

Traffic Modeling of IP Networks Using the Batch Markovian Arrival Process

Alexander Klemm, Christoph Lindemann, and Marco Lohmann
University of Dortmund
Department of Computer Science
August-Schmidt-Str. 12
44227 Dortmund, Germany
<http://www4.cs.uni-dortmund.de/~Lindemann/>

Abstract

In this paper, we identify the batch Markovian arrival process (BMAP) as analytically tractable model of choice for aggregated traffic modeling of IP networks. The key idea of this aggregated traffic model lies in customizing the batch Markovian arrival process such that the different lengths of IP packets are represented by rewards (i.e., batch sizes of arrivals) of the BMAP. The utilization of the BMAP is encouraged by the observation that IP packet lengths follow to a large extent a discrete distribution. A comparative study with the MMPP and the Poisson process illustrates the effectiveness of the customized BMAP for IP traffic modeling by visual inspection of sample paths over four different time-scales, by presenting important statistical properties, and by analysis of traffic burstiness using R/S statistics. Additionally, we show that the BMAP model outperforms MMPP and Poisson traffic models by comparison of queuing performance.

Key words:

Analytical modeling techniques,
applications including high speed networking, mobile communication,
and the Internet,
analytical/numerical models of aggregated IP traffic,
parameter estimation for stochastic processes.

1 Introduction

Traffic characterization and modeling constitute important steps towards understanding and solving performance-related problems in future IP networks. In order to perform reliable traffic characterization and traffic modeling accurate and detailed IP network measurements have to be conducted. Various detailed measurements have been performed in local area networks (LAN), e.g. [6], [11], as well as in wide area networks (WAN), e.g. [7], [10]. The central idea of traffic modeling lies in constructing models that capture important statistical properties of the underlying measured trace data [2]. For IP traffic, important statistical properties are *burstiness* and *self-similarity*. Intuitively, this means that measured IP traffic shows noticeable sustained periods with arrivals above the mean (i.e., bursts) over a wide range of different time-scales [15]. Aggregated traffic models capture the entire traffic stream without explicitly considering individual traffic sources, e.g. the traffic originated by individual users. The problem of accurately capturing these properties in aggregated traffic models has been solved for non-analytically tractable models but is still subject of current research interest for analytically tractable models. Non-analytically tractable models, e.g. fractional Gaussian noise (fGN) and fractional autoregressive integrated moving average (fARIMA), naturally capture burstiness as well as self-similarity. Various research papers have subjected these models, e.g., Ledesma and Liu reported the effective construction of fGN in [8].

For analytically tractable models, e.g. the Markov-modulated Poisson process (MMPP, [4]), recent work has been proposed that utilizes the MMPP in order to mimic self-similar behavior [1], [16]. Skelly, Schwartz, and Dixit [13] utilized the MMPP for video traffic modeling. They described a simple and efficient method for parameter estimation of a general MMPP based on the match of the marginal distribution of the real arrival rate. The class of batch Markovian arrival process (BMAP, [9]) includes the well known Poisson-process, MMPP, and Markovian arrival process (MAP, [9]) as special cases and additionally associates rewards (i.e., batch sizes of arrivals) to arrival-times. However, due to the additional rewards the BMAP provides a more comprehensive model for representing IP traffic than the MMPP or the MAP, while still being analytically tractable.

The challenge for employing BMAPs to model IP traffic constitutes the proper parameter estimation for this arrival process from the given trace data. In fact, measured trace data does not contain all statistical properties required for the unique specification of a corresponding BMAP. Due to this incomplete data, the parameters for a BMAP cannot be properly estimated by standard statistical techniques, e.g. moment matching. Dempster, Laird, and Rubin introduced the expectation-maximization (EM) algorithm [3] for computing maximum likelihood estimates from incomplete data. Ryden tailored the EM algorithm for the MMPP and developed an implementation [12]. To the best of our knowledge, tailoring the EM

algorithm for BMAPs, developing a numerical stable implementation and utilizing the BMAP for traffic modeling is an open research problem.

In this paper, we present a detailed analysis of IP traffic measurements recently conducted at an Internet service provider (ISP) dial-up link. We derive parameterized general distributions for session-level, connection-level, and packet-level characteristics according to different application-types. Because of the nature of generally distributed sources, this detailed synthetic traffic model constitutes a non-analytically tractable traffic model. According to [14], our analysis confirms that 80% of the mass of the packet-length distribution is concentrated on the three packet length 40 bytes, 576 bytes, and 1500 bytes. Based on these observations, we introduce an aggregated traffic model for IP networks that is both analytically tractable and closely captures the statistics of the measured traffic data. The key idea of this aggregated traffic model lies in customizing the batch Markovian arrival process such that these different lengths of IP packets are represented by rewards of the BMAP. We introduce an efficient method for estimating the parameters of a BMAP with the EM algorithm. Furthermore, we present computational formulas for the E-step of the EM algorithm and show how to utilize the EM algorithm for the effective parameter estimation of BMAPs. In order to show the advantage of the BMAP modeling approach over other widely used analytically tractable models, we compare the customized BMAP with the MMPP and the Poisson process by means of visual inspection of sample paths over four different time-scales, by presenting important statistical properties, by formal analysis of traffic burstiness using R/S statistics, and by queuing system analysis.

The paper is organized as follows. Section 2 presents the analysis and characterization of the measured IP traffic. To make the paper self-contained, Section 3 recalls the definition and properties of the BMAP and provides a primer to the EM algorithm. Moreover, we introduce effective computational formulas for the expectation step (E-step) and the maximization step (M-step) tailored to the BMAP, and present a framework for traffic modeling of aggregated IP traffic that utilizes the BMAP. In Section 4, a comparative study illustrates the effectiveness and accuracy of the proposed traffic model. Finally, concluding remarks are given.

2 IP Traffic Measurement and Characterization

2.1 Dial-up Modem/ISDN Traffic Measurements

We conducted detailed traffic measurements at the ISP dial-up modem/ISDN link of the University of Dortmund. At the time of performing these measurements in January 2001 the university offered free Internet access for students and employees, so the users' costs depend on the telecommunication-tariffs and session duration only. Therefore, measurements in an

ISP environment can be regarded as characteristic for session-oriented Internet traffic. In the four-week measurement period approximately 110,000 user sessions with a total data volume of 120 GB have been logged.

The measurements were performed by the software package *TCPdump* running on a Linux host that sniffs all IP packets in the Ethernet segment between the MaxTNT dial-up routers and the Internet router (see Figure 1). For all IP datagrams sourced or targeted by dial-up modems the TCP/IP header information in conjunction with a timestamp of the arrival-time have been recorded and stored for offline processing. The header information includes source and target IP addresses, the port numbers, the packet length, and the TCP header flags. In addition to the header trace we use the log-files generated by the MaxTNT dial-up routers. For each dial-up session they provide information about session start- and end-time, the assigned IP address, and the link bandwidth. By aligning the measured trace data with the MaxTNT log-files, the header trace can be split into separate trace files for each dial-up session. Furthermore, the session interarrival-time and session volume distributions are determined.

In order to derive the connection interarrival-time and connection volume distributions, for each application-type all TCP connections within each session are reconstructed by means of IP address and port number pairs. We observed that many HTTP connections persist a relatively long period related to the transmitted data volume. Further investigations of this phenomenon showed, that many HTTP connections are closed with a *reset* packet, which is transmitted a very large time after the other packets of this connection. This is caused by the HTTP implementations of most Internet browsers, which keep connections open in order to transmit several documents within the same connection and thus avoid overhead for connection establishment. The reset packets abortive release the connections when the user terminates the HTTP application. For our traffic analysis we ignore these “late reset” packets for the following reasons: (1) they do not contribute to the transmission of documents and (2) they affect the packet interarrival-time distribution in a way that the measured data cannot be well represented by fittings to general distributions.

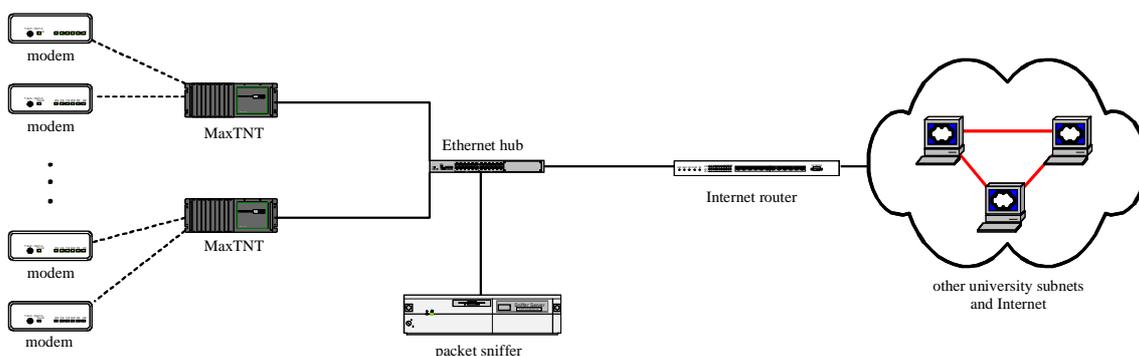


Figure 1. Measurement environment at the dial-up modem/ISDN links

2.2 Traffic Analysis and Traffic Characterization

This section presents fundamental characteristics of the IP traffic measured at our university. First, we analyzed the trace data in order to obtain the application usage pattern of the modem/ISDN users. The data volume fractions are broken down to HTTP, FTP, e-mail, Napster, UDP, and other TCP applications. As Figure 2 shows, HTTP applications, e.g. Web browsing, dominate with a fraction of 73%, followed by the popular music download application Napster with 9% and e-mail with 6% of the overall transmitted data volume. Note, that the court decision against Napster was not effective when conducting the measurements. The relatively small fraction of 2% for FTP applications and the observation, that there is a significant number of HTTP connections comprising of very large data volume, show that file downloads are increasingly performed via HTTP. With 4% of the total data volume UDP applications form a relatively small part of the overall application usage. Taking into account that DNS lookups are performed via UDP the fraction of real-time UDP applications is so small, that realistic statistical measures about real-time traffic cannot be derived from this data. Therefore, we focus our investigations on non real-time traffic.

The analysis is performed at three different traffic-levels: session-level, connection-level, and packet-level.

- (1) The session-level describes the dial-up behavior of the individual users, characterized by the session interarrival-time distribution and the session data-volume distribution.
- (2) The connection-level describes for each individual application the corresponding distribution of connection interarrival-times within a user-session as well as the distribution of connection data volume.
- (3) The packet-level characterizes the packet interarrival-time distribution and the packet length distribution within the application specific connections.

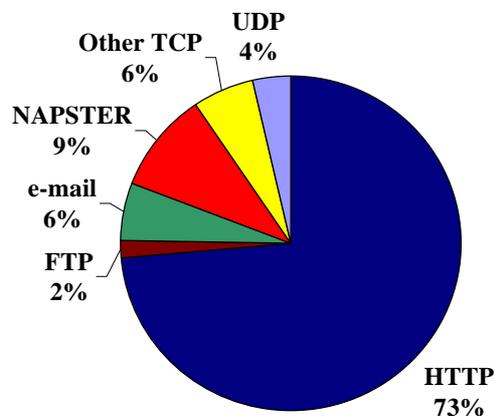


Figure 2. Application usage distribution with respect to data volume

According to the users' different application-types parameterized distributions have been derived for the statistical measures at each traffic-level. We observe that each statistical measure comprises of a *characteristic distribution*, e.g. the HTTP connection interarrival-times are distributed according to a lognormal distribution. In order to find such a characteristic distribution for a specific statistical measure we use a least-squares regression, utilizing the following set of probability density functions (pdf): Lognormal, Pareto, Weibull, Gamma, and Exponential. The detailed statistics of this analysis are omitted because of space limitations.

The analysis of the packet length distributions revealed that the packet lengths of all relevant TCP applications follow to a large extend a discrete distribution. In Figure 3 the probability density of packet lengths is depicted and shows, that the packet lengths 40 bytes, 576 bytes and 1500 bytes dominate with an overall percentage of 80% of all TCP packets. This observation can be explained with the maximum transfer units (MTU) of the used network protocols. Most application protocols like FTP, HTTP, POP3, and SMTP are used to transfer relatively large data blocks (opposed to many small packets in real-time applications). Therefore, in order to reduce overhead, as many packets as possible are filled up to the MTU of the underlying protocol, which typically comprises of 1500 bytes in the Ethernet protocol and 576 bytes in the serial line Internet protocol (SLIP). The choice between a MTU of 576 bytes or 1500 bytes depends on the network configuration of the dial-up client. The huge amount of 40 bytes packets is to a large extend caused by TCP acknowledgments with an empty data field. Recall, that the TCP and IP headers without any options consist of 40 bytes. Furthermore, we observe that the remaining packet lengths are distributed uniformly between 40 bytes and 1500 bytes. Table 1 presents the fractions of these discrete packet lengths for

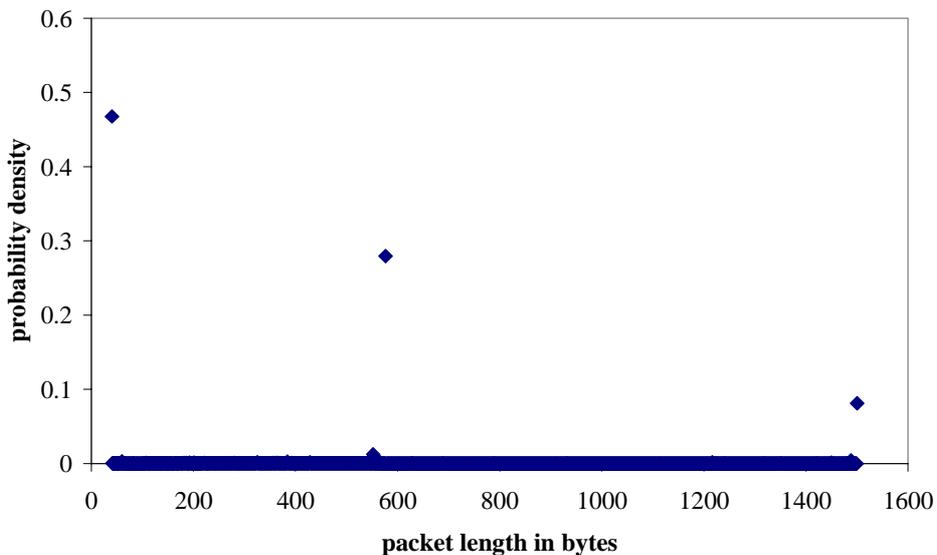


Figure 3. Packet length distribution of TCP packets in dial-up ISP networks

	packet length 40 byte	packet length 576 byte	packet length 1500 byte	other packet lengths
HTTP	46.77 %	27.96 %	8.10%	17.17 %
e-mail	38.25 %	25.98 %	9.51 %	26.26 %
Napster	34.98 %	45.54 %	4.18 %	15.30 %
FTP	40.43 %	18.08%	9.33 %	32.16 %

Table 1. Fractions of different packet lengths in overall traffic

HTTP, e-mail, Napster, and FTP. In contrast to TCP packets, UDP packet lengths follow lognormal distribution. Note, that similar characteristics concerning discrete packet size distributions have been observed for local area networks (LAN) and wide area networks (WAN), [14]. Thus, the basic ideas outlined in the next section can also be applied for traffic modeling in LAN and WAN.

The statistical properties on the three different traffic-levels can be employed for synthetic traffic generation that an individual user generates: A single user can run different applications that may be concurrently active, e.g. WWW browsing while downloading Napster music files. Each application is completely described by its statistical properties. These statistical properties comprise of an alternating process of ON- and OFF-periods with some application specific length or data volume distribution, respectively. During an ON-period, i.e. an application specific connection, the user applies the appropriate application in an active fashion. The interarrival-time between two successive connection starting points of the same application-type and the data volume of each connection are drawn from the parameterized general distributions. Within each ON-period the packet arrival process is completely captured by the packet interarrival-times, drawn according to an application dependent distribution, and the corresponding packet lengths. The overall traffic stream of one user constitutes of the superposition of the packet arrival processes of all application connections within the user's session. Moreover, new users enter the considered system environment according to a session interarrival-time distribution and leave the system after transferring a specific data-volume drawn according to a session volume distribution.

Because of the nature of generally distributed sources, this detailed synthetic traffic model constitutes a non-analytically tractable traffic model. Thus, this traffic model can be utilized as traffic generation component in simulation studies on a per-user basis. Based on this model, performance studies for changing traffic characteristics can easily be conducted by new parameterization of the characteristic distribution for the considered statistical measure. This is easily accomplished by changing the values of expectation and standard deviation of the characteristic distribution of the considered statistical measure. For example, experiments for increasing file sizes can be easily performed by changing the parameters of the connection data volume distribution. In order to get an analytically tractable traffic model that can be integrated as a traffic generation component within analytical models, we identified the batch Markovian arrival process. The utilization of the BMAP is encouraged by the observation that

IP packet lengths follow to a large extent a discrete distribution. The key idea of this aggregated traffic model lies in customizing the batch Markovian arrival process such that the different lengths of IP packets are represented by rewards of the BMAP. Recall, that aggregated traffic models capture the entire traffic stream without explicitly considering individual traffic sources. Thus, the aggregated traffic stream comprises of a sequence of interarrival-times of packet arrivals and packet lengths.

3 Traffic Modeling using the Batch Markovian Arrival Process

3.1 The Batch Markovian Arrival Process

The batch Markovian arrival process (BMAP) belongs to the class of Markov renewal processes. Consider a continuous-time Markov chain (CTMC, [9]) with $(N + 1)$ states $\{0, 1, \dots, N\}$, where the states $\{1, 2, \dots, N\}$ are transient states and 0 is the absorbing state. Moreover, $\boldsymbol{\pi}$ denotes the initial state probability vector of the CTMC. Based on this governing CTMC, the BMAP can be constructed as follows: The CTMC evolves until an absorption in state 0 occurs. The chain is then instantaneously restarted in one of the transient states $\{1, 2, \dots, N\}$. When restarting the BMAP after absorption in a transient state j , the probability for selecting state j is allowed to depend on state i from whom absorption has occurred. Thus, the distribution of the next arrival may depend on the previous history. Furthermore, there may exist multiple paths between states i and j corresponding to different rewards, i.e., batch sizes of arrivals. Due to the additional rewards the BMAP provides a more comprehensive model for representing IP traffic than the MMPP and the MAP, while still being analytically tractable.

Formally, assume the BMAP is in a transient state i for an exponentially distributed time with rate λ_i . When the sojourn time has elapsed, there are $(M + 1)$ possible cases for state transitions: With probability $\mathbf{P}(m)_{i,j}$ the BMAP enters the absorbing state 0 and an arrival of batch size m occurs. Then, the process is instantaneously restarted in state j . Note that the selection of state j ($1 \leq j \leq N$) and batch size m ($1 \leq m \leq M$) is uniquely determined by $\mathbf{P}(m)_{i,j}$. On the other hand, with probability $\mathbf{P}(0)_{i,j}$ the BMAP enters another transient state j , $j \neq i$, without arrivals. We can define $\mathbf{D}(0)_{i,j} = \lambda_i \cdot \mathbf{P}(0)_{i,j}$ for $i \neq j$, $\mathbf{D}(0)_{i,i} = -\lambda_i$, and $\mathbf{D}(m)_{i,j} = \lambda_i \cdot \mathbf{P}(m)_{i,j}$. Here, $\mathbf{D}(0)$ defines the rate matrix of transitions without arrivals, whereas the matrices $\mathbf{D}(m)$ define rate matrices of transitions with arrivals of batch size m ($1 \leq m \leq M$). Summing up $\mathbf{D}(0)$ and $\mathbf{D}(m)$ ($1 \leq m \leq M$) leads to $\mathbf{D} = \mathbf{D}(0) + \sum_{m=1}^M \mathbf{D}(m)$, where \mathbf{D} is the infinitesimal generator matrix of the CTMC underlying the BMAP. Furthermore, the matrix $\mathbf{f}_m(t)$ of probability density functions (pdf) defines probability laws for state changes in the CTMC from i to j with an arrival of batch size m at time t . The matrices $\mathbf{f}_m(t)$ are given by $\mathbf{f}_m(t) = e^{\mathbf{D}(0)t} \cdot \mathbf{D}(m)$.

3.2 Parameter Estimation Procedure

We derived a computational efficient and numerical robust EM (expectation-maximization) algorithm for the parameter estimation process of BMAPs, i.e., estimation of the parameter set ϕ comprising of the probability vector π and the transition rate matrices $\mathbf{D}(0), \dots, \mathbf{D}(M)$. The EM (expectation-maximization) algorithm [3] implements maximum likelihood estimation in case of incomplete data. Such incomplete data can be thought of as partial observations of a larger experiment. In fact, for a BMAP only arrival-times and batch sizes of arrivals, e.g. arrival-times of IP packets and their packet lengths, are observable. All state changes in the governing CTMC not hitting the absorbing state are not observable and, thus, *cannot* be derived from measured trace data.

Formally, suppose that \mathbf{y} is the observable part of a considered experiment. This experiment can be described completely by \mathbf{y} and the non-observable data \mathbf{x} denoted as the *missing data*. Let $\mathcal{L}(\phi, \mathbf{y})$ be the *likelihood* of a parameter set ϕ given the observation \mathbf{y} and let $\mathcal{L}^c(\phi, \mathbf{x}, \mathbf{y})$ be the so-called *complete likelihood* of the parameter set ϕ including the missing data \mathbf{x} . Assume $\mathbf{y} = \{(t_1, b_1), (t_2, b_2), \dots, (t_n, b_n)\}$ is the observed sequence of interarrival-times t_k and the corresponding batch sizes b_k . Define $\Delta t_k = t_k - t_{k-1}$ for $1 \leq k \leq n$. Then, the likelihood of a BMAP with parameter set ϕ is given by:

$$\mathcal{L}(\phi, \mathbf{y}) = \pi \cdot \prod_{k=1}^n \mathbf{f}_{b_k}(\Delta t_k) \cdot \mathbf{1} \quad (1)$$

Recall, that in Eq. (1), π denotes the initial state probability vector of the CTMC, $\mathbf{f}_{b_k}(t)$ defines the matrix of probability density functions, and ϕ is a specific parameter set for the BMAP comprising of π and the transition rate matrices $\mathbf{D}(0), \dots, \mathbf{D}(M)$. The vector $\mathbf{1}$ represents a vector of appropriate dimension comprising of 1s in each entry. Note, that the (logarithm of the) likelihood measures the quality of the estimated parameter set.

The EM algorithm iteratively determines estimates of the missing parameter set ϕ of the BMAP. Denote by $\phi(r)$ the parameter set calculated in the r -th iteration of the EM algorithm. We denote further by \mathbb{P}_ϕ and \mathbb{E}_ϕ the conditional probability and the conditional expectation given the estimate ϕ , respectively. As shown in [3], the estimate

$$\hat{\phi} = \arg \max_{\phi} \left\{ \mathbb{E}_{\phi(r)}(\log \mathcal{L}^c(\phi, \mathbf{x}, \mathbf{y}) | \mathbf{y}) \right\}, \text{ for } r = 0, 1, 2, \dots, \quad (2)$$

satisfies $\mathcal{L}(\hat{\phi}, \mathbf{y}) \geq \mathcal{L}(\phi(r), \mathbf{y})$ and $\phi(r+1) = \hat{\phi}$ is the estimate for the parameter set determined in the $(r+1)$ -th step of the algorithm. This iterative procedure is repeated until a predefined maximum number of iterations is reached or until convergence criteria holds. That is, each component of $\phi(r)$ and $\phi(r+1)$ differs only up to a predefined ε , respectively. The computation of the conditional expectation in (2) is called the E-step whereas the derivation of the maximum in (2) constitutes the M-step of the EM algorithm. As described in [3], the likelihood $\mathcal{L}(\phi, \mathbf{y})$ is highly non-linear in ϕ and is difficult to maximize, while the complete likelihood $\mathcal{L}^c(\phi, \mathbf{x}, \mathbf{y})$ employed in the M-step can often be computed in closed form. This is

the main reason for the widespread use of the EM algorithm. A further advantage of the EM algorithm over other maximum likelihood methods lies in the good convergence behavior of the iterative scheme.

3.3 Effective Computational Formulas for the BMAP

Recall, the observed data in a BMAP is $\{(t_1, b_1), (t_2, b_2), \dots, (t_n, b_n)\}$. The generator of the CTMC $\{X(t): t \geq 0\}$ underlying the BMAP constitutes the missing data. We assume that $N(t)$ is the counting process of the batch sizes for arrivals. Let $\{t_k: 1 \leq k \leq n\}$ be the sequence of arrival-times. Without loss of generality, we assume $t_0 = 0$ and $t_n = T$. Considering the likelihood estimates of the EM algorithm introduced above, we show in the following how to maximize the likelihood for the parameter set of a BMAP.

First of all, we have to define the complete likelihood of the BMAP using the observed data \mathbf{y} and the non-observable data \mathbf{x}

$$\begin{aligned} \mathcal{L}^c(\boldsymbol{\phi}, \mathbf{x}, \mathbf{y}) = & \pi_{i_0(1)} \times \prod_{k=1}^n \prod_{l=0}^{m(k)-1} \lambda_{i_l(k)} \cdot e^{-\lambda_{i_l(k)} \cdot s_l(k)} \cdot \mathbf{P}(0)_{i_l(k), i_{l+1}(k)} \\ & \times \lambda_{i_{m(k)}(k)} \cdot e^{-\lambda_{i_{m(k)}(k)} \cdot s_{m(k)}(k)} \cdot \mathbf{P}(b_k)_{i_{m(k)}(k), i_0(k+1)} \end{aligned} \quad (3)$$

where $\mathbf{P}(m)$ and λ_i are defined as above, $\sum_{j=0}^{m(k)} s_j(k) = \Delta t_k$, and b_k is the batch size of the k -th arrival. The first term of the right hand side of Eq. (3) specifies the probability of starting the CTMC in state $i_0(1)$. For each arrival epoch of the BMAP, the second term describes the transient trajectory up to a state $i_{m(k)}(k)$ from which absorption occurs. The last portion of Eq. (3) represents the transition from the transient state $i_{m(k)}(k)$ to the absorbing state 0 and the restarting of the process in state $i_0(k+1)$ with an arrival of batch size b_k .

In order to simplify the notation in the estimation step, we define the *sufficient statistics* \mathbf{T} , $\mathbf{A}(1), \dots, \mathbf{A}(M)$ and \mathbf{s} as follows. For $1 \leq i, j \leq N$ and $i \neq j$ we define $\mathbf{T}_{i,j}$ as:

$$\mathbf{T}_{i,j} = \#\{t \mid 0 \leq t \leq T, X(t-) = i, X(t) = j, N(t) = N(t-)\} \quad (4)$$

$\mathbf{T}_{i,j}$ is the number of *transient state transitions* from state i to state j without an arrival. For $1 \leq m \leq M$ and $1 \leq i, j \leq N$ we define $\mathbf{A}(m)_{i,j}$ as:

$$\mathbf{A}(m)_{i,j} = \#\{t_k \mid 1 \leq k \leq n, X(t_k-) = i, X(t_k) = j, N(t_k) = N(t_k-) + m\} \quad (5)$$

$\mathbf{A}(m)_{i,j}$ is the number of *absorbing state transitions* from state i to state j with an arrival of batch size m at arrival-times t_k . Finally, for $1 \leq i \leq N$ we define \mathbf{s}_i as:

$$\mathbf{s}_i = \int_0^T I(X(t) = i) dt, \text{ where } I(\cdot) \text{ is the indicator function.} \quad (6)$$

\mathbf{s}_i captures the total time the CTMC resides in state i .

The key idea of these sufficient statistics is to capture the complete likelihood expression in a more intuitive fashion. Typically, the sufficient statistics can be determined by numerically tractable characteristics of the CTMC. These sufficient statistics can easily be used to rewrite and simplify the expression of the complete likelihood in Eq. (3). That is:

$$\mathcal{L}^c(\boldsymbol{\phi}, \mathbf{x}, \mathbf{y}) = \prod_{i=1}^N \pi_i^{I(X(0)=i)} \cdot \prod_{i=1}^N e^{\mathbf{D}(0)_{i,j} \cdot S_i} \cdot \prod_{i=1}^N \prod_{j=1, j \neq i}^N \mathbf{D}(0)_{i,j}^{T_{i,j}} \cdot \prod_{m=1}^M \prod_{i=1}^N \prod_{j=1}^N \mathbf{D}(m)_{i,j}^{A(m)_{i,j}} \quad (7)$$

Intuitively, the second product of Eq. (7) symbolizes the sojourn time of the CTMC for each state i . The third product of Eq. (7) captures the behavior for all transient state transitions between states i and j . Similarly, the last product of Eq. (7) represents the absorbing state transitions between states i and j with arrivals of batch size m .

When adopting (2) to the considered case of a BMAP, we have to recognize that $\mathbf{y} = \{(t_1, b_1), (t_2, b_2), \dots, (t_n, b_n)\}$ is completely characterized by the counting process $N(t)$ introduced above. This leads directly to Eq. (8) with the following abbreviations (9) to (12) for ease of notation. In Eq. (8), the maximization $\hat{\pi}_i$ of π_i is already given in a natural way. Furthermore, it can be shown that using the definition of a BMAP and appropriate partial differentiation, Eq. (8) is maximized by Eq. (13). Additionally to the maximization by partial differentiation, each of the expressions in Eq. (13) utilizes the *maximized sufficient statistics* (9) to (12) in a very intuitive manner. For example, $\hat{\mathbf{D}}(0)_{i,j}$ is the ratio of the total number of transient state transitions between states i and j and the total time spent in state i .

The expressions (8) to (12) in Figure 4 represent the E-step of the EM algorithm, while the expressions in Figure 5 represent the M-Step of the EM algorithm. For ease of notation, we define $\mathbf{R}(k)$ for $k = n+1, \dots, 1$ as $\mathbf{R}(n+1) = \mathbf{1}$ and $\mathbf{R}(k) = \mathbf{f}_{b_k}(\Delta t_k) \cdot \mathbf{R}(k+1)$. Let $\mathbf{1}_i$ denote the i -th unity column vector. Eqs. (10) and (11) can be transformed to integrals over matrix exponentials by means of probability laws. We omit these transformations as well as the detailed evaluation of Eqs. (9) and (12) because of space limitations. Already for an MMPP, the problem with the practical applicability of the EM algorithm for parameter estimation lies in the stable numerical computation of integrals over matrix exponentials as specified in such equations. Ryden [12] proposed a diagonalization method to compute $e^{\mathbf{Q}t}$, but this approach relies on the diagonalization property of the matrix \mathbf{Q} . It is known that decomposition techniques like diagonalization are in general not stable numerical methods for computing matrix exponentials.

For efficient and reliable calculation of (integrals over) matrix exponentials of the transformed equations of (9) to (12), we have derived effective computational formulas based on the randomization technique [5] enhanced by a stable calculation of Poisson probabilities. Again, we omit these formulas because of space limitations. Furthermore, we adopted the scaling procedure proposed in [12] for calculating the sufficient statistics $\hat{\mathbf{T}}_{i,j}$, $\hat{\mathbf{A}}(1)_{i,j}, \dots, \hat{\mathbf{A}}(M)_{i,j}$, \hat{s}_i , $\hat{\pi}_i$, and the likelihood estimate $\mathcal{L}(\boldsymbol{\phi}(r), \mathbf{y})$. This scaling procedure is necessary because these quantities can take extremely small or extremely large values.

E-step:

$$\begin{aligned} & \mathbb{E}_{\phi(r)} \{ \log \mathcal{L}(\phi, \mathbf{x}, \mathbf{y}) | N(u), 0 \leq u \leq T \} \\ &= \sum_{i=1}^N \hat{\pi}_i \cdot \log \pi_i + \sum_{i=1}^N \mathbf{D}(0)_{i,i} \cdot \hat{\mathbf{s}}_i + \sum_{i=1}^N \sum_{j=1, j \neq i}^N \hat{\mathbf{T}}_{i,j} \cdot \log \mathbf{D}(0)_{i,j} + \sum_{m=1}^M \sum_{i=1}^N \sum_{j=1}^N \hat{\mathbf{A}}(m)_{i,j} \cdot \log \mathbf{D}(m)_{i,j} \end{aligned} \quad (8)$$

where:

$$\hat{\pi}_i = \mathbb{P}_{\phi(r)} \{ X(0) = i | N(u), 0 \leq u \leq T \}. \quad (9)$$

$$\hat{\mathbf{s}}_i = \mathbb{E}_{\phi(r)} \{ s_i | N(u), 0 \leq u \leq T \} = \int_0^T \mathbb{P}_{\phi(r)} (X(t) = i | N(u), 0 \leq u \leq T) dt. \quad (10)$$

$$\hat{\mathbf{T}}_{i,j} = \int_0^T \mathbb{P}_{\phi(r)} \{ X(t-) = i, X(t) = j, N(t) = N(t-) | N(u), 0 \leq u \leq T \} dt. \quad (11)$$

$$\hat{\mathbf{A}}(m)_{i,j} = \sum_{k=1}^n \mathbb{P}_{\phi(r)} \{ X(t_k -) = i, X(t_k) = j, N(t_k) = N(t_k -) + m | N(u), 0 \leq u \leq T \}. \quad (12)$$

Figure 4. E-step for parameter estimation of BMAPs**M-step:**

$$\hat{\pi}_i = \pi_i(r) \cdot \mathbf{1}_i \cdot \mathbf{R}(1) / \mathcal{L}(\phi(r), \mathbf{y}), \text{ for } i = 1, \dots, N.$$

$$\hat{\mathbf{D}}(0)_{i,j} = \hat{\mathbf{T}}_{i,j} / \hat{\mathbf{s}}_i \text{ and } \hat{\mathbf{D}}(m)_{i,j} = \hat{\mathbf{A}}(m)_{i,j} / \hat{\mathbf{s}}_i, \text{ for } i, j = 1, \dots, N. \quad (13)$$

$$\hat{\mathbf{D}}(0)_{i,i} = - \sum_{j=1, j \neq i}^N \hat{\mathbf{D}}(0)_{i,j} - \sum_{m=1}^M \sum_{j=1}^N \hat{\mathbf{D}}(m)_{i,j}, \text{ for } i = 1, \dots, N.$$

Figure 5. M-step for parameter estimation of BMAPs

3.4 BMAP Traffic Modeling Framework

As stated above, we customize the batch Markovian arrival process such that different lengths of IP packets are represented by rewards of the BMAP. In order to represent an aggregated traffic stream utilizing the BMAP, we apply the parameter estimation procedure introduced above for a BMAP with N transient states and a maximum batch size of M . The choice of N and M is crucial for an accurate capturing of the interarrival process and the reward process of the aggregated traffic, respectively. Note, that the run-time of the EM algorithm scales linearly with respect to the number of samples n of the considered trace, but is independent of the choice of M . Thus, this modeling approach can be effectively applied for arbitrary packet length (i.e., reward) distributions by an increasing value of M .

The underlying trace file, which constitutes the aggregated traffic, comprises of packet interarrival-times as well as the corresponding packet lengths. Recalling the BMAP

definition, the mapping process of packet lengths to BMAP rewards results in a BMAP parameter set, i.e., π and $\mathbf{D}(0), \dots, \mathbf{D}(M)$, of reasonable size $(M+1)N^2 + N$. We map the packet lengths onto the discrete packet lengths s_m , for $1 \leq m \leq M$, where s_m is the average packet length of all packets of the considered trace comprising of packet lengths between $L \cdot (m-1)/M$ bytes and $L \cdot m/M$ bytes, where L denotes the maximum packet length of the considered IP network. In the case of Ethernet LAN the maximum transfer unit (MTU) of 1500 bytes determines this maximum packet length. Therefore, arrivals with batch size m , $1 \leq m \leq M$, represent packet arrivals with a packet length of s_m bytes. Note, that this mapping process is applied for ease of notation only. Without this mapping, the proposed estimation procedure can be applied for estimating the rate matrices $\mathbf{D}(0), \mathbf{D}(s_1), \dots, \mathbf{D}(s_M)$, where rate matrices $\mathbf{D}(m)$, $m \notin \{0, s_1, \dots, s_M\}$, are empty. The estimation of these matrices obviously requires the same computational effort. Moreover, the estimated matrices are identical to the matrices $\mathbf{D}(0), \dots, \mathbf{D}(M)$, which are computed when the mapping process is applied.

4 Comparative Study of Aggregated IP Traffic Modeling

As stated above, the synthetic traffic model comprising of generally distributed sources is not analytically tractable. Thus, it can be employed for simulation studies only. To overcome this restriction, we utilize the traffic modeling framework introduced in the previous section with a trace file comprising of 1,500,000 samples (measured 10.00 a.m. 13 December 2000 at the dial-up modem/ISDN link). This trace file comprises of packet interarrival-times and the corresponding packet lengths. Based on this trace file the BMAP parameter estimation procedure is applied for a 3-state BMAP $N=3$ with a maximum batch size of $M=3$. Recall, the choice of M is crucial for the mapping process of packet lengths to BMAP rewards and corresponds to, but is not restricted by the fact that a large amount of packets comprise of three different packet lengths (see Table 1 and Figure 3). Recalling the mapping process of packet lengths to BMAP rewards, arrivals with batch size m , $1 \leq m \leq M$, represent packet arrivals with a packet length of s_m bytes. The average packet lengths s_m ($M=3$) of our measurements are as follows: $s_1 = 94$ bytes, $s_2 = 575$ bytes, and $s_3 = 1469$ bytes. The considered estimation procedure is quite effective and requires less than 20 minutes of CPU time on a Pentium III PC with 128 MB of main memory.

The following shows the effectiveness of our BMAP modeling approach, compared with MMPP and Poisson process, by means of visual inspection of sample paths over multiple time-scales, by presenting important statistical properties, by formal analysis of traffic burstiness as well as by investigations of queuing behavior. Figure 6 and 7 plot sample paths of the measured traffic (Figure 6, left) compared with the sample paths of the aggregated traffic streams of the customized BMAP (Figure 6, right), the MMPP (Figure 7, left), and the Poisson process (Figure 7, right), respectively. For sample-path construction of the MMPP

and the Poisson process, we associate the average packet length of all IP packets comprising of 315 bytes to the arrival-times of the MMPP and the Poisson process. The aggregated traffic streams of the MMPP and the Poisson utilized for sample-path construction comprise of the same number of samples as the measured trace file. Note, that the parameter matrices $\mathbf{D}(0)$ and $\mathbf{D}(1)$ of the MMPP have also been estimated by application of the EM algorithm for BMAP. This is accomplished by restriction of $\mathbf{D}(1)$ to diagonal entries that are associated with the state-dependent Poisson arrival-rates of the MMPP. The arrival-rate of the Poisson process is naturally given by the mean arrival-rate of the measured trace file. In order to show the effectiveness of our approach these sample paths are plotted on four different time-scales, i.e. 0.001 sec, 0.01 sec, 0.1 sec, and 1.0 sec. Figure 6 and 7 evidently show that the customized BMAP authentically captures the average transferred data volume per time unit and exhibits traffic burstiness over multiple time-scales in the considered scenario. Moreover, these sample paths show the clear advantage of the customized BMAP over the MMPP and the Poisson process, which fail to capture the original sample path over almost all time-scales.

Table 2 presents additionally statistical properties for the data rates of the measured traffic, the BMAP, the MMPP, and the Poisson process, on different time-scales in terms of mean, standard deviation, skewness, and kurtosis. Recall, that the mean gives the center of the distribution and the standard deviation measures the dispersion about the mean. The third moment about the mean measures skewness, the lack of symmetry, while the fourth moment measures kurtosis, the degree to which the distribution is peaked. In Table 2, skewness and kurtosis are standardized by an appropriate power of the standard deviation. We observe, that mean and standard deviation of the measured traffic and the customized BMAP perform quite similar over the considered time-scales, with exception of the BMAPs standard deviation on

time unit in sec.		mean	standard deviation	skewness	kurtosis
0.001	measured traffic	27.63	153.02	6.72	52.77
	customized BMAP	27.70	157.90	8.12	84.34
	MMPP	27.65	118.39	5.58	42.73
	Poisson process	27.65	93.37	3.38	14.46
0.01	measured traffic	276.27	697.46	4.21	29.03
	customized BMAP	277.01	662.32	3.82	23.14
	MMPP	276.47	491.13	2.81	13.84
	Poisson process	276.49	295.82	1.07	4.14
0.1	measured traffic	2762.64	2240.81	1.44	6.15
	customized BMAP	2770.10	2191.71	1.27	5.15
	MMPP	2764.61	1672.45	0.90	4.21
	Poisson process	2764.86	933.00	0.33	3.11
1	measured traffic	27621.50	10954.80	0.60	3.21
	customized BMAP	27697.30	6929.95	0.41	3.19
	MMPP	27643.70	5413.94	0.28	3.05
	Poisson process	27643.70	2972.29	0.03	2.94

Table 2. Statistical properties of data rates on different time-scales

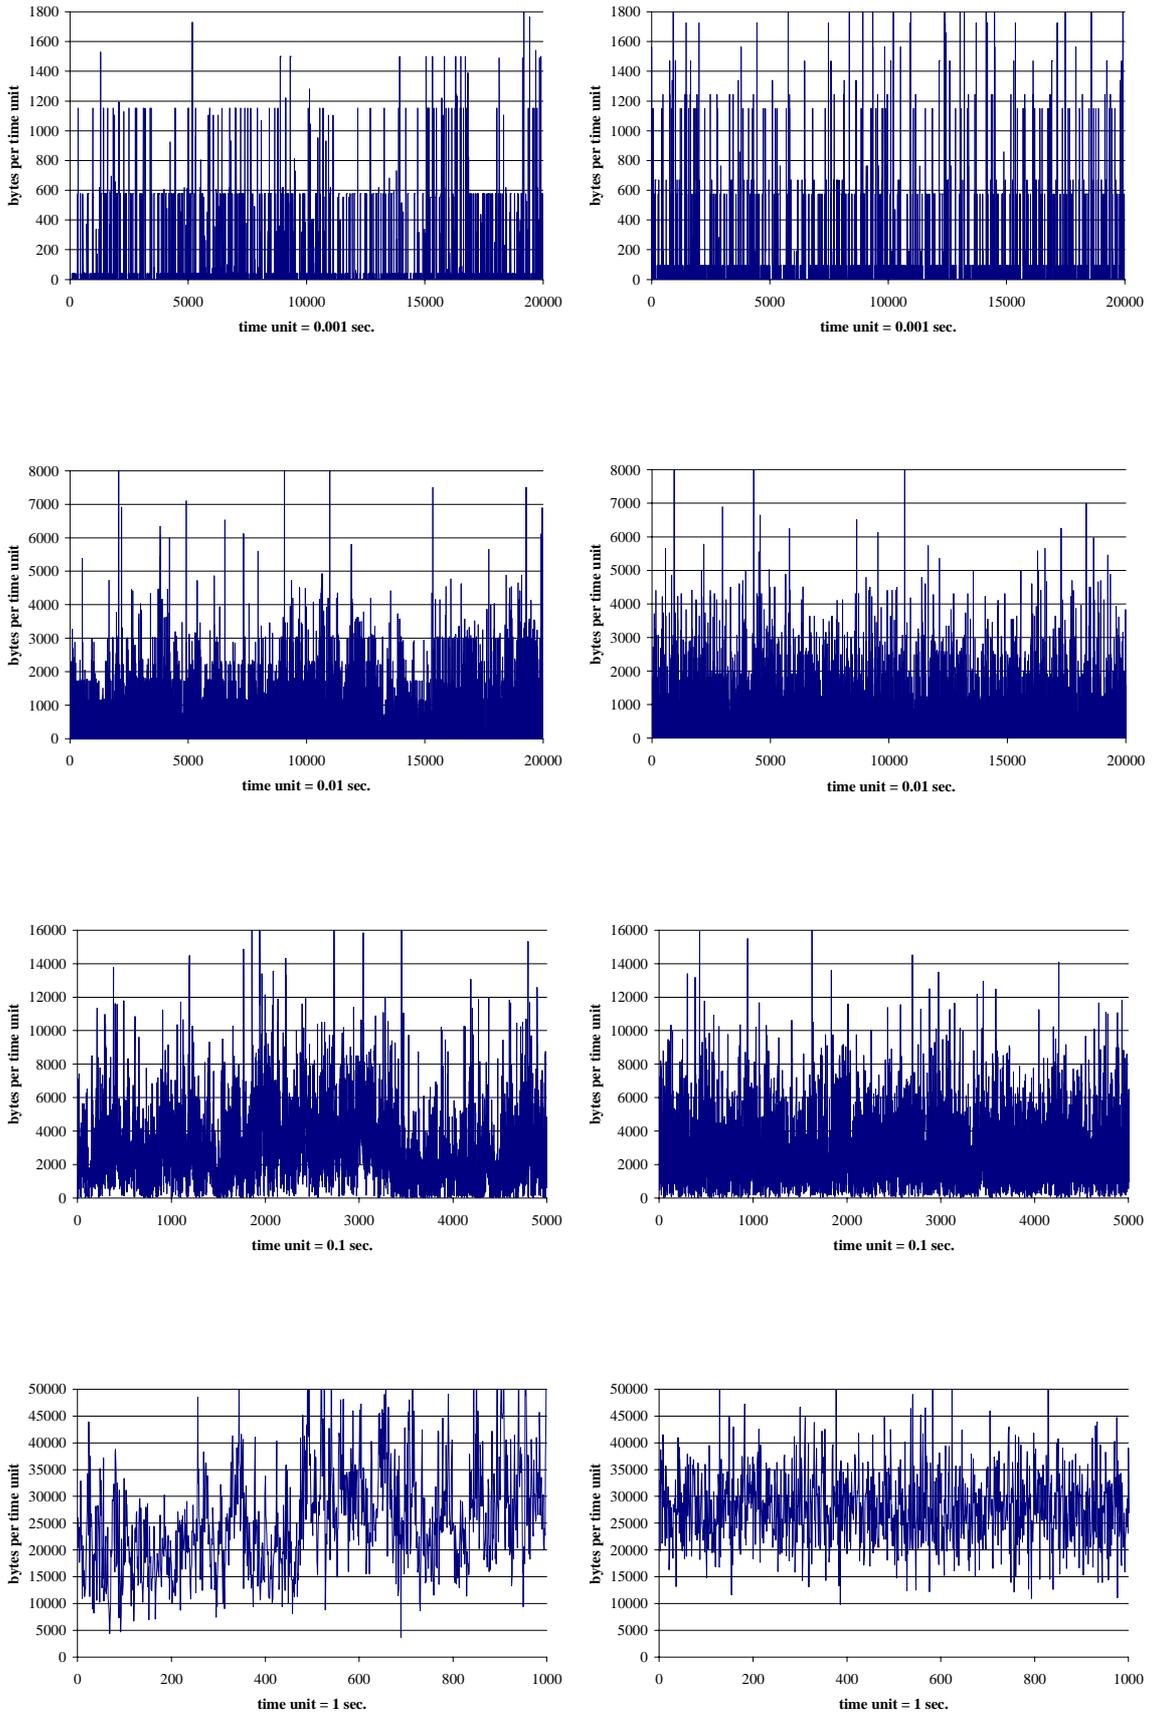


Figure 6. Sample paths of the measured traffic (left) and the customized BMAP (right) at different time-scales

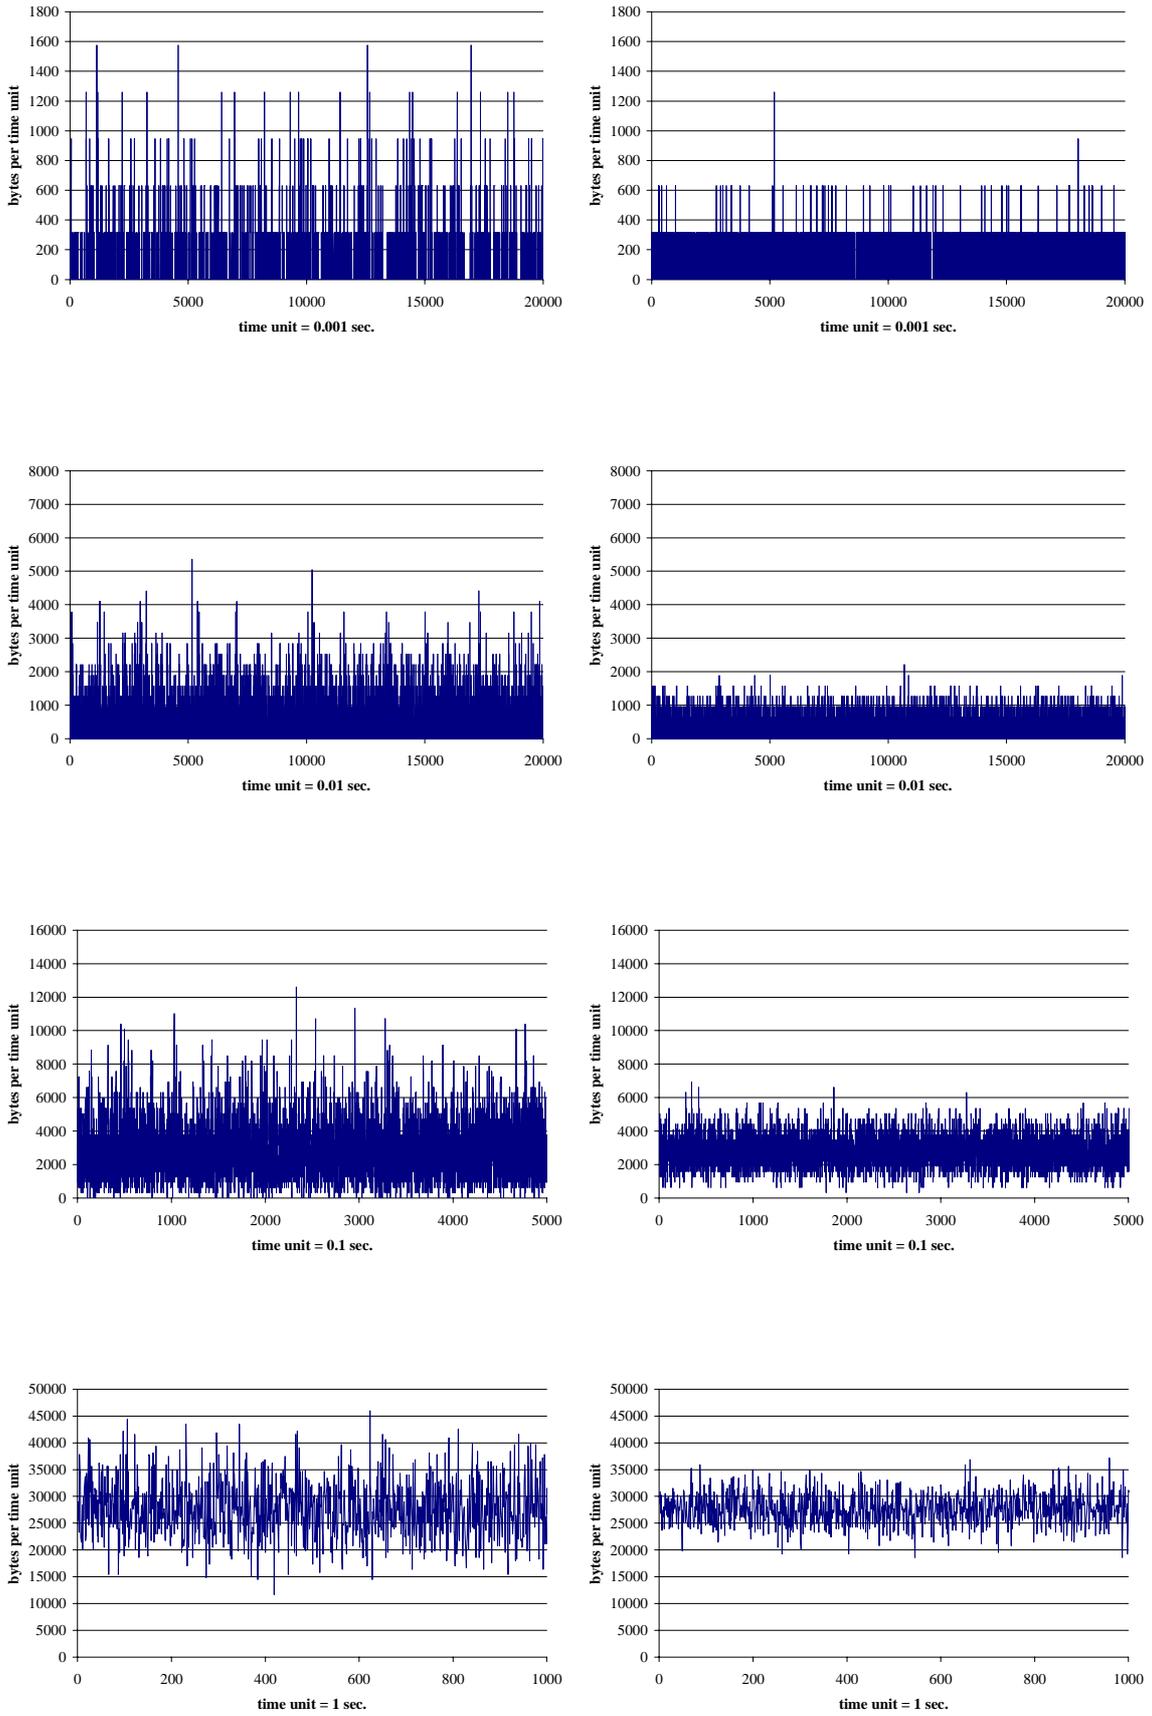


Figure 7. Sample paths of the MMPP (left) and the Poisson process (right) at different time-scales

the largest time-scale. The skewness of the measured traffic is quite similar on medium time-scales, i.e., 0.01 sec and 0.1 sec, while the customized BMAP overestimates the skewness on the smallest time-scale and underestimates it on the largest time-scale. Furthermore, the last column of Table 2 indicates, that kurtosis, i.e., peakedness, is well captured on the three largest time-scales, whereas the BMAP significantly exceeds the measured traffic on the smallest time-scale. This is, because on the smallest time-scale the various packet lengths of the measured traffic cannot be represented exactly by only three ($M = 3$) different reward values, i.e., packet lengths. This effect diminishes with increasing value of M . Moreover, Table 2 evidently shows, that the MMPP as well as the Poisson process are clearly inferior compared with the customized BMAP and loosely capture standard deviation, skewness and kurtosis, over all considered time-scales.

These observations are emphasized by the analysis of traffic burstiness, which can be expressed in terms of the Hurst parameter H . Figure 8 plots the R/S statistics [15] of the measured traffic, the customized BMAP, the MMPP as well as the Poisson process. The degree of traffic burstiness H , can easily derived by the slopes of linear regression plots of the R/S statistics. As expected, the Poisson process ($H = 0.5558$) fails to capture the traffic burstiness, while the MMPP ($H = 0.6408$) and the customized BMAP ($H = 0.6418$) both indicate a significant amount of traffic burstiness compared with the Hurst parameter of the measured traffic ($H = 0.6785$).

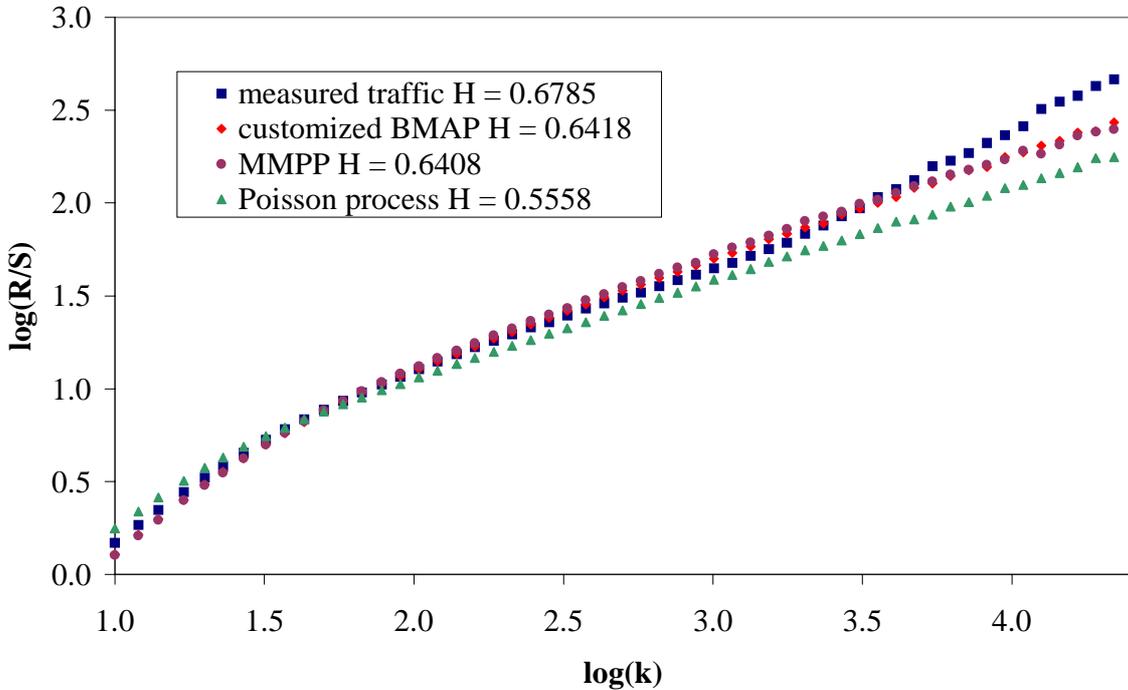


Figure 8. R/S statistic plot of the trace and the analytically tractable models

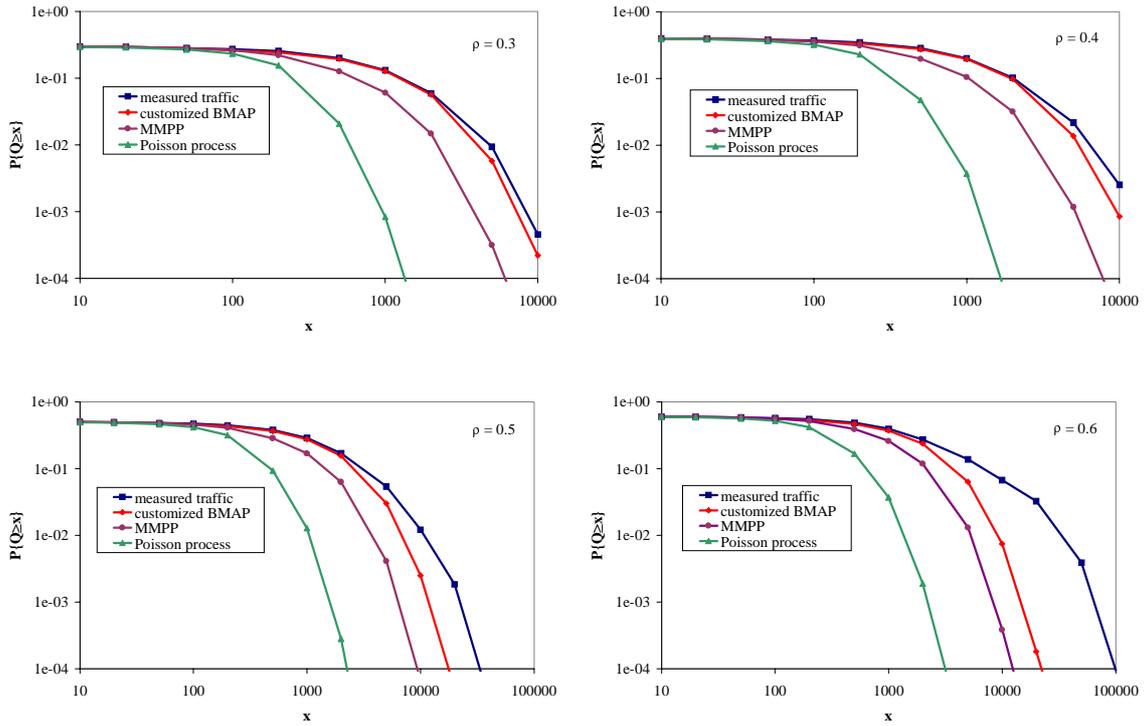


Figure 9. Complement distribution of queue length Q of single server queue with deterministic service time for different traffic intensities ρ

The practical applicability of our BMAP modeling approach can be emphasized by the analysis of the queuing performance. As proposed in [1], we utilize a simple queuing model with deterministic service time and unlimited capacity for investigations of the complement distribution of the queue length. Figure 9 depicts the complement distribution of the queue length Q of the BMAP/D/1 queuing system, the MMPP/D/1 queuing system and the M/D/1 queuing system (using the Poisson process), compared with the simulations performed with the measured traffic for different traffic intensities ρ . It is obvious that the BMAP model shows a similar behavior in terms of queuing performance for low traffic intensities, i.e., $\rho = 0.3$ and $\rho = 0.4$. For traffic intensities of $\rho = 0.5$ and $\rho = 0.6$ the customized BMAP matches the distribution of the measured traffic up to medium load. As expected, the Poisson process loosely performs for all considered traffic intensities. Obviously, the MMPP outperforms the Poisson process, but is significantly inferior in capturing the complement distribution of the queuing length, compared with the customized BMAP for all considered traffic intensities.

Conclusions

We present a detailed analysis of IP traffic measurements recently conducted at an Internet service provider (ISP) dial-up link and derive parameterized general distributions for session-level, connection-level, and packet-level characteristics. Moreover, we observe that almost all IP packets comprise of three dominating packet lengths. Based on these observations, we introduce an aggregated traffic model for IP networks that is both analytically tractable and closely captures the statistics of the measured traffic data. The key idea of this aggregated traffic model lies in customizing the batch Markovian arrival process such that these different lengths of IP packets are represented by rewards (i.e., batch sizes of arrivals) of the BMAP. We introduce an efficient method for estimating the parameters of a BMAP with the EM algorithm. In fact, we present efficient computational formulas for the E-step of the EM algorithm and show how to utilize the EM algorithm for the effective parameter estimation of BMAPs. In order to demonstrate the advantages of the BMAP modeling approach over other widely used analytically tractable models, we compare the customized BMAP with the MMPP and the Poisson process by means of visual inspection of sample paths over four different time-scales, by presenting important statistical properties, by formal analysis of traffic burstiness, and by queuing system analysis.

References

- [1] S. Andersson and T. Ryden, Maximum Likelihood Estimation of a Structured MMPP with Applications to Traffic Modeling, *Proc. 13th ITC Specialist Seminar on Measurement and Modeling of IP Traffic, Monterey CA*, 20.1-20.10, 2000.
- [2] G.R. Ash, Traffic Engineering & QoS Methods for IP-, ATM-, & TDM-Based Multiservice Networks, *Internet Draft draft-ietf-tewg-qos-routing-01.txt*, 2001.
- [3] A.P. Dempster, N.M. Laird, and D.B. Rubin, Maximum Likelihood from Incomplete Data via the EM Algorithm, *Journal of the Royal Statistical Society* **39**, 1-38, 1976.
- [4] W. Fischer and K. Meier-Hellstern, The Markov-modulated Poisson Process (MMPP) Cookbook, *Performance Evaluation* **18**, 149-171, 1993.
- [5] D. Gross and D.R. Miller, The Randomization Technique as a Modeling Tool and Solution Procedure for Transient Markov Processes, *Operations Research* **32**, 345-361, 1984.
- [6] J. Kilpi and I. Norros, Call Level Traffic Analysis of a Large ISP, *Proc. 13th ITC Specialist Seminar on Measurement and Modeling of IP Traffic, Monterey CA*, 6.1-6.9, 2000.

- [7] H. Kruse, M. Allman, J. Griner, S. Ostermann, and E. Helvey, Satellite Network Performance Measurements Using Simulated Multi-User Internet Traffic, *Proc. 7th Int. Conf. on Telecommunication Systems*, 1999.
- [8] S. Ledesma and D. Liu, Synthesis of Fractional Gaussian Noise using Linear Approximation for Generating Self-similar Network Traffic, *Computer Communication Review* **30**, 4-17, 2000.
- [9] D. M. Lucantoni, New Results on the Single Server Queue with a Batch Markovian Arrival Process, *Comm. in Statistics: Stochastic Models* **7**, 1-46, 1991.
- [10] S. McCreary and K.C. Claffy, Trends in Wide Area IP Traffic Patterns: A View from Ames Internet Exchange, *Proc. 13th ITC Specialist Seminar on Measurement and Modeling of IP Traffic*, Monterey CA, 1.1-1.12, 2000.
- [11] C. J. Nuzman, I. Saniee, W. Sweldens, and A. Weiss, A Compound Model for TCP Connection Arrivals, *Proc. 13th ITC Specialist Seminar on Measurement and Modeling of IP Traffic*, Monterey CA, 1-9, 2000.
- [12] T. Ryden, An EM Algorithm for Parameter Estimation in Markov Modulated Poisson Processes, *Computational Statistics and Data Analysis* **21**, 431-447, 1996.
- [13] P. Skelly, M. Schwartz, and S. Dixit, A Histogram-Based Model for Video Traffic Behavior in an ATM Multiplexer, *IEEE Trans. on Networking* **1**, 446-459, 1993.
- [14] K. Thompson, G.J. Miller, and R. Wilder, Wide-Area Internet Traffic Patterns and Characteristics, *IEEE Network Magazine* **11**, 10-23, November/December 1997.
- [15] W. Willinger, V. Paxson, and M.S. Taqqu, Self-similarity and Heavy Tails: Structural Modeling of Network Traffic, *In: A Practical Guide to Heavy Tails*, 27-53, Chapman & Hall 1998.
- [16] T. Yoshihara, S. Kasahara, and Y. Takahashi, Practical Time-Scale Fitting of Self-Similar Traffic with Markov-Modulated Poisson Process, *Telecommunication Systems* **17**, 185-211, 2001.