Statistical Issues in Computer Networks and Traffic Analysis

By:

Bowei Xi, Xiadong Yang, Vijayan N. Nair and George Michailidis

Technical Report #15-01

Bowei Xi is an Assistant Professor in the Department of Statistics at Purdue University.

Xiadong Yang is a Postdoctoral Fellow at Kansas State University.

Vijayan N. Nair is D.A. Darling Professor of Statistics and Professor of Industrial and Operations Engineering at the University of Michigan.

George Michailidis (corresponding author) is Professor of Statistics, Electrical Engineering and Computer Science at the University of Michigan.

Department of Statistics
Purdue University

March 2015

Statistical Issues in Computer Networks and Traffic Analysis

Bowei Xi, Xiadong Yang, Vijayan N. Nair and George Michailidis*

1 Introduction

Advancements in computer technology and storage capabilities have allowed network engineers to collect very large amounts of data obtained from computer networks to address a number of engineering tasks, including network provisioning, providing high quality-of-service to users in advanced applications such as Internet telephony and television and online gaming, configuring network protocols, fault diagnosis, traffic forecasting, just to name a few [39].

Different types of network data can be collected that differ in their granularity, accuracy, volume and delay [32]. We start by providing a brief high level description of computer network operations. Networks consist of nodes (routers and switches) connected by physical links (optical or copper wires). Data, in the form of packets, are transmitted over the network from one node (called a source) to another node (called the destination) on predetermined paths, or *routes*. A stream of packets from a particular source to a particular destination defines a flow. In many applications, flows are examined at a more granular level, such as the protocol level (e.g. http, ftp) [39]. Data on flow-level traffic can be obtained from NetFlow [14] or similar (tcpdump) technologies that provide very detailed about the flow, including the application, packet and byte volumes, transmission protocol and delays, etc. In principle, such data can be collected for all packets of all flows. However, this is impractical in today's high speed networks, which has led to the implementation of sampling strategies for data collection purposes (see Section 2).

*Bowei Xi is an Assistant Professor in the Department of Statistics at Purdue University. Xiadong Yang is a Postdoctoral Fellow at Kansas State University. Vijayan N. Nair is D.A. Darling Professor of Statistics and Professor of Industrial and Operations Engineering at the University of Michigan. George Michailidis (corresponding author) is Professor of Statistics, Electrical Engineering and Computer Science at the University of Michigan.

The analysis of computer network traffic on a single link has been the focus of a number of studies over the years. Originally, the majority of traffic consisted of transfers of files containing data. With the explosion of the Web and the introduction of new applications (audio and video streaming, etc.) most of the traffic shifted to video transfers through new protocols such as peer-to-peer that altered its characteristics. In Section 3, we provide a review of the evolving characteristics, focusing in particular on Voice-over-IP (VoIP) technologies. Finally, a number of other interesting statistical issues arising from network traffic are discussed in [17].

Another source of data is the aggregate volume on a link. Such data are of limited resolution, since they provide information about the total number of packets (bytes) over all flows that traverse a link over a prespecified period of time. Their temporal resolution ranges from a few seconds (about 20-30 secs) to a few minutes (about 2-10 mins). Their main advantage is their accuracy (no sampling required) and ease of collection and storage. However, if one is interested in extracting information about individual flows, she should solve an inverse type of problem that decomposes the aggregate data into their flow constituent parts. This is the network tomography problem reviewed in Section 4.

Note that our goal has been to provide a brief summary of the main developments in the area under consideration and provide pointers to references that discuss in more technical depth particular statistical models and issues. Further, as with any overview paper, the discussion is heavily influenced by our own research interests in this area.

2 Sampling Issues for Network Traffic

As mentioned in the introductoy section, the collection of the necessary information on every packet traversing a computer network is prohibitive in terms of processing capacity, cache memory and required bandwidth in today's high speed links. Hence, packet sampling techniques have emerged as a scalable alternative to address this problem. The Internet Engineering Task Force working groupa have made a number of recommendations that have been implemented in highspeed routers. Specifically, a request for comments (RFC) was submitted [53] describing packet selection schemes and the parameters needed for them. Another report [18] describes a framework for the packet sampling protocols to be employed, details of the protocols and their tuning parameters are given in [15].

An overview of networking application where sampling proves useful is provided in [19]. In addition, the basic mechanisms of systematic (deterministic) and random sampling as applied to networks is briefly reviewed. An important application area for sampling is to understand the characteristics of traffic flows; specifically, estimate the distribution of the

flow lengths (in number of packets), together with the distribution of their sizes (in number of bytes), as well as the number of active flows over a fixed time period. One issue that arises due to the very low sampling rate (.001-.01) is that not packets from all *active flows* traversing a link during a time period of interest would be observed, which introduces a potential bias, since "small" flows would be missed. In a series of papers, Duffield and his collaborators [20, 21] examined the problem of estimating the length of a single flow and proposed an estimator that scales-up (multiplies) the sampled flow length by the inverse of the sampling factor (e.g. if one employs systematic sampling and selects one out of 100 packets, then the multiplier is 100). They also introduced a model for inferring the flow length distribution that was refined in [49]. In [35], a Bayesian approach for addressing this problem is introduced. We follow the presentation in [49] to introduce the model.

Suppose that there are M active flows on a link, comprised of $N_m, m = 1, \cdots, M$ packets each. The number of packets in each flow gives the flow length. In addition, a Bernoulli sampling mechanism is employed; namely, each packet is selected with probability p, independent of its characteristics (e.g. origin, protocol, application, etc.). Each sampled packet can be uniquely assigned to a particular flow, by observing its flow key obtained from information available in the packet header. Hence, the available data are sampled flow length $n_1, n_2, \cdots, n_r$, where r is the number of sampled flows over a prespecified time interval.

Let $\phi = \{\phi_i\}$ with ϕ_i denoting the probability that a flow contains i packets. Further, let $g_j, j = 0, 1, \cdots, J$ be the frequency of sampled flows of size j, with J being the total number of different sampled flow sizes in all the observed flows. The unobserved quantity g_0 corresponds to the frequency of unsampled flows. An estimate of the total number of active flows is then given by $M = \sum_{j=0}^{J} g_j$, while the observable quantity $r = \sum_{j=1}^{J} g_j$ gives the total number of sampled flows. Letting c_{ij} denote the probability of having j packets sampled, given that the true flow length is i packets, we get that $p_{ij} = \phi_i c_{ij}$ represents the probability that an original flow contains i packets and $j \leq i$ of them have been sampled. Finally, denoting by f_{ij} the frequency of flows of length i with j packers sampled we get the following relationship: $g_j = \sum_i f_{ij}$. We can then postulate the following joint model for the number of original flows and the probability that they contain i packets:

$$L(\phi, M) = \binom{M}{g_0, g_1, ...g_J} \prod_{j \geq 0} (\sum_{i \geq j} \phi_i c_{ij})^{g_j} \tag{1}$$

where $M = \sum_{j=0}^{J} g_j$. The objective becomes to maximize this likelihood function subject to the following constraints: $\sum_i \phi_i = 1$, and $\phi_i \geq 0$, where $i \in S_I = \{i(0), i(1), ..., i(J)\}$, with $i(j)$ denoting the length of a flow being i packets when j of them have been sampled. In [49], $i(0)$ was set to $\frac{1}{2p}$ instead of 0, since an original flow containing 0 packets is rather meaningless. Also, the possible flow lengths values are restricted to integer values closest

to j/p.

In [20] two different likelihood functions were considered, one for ϕ (the flow length distribution) and a separate one for the total number of flows M. However, as noted in [49], the proposed estimate of M in [20] given by the number of observed flows divided by the sampling probability (i.e. r/p) is not the maximum likelihood estimate. The integrated framework for this joint estimation problem that treats M as a nuisance parameter discussed in [49] gives maximum likelihood estimates based on the EM algorithm for both quantities of interest, namely ϕ and M.

In subsequent work [50], an extension of the above framework is provided for estimating the size of the flow distribution in terms of bytes. Specifically, the estimation of the flow sizes (in bytes) is accomplished through a random effects regression model that utilizes the flow length information $\hat{\{\phi\}}$ previously obtained. This problem is also addressed in [43]. Extensive empirical evidence from real network traffic traces suggests that flow length distributions and consequently flow size ones are bimodal, with one mode corresponding to the short flows and the second one to long flows. An adaptation of the likelihood framework presented in (1) is introduced in [50] for this scenario. Finally, two-stage sampling strategies are introduced and their properties discussed. Under such a scheme, in the first stage flows are sampled uniformly with probability p_f irrespective of their lengths, while in the second stage, packets are sampled uniformly with probability p_p from the selected during the first stage flows. Such a mechanism has recently become technically feasible to deploy and it shown that this two-stage sampling scheme overcomes the difficulty posed by length biased sampling, since each flow has an equal probability of being selected.

In [22], trajectory sampling is introduced, where a sampled packet is followed at all the routers on the path it travels and its properties examined. In [24], the spectral properties of the sampled packets are studied and an algorithm to reconstruct the spectrum of network traffic proposed. Finally, in [1] flow sampling is used for addressing efficient anomaly detection issues.

3 Traffic Characteristics and Statistical Modeling of Network Traffic

Since the early 1990s it has been well established that the traffic over a *single link* exhibits intricate temporal dependence, known as *burstiness*, which could not be explained by traffic models developed for telephone networks [31]. To account for these empirical facts,

network researchers, statisticians and probabilists introduced models that exhibited long–range dependence and self–similarity [23], which in turn are affected by the presence of heavy tails in the distribution of file sizes [37]. Further, a *mechanistic* model was presented in [42]. A competing model based on queueing ideas was studied in [33]. These works led to further developments; see eg [25].

We give a brief overview of the main features of such models. Suppose on a fixed network route there are M independent users. Let $\{X(t)\}_{t\geq 0}$ denote the traffic intensity of one such user in bytes per unit time. Thus $\int_a^b X(t)dt$ is the total traffic (bytes) generated by the user during the time interval (a, b). It is assumed that $\{X(t)\}_{t\geq 0}$ is a strictly stationary stochastic process with finite mean. Further, let $\{(T_j, Z_j)\}_{j\in}$ be a stationary marked point process of arrival times T_j's in with marks Z_j's. At time T_j, the user initiates a transmission at constant unit rate, which lasts for a time Z_j. Thus, the traffic intensity at time t equals:

$$X(t) = \sum_{j\in} I(T_j \leq t < T_j + Z_j), \tag{2}$$

where $\cdots \leq T_0 \leq 0 \leq T_1 \leq \cdots$. The process $X(t)$ comes about from two popular traffic models: (1) an $M/G/\infty$ *model*, where the T_j's are arrival times of a Poisson point process with constant intensity, independent of the marks Z_j's or (2) an *On/Off model*, where the Z_j's and the T_j's are *dependent* and the *durations* of the user activity Z_j's are modeled with heavy tailed distributions with finite mean but infinite variance, due to the agreement with a large body of empirical work; see e.g. [16]. The heavy tailed nature of the durations, implies that the process $X(t)$ of user activity is long–range dependent (LRD). For our presentation, we focus on the *On/Off model* and suppose that the tails of the On and Off durations are heavy.

Let now $\{X^{(i)}(t)\}$, $1 \leq i \leq M$ be independent and identically distributed stationary processes modeling the traffic intensities of M users sharing a given route. Then, the cumulative traffic over the route generated by the users is:

$$X^*(T, M) := \int_0^T \sum_{i=1}^{M} X^{(i)}(t)dt.$$

We are interested in the asymptotic behavior of the cumulative traffic fluctuations about the mean:

$$X_0^*(T, M) := X^*(T, M) - X^*(T, M).$$

As shown in [42], if the $X^{(i)}(t)$'s are *On/Off processes*, then

$$\mathcal{L} \lim_{T\to\infty} \frac{1}{T^H} \Big\{ \mathcal{L} \lim_{M\to\infty} \frac{1}{\sqrt{M}} X_0^*(Tt, M) \Big\}_{t\geq 0} = \{B_H(t)\}_{t\geq 0}, \tag{3}$$

where $B_H = \{B_H(t)\}_{t\geq 0}$ is a fractional Brownian motion (fBm) with self–similarity parameter H and '$\mathcal{L}\lim$' denoting finite–dimensional distributions convergence.

Relation (3) shows that the fluctuations of the cumulative traffic about its mean behave asymptotically like the fractional Brownian motion, as the number of users M and the time scale T are sufficiently large. The increments $G(k) := B_H(k) - B_H(k-1),\ k = 1, 2, \ldots,$ of fBm can then serve as a model for the traffic traces of the number of bytes transmitted over the network at certain, sufficiently large time scales.

The order of the limits in (3) is important. If one takes $T \to \infty$ first and then $M \to \infty$, as shown in [42], one obtains:

$$\mathcal{L}\lim_{M\to\infty} \frac{1}{M^{1/\alpha}} \Big\{ \mathcal{L}\lim_{T\to\infty} \frac{1}{T^{1/\alpha}} X_0^*(Tt), M \Big\}_{t\geq 0} = \{\Lambda_\alpha(t)\}_{t\geq 0}. \tag{4}$$

Now the limit process $\Lambda_\alpha = \{\Lambda_\alpha(t)\}_{t\geq 0}$ has *independent* and stationary increments with $\alpha-$stable distributions, with α denoting the tail index of the heavy tailed marginal distribution. It is the Lévy stable motion – the infinite variance counterpart to the Brownian motion.

Relations (3) and (4) show two different regimes for the network. The first involves many users relative to the time scale and the second, just a few users relative to the time scale. Similar results were shown to hold for the $M/G/\infty$ and other activity rate models (see e.g. [34]).

In [**?**], an integration of the above presented single-link flow models together with the underlying routing mechanism is employed to come up with a *network wide global* traffic model. The proposed model arises from a limit approximation of the traffic fluctuations as the timescale and the number of users sharing the network grow. The resulting probabilistic model is comprised of a Gaussian and/or a stable, infinite variance components, depending on the growth regime. It can be succinctly described and handled by certain spacetime random fields.

We discuss next an interesting application related to network traffic modeling, that is becoming prevalent due its cost advantages over classical telephony. Three decades ago nearly all voice communication was carried by the Public Switched Telephone Network (PSTN). Voice communication over the Internet is rapidly gaining popularity nowadays. Voice over IP (VoIP) technology offers less expensive and flexible telephone service to the end users. It is also cost effective for the service providers to maintain one network that transmits both voice and data traffic. To ensure the quality of a VoIP call, stringent Quality of Service (QoS) criteria are imposed. Statistical models and methods prove crucial in understanding the properties of network traffic induced by this new application and in determining the different engineering factors for efficient network resource allocation.

In [48], data were collected on a 100 megabits/sec link of the Global Crossing (GBLX) network in Newark, New Jersey over 48 hrs, resulting in 1.315 billion VoIP packet timestamps and headers. This specific link is between an IP-PSTN gateway and an IP network edge router. The VoIP traffic consists of multiple applications, such as voice calls, faxes, and credit card processing [27]. A VoIP call produces two semi-calls, transmitted over two separate cables. For this data, signals are captured at 64 kilobits/sec and the bits accumulated over 20 ms intervals, resulting in packet sizes of 200 bytes including the headers. As with most other applications, VoIP traffic exhibits strong diurnal patterns, ranging from 2 to 11 megabits/sec. Detailed call records provide additional information, such as whether a call was successfully connected or merely an attempt. Finally, note that VoIP packets have priority over those from other, less time sensitive applications and there is regular packet transmission only when an algorithm detects a signal. Consequently a semi-call consists of alternating transmission (on) intervals and silence (off) intervals.

There have been few studies of live VoIP traffic in the past; see e.g. [7, 13, 52], focusing on estimating Hurst parameters for VoIP traffic. In [48],5 successive subsets of the data, each being a subset of the previous one, to study the properties of the VoIP traffic. They are labeled as *full, processed, arrival, complete-call*, and *detail-augmented*. The *full data* contain 1.315 billion packets from 332018 calls (664036 semi-calls). The *processed data* consist of 144185 calls. Calls that have few packets or large gaps are removed. Although there are a large number of small calls, they contribute little to the traffic bit-rate and have negligible impact on QoS study. The *arrival data* have 144046 calls that arrived during the data collection period. Calls that were in progress at the beginning of the data collection period are removed. The *complete-call data* have 138770 calls that were concluded by the end of the data collection period with an estimated probability 0.9999. The *detail-augmented data* contain 78050 complete calls from the first day, for which the call detail records are available. 50% of the calls in the *complete-call data* with call duration greater than the median duration contribute 97.15% of the bits. 16.50% of the longest calls with call duration greater than 128 sec contribute 87.90% of the bits. Due to the massive size of the data no model assumption is true. We can tolerate deviations from model assumptions for the shorter calls since they have insignificant impact on QoS.

The modeling is for the IP inbound traffic as it is first seen on the network, the offered load not altered by network processing. The measured VoIP traffic is employed as the IP inbound traffic. This assumption is validated by examining the timestamp accuracy and the delay jitter. The hardware timestamp accuracy is reported to be $\pm 0.1\mu$s [2]. The measured packet inter-arrival times confirm the level of the accuracy. Further, examination of the delay jitter of the calls in the *processed data* is undertaken. The voice packets in a transmission interval are generated at a constant rate – 20 ms. The difference between the actual voice packet inter-arrival times and 20 ms is the delay jitter. The measured jitter is

small and shows a 9-jitter cycle. A regression model is fitted to the jitter of the semi-calls that have at least 25 transmission intervals and 20 cycles within each transmission intervals. We use the bisquare robust estimator and 9 explanatory variables, $\cos(2\pi ki/9)$, $i = 0, ..., 4$ and $\sin(2\pi ki/9)$, $i = 1, ..., 4$. The amplitude of the waveform fit does not depend on the sample size of each transmission interval. Both inbound and outbound traffic show similar amplitude of the fitted waveforms. It is highly likely that the gateway packetization algorithm produces the cycle in both directions. Hence the residuals are the variation of the packet inter-arrival times caused by network processing. The 0.999 quantiles of the residuals from both inbound and outbound directions are smaller than 0.4 ms. End-to-end delay jitter should be less than 30 ms [27]. The observed jitter is very small compared with the QoS requirement, confirming that the measured data is a good approximation of the IP-bound traffic.

Due to the presence of diurnal effects, the data were further subdivied to 15 minute intervals where the call-rate is relatively stable for studying the call arrival process. Note that the call arrival process of the *full data* is not a Poisson process, which is a well established model for telephone calls [3]. It is bursty with positively auto-correlated call inter-arrivals due to rapid dialing and other technological innovations [8]. By removing very short calls, leads to an approximately exponential inter-arrival process for the 15 minute interval; hence, a non-homogeneous Poisson process is an accurate approximation of the arrival process over longer time periods.

As noted above, silence suppression is a prevalent feature in VoIP, but has not been studied in the literature. To examine its effect, 1000 calls from the *detail-augmented data* with durations from 2^8 sec to 2^{10} secs are sampled, equally spaced on a logarithmic scale. In addition, a holdover of transmission is imposed when a silence interval is shorter than 25 ms and the interval lengths for a low frequency trend are adjusted and a 0 origin for modeling purposes is used. Let t_r be the adjusted transmission interval length and s_r be the adjusted silence interval length. t_r and s_r of one semi-call form an uncorrelated bivariate time series. $\sqrt{t_r}$ and $\sqrt{s_r}$ are well approximated by a gamma distribution. It was noted that different calls show different interval properties because the live VoIP traffic is a mixture of various VoIP applications. The past studies did not involve live traffic carried by an operational network and did not take the effect of the mixture into consideration [6, 9, 11, 26, 30, 36]. The varying interval properties is modeled through a random effect model for the square-root gamma distribution parameters and the sampling variability is estimated by bootstrapping.

Both the *complete-call data* and the *detail-augmented data* for the call duration distribution were examined. The call duration distribution is complex and can not be fitted by a simple standard parametric distribution. The *detail-augmented data* revealed that the calls can be categorized as attempts and connects. The attempt duration and the connect

duration can be successfully modeled as piece-wise linear Weibull and the overall duration distribution as a mixture of piece-wise linear Weibull distributions. It is likely the linear pieces are caused by the mixture of different applications in the traffic, such as voice calls and various machine-to-machine connections. The call duration cumulative distribution function D is a critical statistical aspect of the VoIP traffic because it affects the time dependence of the packet arrival process. Assume the call arrival rate is a constant C. Then, the auto-correlation function of the number of simultaneous active calls at time lag τ is $\rho(\tau) = C\int_{\tau}^{\infty}(c-\tau)d(c)dc$. When $\sum_{\tau}\rho(\tau) = \infty$ the process is long range dependent. Different from the best effort traffic models for file transfers [4, 38, 45], the VoIP traffic is not formally long range dependent, because it duration distribution does not have a sufficiently heavy tail. Packet counts in 20 ms intervals of the measured data were examined, which are adjusted for non-stationarity. Further analysis of packet counts of a 96 hour synthetic stationary series was undertaking. The spectra plots and variance time plots [5] of the measured and synthetic series confirm the above conclusion. Nevertheless the VoIP traffic is strongly persistent over a long period of about 75min/cycle (99.8% of the calls end within 75 minutes). Such strong persistence has a major impact on QoS.

This detailed analysis of VoIP traffic undertaken in [48] provided valuable insight for modeling purposes. Specifically, two models were considered that use a superposition of sampled or synthetic semi-calls to generate the multiplexed traffic. One model is a semi-empirical model: 1) The call arrivals are generated from a non-homogeneous Poisson process where the arrival rate is determined by the target traffic bit-rate; 2) A semi-call is sampled from the *complete-call data*; 3) The packet arrivals from the semi-calls are superposed to generate the simulated VoIP traffic. The second model is a mathematical model. Instead of sampling semi-calls from the *complete-call data*, different components of a semi-call are generated from parametric models: 1) Call duration is generated from the mixture of piecewise Weibull distributions for attempts and connected calls; 2) Generate the shape and scale parameters of the square-root gamma distributions for the transmission and silence intervals; 3) Generate alternating transmission and silence interval lengths until we reach the generated call duration, and insert packets. The independent and identical distribution assumptions needed for the proposed models are verified from the data. No distinction between the two semi-calls of the same VoIP call occurs, because they have similar bit-rate distributions except for the very short calls. The proposed models can be used for QoS simulation studies for both wireless and wireline networks.

4 Network Tomography

It can be seen that a major effort has been undertaking in collecting network traffic data, understanding its properties and modeling them for a single link. However, understanding the impact of traffic on the whole network in terms of QoS, as well as estimating customers demand are two important *network-wide* problems. Network tomography techniques are designed to address such problems.

As pointed out in [28], there are types of tomography methods: (i) link-oriented methods that collect *passively* packet and network flow information at network devices and (ii) path-oriented methods that collect information about connectivity and latency in a network by *actively* sending probe packets through the network from nodes located on its periphery. The first type of methods help address primarily capacity planning and network routing issues, while the second type of methods are geared towards addressing QoS issues.

The first problem originated in the work of Vardi [44], who also coined the term network tomography, while the second one in [10]. Both types of problems have received a lot of attention in the engineering and statistical literature. Two comprehensive reviews, summarizing the work in the area up to 2004 and 2006, are given in [12] and [28], respectively. We provide a brief introduction to the two tomography problems and provide a summary of some recent developments, past 2006.

We present next the link-oriented problem of network tomography. Let Y_t denote the vector containing the total number of packets traversing all network links over a fixed period of time t. Let A denote the routing matrix of the network; i.e. it is a binary (0/1) matrix with rows corresponding to links and columns to the origin-destination flows. A column of A indicates the path that a flow takes from its origin to its destination. Finally, let X_t denote the vector containing the total number of packets belonging to all the flows. We can then relate X_t and Y_t through the following relationship:

$$Y_t = AX_t, \;\; t = 1, 2, \cdots \tag{5}$$

The goal of network tomography is to infer the distribution of the flows (X) from the observed data (distribution of Y). A moment of reflection shows that this is an ill-posed linear inverse problem, since in general there are many more flows than link measurements, or in other words the routing matrix A is not of full column rank.

To overcome this technical difficulty special statistical parametric models were introduced; specifically, a mean-variance relationship that generates a full rank system of linear equations was used, where flow volumes are assumed *independent* of each other and were modeled either as Poisson or as normally distributed with flow variances proportional to their means. The proportionality assumption leads to identifiability of means through

identifiability of variances (details given in the review papers of [12, 28]). Another class of models motivated by ideas from transportation networks, aims to introduce enough constraints to regularize the inverse problem, so as to obtain a unique solution. For example, gravity models [12], assume that a flow f_j between two nodes k and ℓ is proportional to the total amount of traffic departing from node k and that entering the node ℓ.

In [40], the identifiability problem is examined in considerable generality. Specifically, models that allow a limited degree of dependence between flows (forward and reverse flows between pairs of nodes) are introduced, as well as models that incorporate measurements for both packets and bytes. It is shown that under mild assumptions on the characteristic function of flow volume distribution all $n-$th order cumulants for $n \geq 2$ of X are estimable (uniquely identifiable) from the observed data Y if a particular matrix B which is a function of the routing matrix A is of full rank. Further, utilizing a linear isometry relatinship between the Gaussian and the symmetric α-stable distribution the parameters of the latter are also identifiable (since stable distribution do not have finite variances). Further, it is shown that the matrix B is provably of full rank, if shortest path routing is used for obtaining the routing matrix A. The results in [40] are the most general for to date for addressing the identifiability problem for this version of network tomography. It is of considerable interest to extend them to more general dependence structures amongst flows.

We now turn our attention to the path-oriented version of network tomography. Most of the literature has dealt with tree topologies, as opposed to general graphs. In this instance of the problem, packets are *injected* on the root node of the tree and routed to its leaf nodes. The data collected in the Y vector correspond to average delays (or losses) of a sequence of such packets destined for all the leaf nodes. It can be seen that a similar relationship, namely $Y = AX$, holds, where A again denotes the routing matrix used (for an example see [28]). It can easily be seen that a unique solution could be obtained if A is of full column rank.

It turns out that for this problem one can design the injection of the packets in a particular way so as to achieve identifiability. Specifically, a collection of multicast packets must be used, such that each internal node of the tree is used as a *splitting node* for the multicast packet and all the leaf nodes are used as receiver nodes of such packets. As shown in [29, 47] this is a necessary and sufficient condition for identifiability for tree topologies. A multicast packet is one that travels as a single packet up to an internal node of the tree and then duplicates itself with different copies destined for different leaf nodes. It can be seen that this mechanism induces correlations between the copies of the multicast packet that prove useful in resolving the identifiability problem. In [47] and [17] it is shown that multicast packets that duplicate are sufficient for the task at hand. Maximum likelihood estimates are developed based on the EM algorithm for packet losses [47] and delays [29], respectively. Further, in [17] a log-linear formulation of the problem for estimating the

loss rates is introduced that leads to fast to compute estimates. Finally, the problem of identifiability for this version of the newtork tomography problem on directed acyclic graph topologies is still open, although some partial results are given in [46].

Given that one needs to inject multicast packets to estimate the parameters of interest in path-oriented tomography, an interesting question is how to *design the collection*, so as to minimize interference with the network operations. In [51] this problem is examined for particular directed acyclic graph network topologies that correspond to collections of rooted trees. The underlying design problem is formulated as a *set covering* problem with constraints corresponding to a sufficient condition for identifiabilty postulated in [46]. An integer program formulation of the set covering problem is introduced and fast greedy heuristic algorithms are developed and evaluated (since the set covering problem is provably NP-hard). The proposed algorithms work well and produce good designs for topologies involving dozens of root nodes and hundreds of leaf ones.

Path-oriented network tomography techniques can be used to monitor the link of a large network for congestion and anomalous behavior. This would require monitoring hundreds of links parameters (e.g. loss rates or mean delays). An alternative approach is to monitor similar path parameters that can be easily estimated (without requiring solving the inverse tomography problem). In [51] this problem is comprehensively studied and it concluded that a two-stage strategy is economical and efficient; namely, path parameters are continuously monitored using a variety of statistical process control techniques, until an alert is detected. Once an alert is raised, one solves the tomography problem and identifies the problematic link.

References

[1] Androulidakis G, Chatzigiannakis V, Papavassiliou S. Network anomaly detection and classification via opportunistic sampling. *IEEE Network* 2009, 23:6-12.

[2] Arlos P, Fiedler M. A comparison of measurement accuracy for DAG, Tcpdump and Windump. *www.its.bth.se/staff/pca* 2003.

[3] Babu TVJG, Hayes JF. *Modeling and analysis of telecommunications networks*. John Wiley & Sons; 2004.

[4] Barford P, Crovella M. Generating representative web workloads for network and server performance evaluation. *Proceedings of the 1998 ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, 151-160, New York, NY.

[5] Beran J. *Statistics for Long-Memory Processes*, 1994, Chapman & Hall, Boca Raton, FL

[6] Biernacki A. VoIP source model based on the hyperexponential distribution. *Proceedings of World Academy of Science, Engineering and Technology* 2006, 11:202-206.

[7] Birke R, Mellia M, Petracca M, Rossi D. Understanding VoIP from backbone measurements. *The 26th IEEE International Conference on Computer Communications (INFOCOM)* 2007, 2027-2035, Anchorage, AK.

[8] Bolotin VA. Modeling call holding time distributions for CCS network design and performance analysis. *IEEE Journal on Selected Areas in Communications* 1994, 12:433-438.

[9] Brady PT. A model for generating on-off speech patterns in two-way conversation. *Bell System Technical Journal* 1969, 48:2445-2472.

[10] Caceres, R., Duffield, N.G., Horowitz, J. and Towsley, D. Multicast Based Inference of Network Internal Loss Characteristics, *IEEE Transactions on Information Theory*, 1999, 45, 2462-2480

[11] Casilari E, Montes H, Sandoval F. Modelling of voice traffic over IP networks. *Third International Symposium on Communications Systems Networks and Digital Signal Processing (CSNDSP)* 2002, 411-414, Staffordshire, UK.

[12] Castro, R., Coates, M., Liang, G., Nowak, R.D. and Yu, B. Network tomography: recent developments, Statistical Science, 2004, 19, 499-517.

[13] Ciullo D, Mellia M, Meo M. Traditional IP measurements: What changes in a today multimedia IP network. *Telecommunication Networking Workshop on QoS in Multiservice IP Networks* 2008, 262-267, Venezia, Italy.

[14] Cisco, Introduction to Cisco IOS NetFlow, 2007.

[15] Claise B, Johnson A, Quittek J. Packet sampling (PSAMP) protocol specifications. *IETF RFC 5476* 2009.

[16] Crovella, M.E., Taqqu, M.S. and Bestavros, A. Heavy-tailed probability distributions in the World Wide Web. In R. Adler, R. Feldman, and M. S. Taqqu, editors, *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, pages 3-25, Boston, 1998. Birkhäuser.

[17] Denby, L., Landwher, J.M., Mallows, C.L., Meloche, J., Tuck, J., Xi, B., Michailidis, G. and Nair, V.N. Statistical Aspects of the Analysis of Data Networks, *Technometrics*, 2007, 49, 318-334

[18] Duffield N, Chiou D, Claise B, Greenberg A, Grossglauser M, Rexford J. A framework for packet selection and reporting. *IETF RFC 5474* 2009.

[19] Duffield, N., Sampling for passive Internet measurement: A Review, *Statistical Science*, 2004, 19, 472-498

[20] Duffield, N.G., Lund, C. and Thorup, M., Properties and Prediction of Flow Statistics from Sampled Packet Streams, *ACM SIGCOMM Internet Measurement Workshop* 2002, Marseille, France

[21] Duffield, N.G., Lund, C. and Thorup, M., Estimating flow distributions from sampled flow statistics, *IEEE/ACM Transactions on Networking*, 2005, 13, 325-336

[22] Duffield N, Grossglauser M. Trajectory sampling with unreliable reporting. *IEEE/ACM Transactions on Networking* 2008, 16:37-50.

[23] Erramilli, A., Pruthi, P. and Willinger, W. Self-similarity in high-speed network traffic measurements: Fact or artifact? In *Proceedings of the 12th Nordic Teletraffic Seminar NTS12*, 1995, Espoo, Finland

[24] Grieco LA, Barakat C. An analysis of packet sampling in the frequency domain. *Proceedings of the ACM SIGCOMM Conference on Internet Measurement* 2009, Chicago, Illinois.

[25] Hohn, N., Veitch, D. and Abry, P. Cluster processes, a natural language for network traffic *IEEE Transactions on Signal Processing*, 2003, 51(8): 2229-2244

[26] Jiang W, Schulzrinne H. Analysis of on-off patterns in VoIP and their effect on voice traffic aggregation. *Proceedings of the Nineth IEEE International Conference on Computer Communication and Network* 2000, 82-87, Las Vegas, Nevada.

[27] Karapantazis S, Pavlidou FN. VoIP: A comprehensive survey on a promising technology. *Computer Networks* 2009, 53:2050-2090.

[28] Lawrence, E., Michailidis, G., Nair, V.N. and Xi, B. Network Tomography: A Review and Recent Developments, in *Frontiers in Statistics*, Fan and Koul (eds), 2006, 345-364, College Press, Hackensack, NJ

[29] Lawrence, E., Michailidis, G., Nair, V. N. Network Delay Tomography Using Felxicast Experiments, *Journal of the Royal Statistical Society*, B, 2006, 785-813

[30] Lee H.H, Un C.K. A study of on-off characteristics of conversational speech. *IEEE Transactions on Communications* 1986, COM-34:630-637.

[31] Leland, L., Taqqu, M.S., Willinger, W. and Wilson, S. On the self-similar nature of Ethernet traffic. *Computer Communications Review*, 1993, 23:183-193, 1993

[32] Liu, Y., Towsley, D., Ye, T. and Bolot, J. C. An information-theoretic approach to network monitoring and measurement. *Proceedings of the ACM SIGCOMM conference on Internet Measurement*, 2005, Berkeley, CA

[33] Mikosch, T., Resnick, S., Rootzén, H. and Stegeman, A. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *The Annals of Applied Probability*, 2002, 12(1):23-68

[34] Mikosch, T. and Samorodnitsky, G. Scaling limits for cumulative input processes, *Mathemtics of Operations Research*, 2007, 32(4):890-918

[35] Mori,T, Uchida, M. and Kawahara, R. Identifying elephant flows through periodically sampled packets, *ACM SIGCOMM Internet Measurement Workshop*, 2004, Taormina, Italy

[36] Norwine AC, Murphy OJ. Characteristic time intervals in telephone conversation. *Bell System Technical Journal* 1938, 17:281-291.

[37] Park, K. and Willinger, W. (eds) *Self-Similar Network Traffic and Performance Evaluation*, 2000, Wiley & Sons, New York, NY

[38] Paxson V, Floyd S. Wide-area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking* 1995, 3:226-244.

[39] Peterson, L.L and Davie, B.S. *Computer Networks, A Systems Approach*, (3rd Ed), 2003, Morgan Kaufmann, New York, NY

[40] Singhal, H. and Michailidis, G. Identifiability of Flow Distributions from Link Measurements with Applications to Computer Networks, *Inverse Problems*, 2007, 23, 1821-1850

[41] Stoev, S., Michailidis, G. and Vaughan, J. On Global Modeling of Backbone Network Traffic, *Proceedings of IEEE Infocom*, 2010, San Diego, CA

[42] Taqqu, M.S., Willinger, W. and Sherman, R. Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review*, 1997, 27(2):5-23

[43] Tune P, Veitch D. Towards optimal sampling for flow size estimation. *Proceedings of ACM SIGCOMM conference on Internet measurement*, 2008, New York, NY.

[44] Vardi, Y. Network tomography: Estimating source-destination traffic intensities from source data, *Journal of the American Statistical Associaiton*, 1996, 91, 365-377

[45] Willinger W, Taqqu MS, Sherman R, Wilson DV. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking* 1997, 5:71-86

[46] Xi, B. Estimating internal link loss rates using active network tomography, Unpublished doctoral dissertation, 2004, Department of Statistics, The University of Michigan

[47] Xi, B., Michailidis, G. and Nair, V. Estimating Network Loss Rates Using Active Tomography, *Journal of the American Statistical Association*, 2006, 1430-1449

[48] Xi, B., Chen, H., Cleveland, W.S. and Telkamp, T. Statistical analysis and modeling of Internet VoIP traffic for network engineering. *The Electronic Journal of Statistics* 2010, 4:58-116.

[49] Yang, L. and Michailidis, G. Estimation of Flow Lengths from Sampled Traffic, *Proceedings of IEEE Globecom Conference*, 2006, San Francisco, CA.

[50] Yang, L. and Michailidis, G. Sampled Based Estimation of Network Traffic Flow Characteristics, *Proceedings of IEEE Infocom Conference*, 2007, Anchorage, AK.

[51] Yang, X. Network Monitoring, Unpublished doctoral dissertation, 2007, Department of Statistics, The University of Michigan,

[52] Zhang, G., Xie, G., Yang, J. and Zhang, D. Self-similar characteristic of traffic in current metro area network. *Proceedings of the 15th IEEE Workshop on Local and Metropolitan Area Networks* 2007, 176-181, Princeton, New Jersey.

[53] Zseby, T., Molina, M., Duffield, N., Niccolini, S. and Raspali, F. Sampling and Filtering Techniques for IP Packet Selection, *IETF RFC 5474* 2009