, "A hierarchical model for teletraffic," in *Proc. of 37th IEEE Conferences on Decision and Control (CDC)*, vol. 2, Tampa, FL, Dec. 1998, pp. 1674–1679.
- [24] S. Resnick and G. Samorodnitsky, "Limits of on/off hierarchical product models for data transmission," *Annals of Applied Probability*, vol. 13, no. 4, pp. 1355–1398, 2003.
- [25] K. Papagiannaki, N. Taft, and A. Lakhina, "A distributed approach to measure IP traffic matrices," in *IMC '04: Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, Taormina, Italy, Oct. 2004, pp. 161–174.
- [26] Y. Tsang, M. Yildiz, P. Barford, and R. Nowak, "Network radar: tomography from round trip time measurements," in *IMC '04: Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, Taormina, Italy, Oct. 2004, pp. 175–180.
- [27] T. Wolf, R. Ramaswamy, S. Bunga, and N. Yang, "An architecture for distributed real-time passive network measurement," in *Proc. of 14th Annual Meeting of the IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, Monterey, CA, Sept. 2006, pp. 335–344.
- [28] D. G. Manolakis, V. K. Ingle, and S. M. Kogon, *Statistical and Adaptive Signal Processing: Spectral Estimation, Signal Modeling, Adaptive Filtering and Array Processing*. Artech House Publishers, Apr. 2005.
- [29] S. L. Marple, *Digital Spectral Analysis with Applications*. Prentice Hall, Jan. 1987.
- [30] *DAG4.3E Network Monitoring Interface Card*, Endace Measurement Systems, 2005, <http://www.endace.com/dag4.3GE.htm>.