

A Hierarchical Approach based on MMPPs for Modeling Self-Similar Traffic over Multiple Time Scales

António Nogueira[†], Paulo Salvador[†], Rui Valadas[†], António Pacheco[‡]

[†]University of Aveiro / Institute of Telecommunications Aveiro
Campus de Santiago, 3810-193 Aveiro, Portugal

e-mail: nogueira@det.ua.pt; salvador@av.it.pt; rv@det.ua.pt

[‡]Instituto Superior Técnico - UTL

Department of Mathematics and CEMAT
Av. Rovisco Pais, 1049-001 Lisboa, Portugal
e-mail: apacheco@math.ist.utl.pt

Abstract

Efficient traffic engineering of IP networks requires the knowledge of the main characteristics of the supported traffic. Several studies have shown that IP network traffic may exhibit properties of burstiness, self-similarity and/or long-range dependence, with significant impact on network performance. In this work, we propose a Markov Modulated Poisson Process (MMPP), and its associated parameter fitting procedure, that is able to incorporate these characteristics over multiple time scales. This is accomplished through a hierarchical construction procedure that, starting from a MMPP that matches the distribution of packet counts at the coarsest time scale, successively decomposes each MMPP state into new MMPPs, that incorporate a more detailed description of the distribution at finer time scales. The traffic process is then represented by a MMPP equivalent to the constructed hierarchical structure. The accuracy of the fitting procedure is evaluated by comparing the Hurst parameter, the probability mass function at each time scale and the queuing behavior (as assessed by the loss probability and average waiting time), corresponding to the measured and to synthetic traces generated from the inferred models. Several measured traffic traces exhibiting self-similar behavior are considered: the well-known pOct Bellcore trace, a trace of aggregated IP WAN traffic, and a trace corresponding to the popular file sharing application Kazaa. Our results show that the proposed model and parameter fitting procedure are very effective in matching the main characteristics of the measured traces over the different time scales present in data.

keywords: Traffic modeling, self-similar, time scale, Markov Modulated Poisson Process.

1 Introduction

Traffic characterization and modeling comprise important steps towards understanding and solving performance-related problems in future IP networks. An efficient design and control of IP networks needs to take into account the main characteristics of the supported traffic, and therefore accurate and detailed measurements need to be carried out. Traffic modeling refers to the construction of (usually stochastic) models that capture the most important statistical properties of the measured data. Since the work by Leland *et al.* [1] several studies have shown that network traffic may exhibit properties of burstiness, self-similarity and/or long-range dependence (LRD) [1, 2, 3, 4, 5, 6, 7], which have significant impact on network performance.

Burstiness is a traffic behavior showing noticeable periods with arrivals above the mean (bursts) and self-similarity refers to the replication of statistical characteristics over a wide range of time scales. Models like the fractional Gaussian noise (fGN) and the fractional autoregressive integrated moving average (fARIMA) have been proposed to capture burstiness and self-similarity but there is still a lack of analytical results, e.g., to assess the queuing behavior.

In general, self-similarity implies LRD, and vice-versa. The impact of LRD on network performance has been addressed by several authors. References [4, 8, 9, 10], for example, study the case of a single queue and conclude that the buffer occupancy is not affected by autocovariance lags that are beyond the so-called critical time scale (CTS) or correlation horizon (CH), which depends on system parameters such as the buffer capacity. Similar conclusions were observed for the case of tandem queues in [11]. Thus, matching the LRD is only required within the time scales specific to the system under study. One of the consequences of this result is that more traditional traffic models, such as Markov Modulated Poisson Processes (MMPPs), can still be used to model traffic exhibiting LRD. Moreover, the use of MMPPs benefits from the existence of several tools for calculating the queuing behavior and the effective bandwidths.

In this work, we consider discrete-time MMPPs (dMMPPs) instead of continuous-time MMPPs, since they are more natural model for data corresponding to the number of arrivals (packet counts) in a sampling interval. Note that discrete-time and continuous-time MMPPs are basically interchangeable (through a simple parameter rescaling) as models for arrival processes, whenever the sampling interval used for the discrete-time version is small compared with the average sojourn times in the states of the modulating Markov chain.

In this paper we propose a dMMPP traffic model, and its associated parameter fitting procedure, that is able to incorporate traffic characteristics of different time scales. This is accomplished through a construction procedure that successively decomposes dMMPP states into new dMMPPs, thus refining the traffic process by incorporating the characteristics offered by finer time scales. We start at the largest time scale by inferring a dMMPP that matches the probability mass function (PMF) of this time scale. At the next finer time scale, each dMMPP state is decomposed into a new dMMPP that matches the contribution of this time scale to the PMF of the state it descends from. In this way, a child dMMPP provides a detailed description of its parent state PMF. This refinement process is iterated until a pre-defined number of time scales are integrated. Finally, a dMMPP incorporating this hierarchical structure is derived. The number of states of each dMMPP is not fixed a priori; it is determined as part of the fitting procedure. The accuracy of the fitting procedure is evaluated by applying it to several measured traffic traces that exhibit self-similar behavior: the well-known pOct Bellcore trace, a trace of aggregated IP WAN traffic, and a trace corresponding to the file sharing application Kazaa. This application was selected due to its present popularity in the Internet. We compare the PMF at each time scale, and the queuing behavior (as assessed by the loss probability and average waiting time), corresponding to the measured and to synthetic traces generated from the inferred models. Our results show that the proposed fitting method is very effective in matching the PMF at the various time scales and leads to an accurate prediction of the queuing behavior.

Several fitting procedures have been proposed in the literature for estimating the parameters of MMPPs from empirical data ([12, 13, 14, 15, 16, 17, 18, 19, 20, 21], among others). However, most procedures only apply to 2-MMPPs (e.g. [12, 14, 15, 18]). This model can capture traffic burstiness but the number of states is not enough to reproduce variability over a wide range of time scales. On the other hand, the fitting procedures for MMPPs with an arbitrary number of states mainly concentrate on matching first- and/or second-order statistics, without addressing directly the issue of modeling over multiple time scales [13, 16, 17, 19, 21]. Yoshihara *et al.* [20] developed a fitting method for self-similar traffic based on the superposition of 2-MMPPs, that matches the variance at each time scale. In this way, the resulting MMPP reproduces the variance-scale curve characteristic of self-similar processes. Our contribution is to develop a procedure that matches the complete distribution at each time scale (and not only the variance) in order to reproduce accurately self-similar behavior.

The paper is organized as follows. Section 2 introduces self-similarity and long-range dependence, motivating the need for a traffic model that matches the different time scales of the data. Section 3 gives the required background on MMPPs. Section 4 describes the proposed model and Section 5 presents the various steps of the parameter fitting procedure. Section 6 briefly describes the data traces used in the numerical evaluation and in Section 7 we discuss the obtained results. Finally, Section 8 presents the main conclusions.

2 Self-similarity, long-range dependence, and time scales

Consider the continuous-time process $Y(t)$ representing the traffic volume (e.g. in bytes) from time 0 up to time t and let $X(t) = Y(t) - Y(t-1)$ be the corresponding increment process (e.g. in bytes/second). Consider also the sequence $X^{(m)}(k)$ which is obtained by averaging $X(t)$ over non-overlapping blocks of length m , that is

$$X^{(m)}(k) = \frac{1}{m} \sum_{i=1}^m X((k-1)m + i), k = 1, 2, \dots \quad (1)$$

The fitting procedure developed in this work will be based on the aggregated processes $X^{(m)}(k)$.

We start by introducing the notion of distributional self-similarity. $Y(t)$ is exactly self-similar when it is equivalent, in the sense of finite-dimensional distributions, to $a^{-H}Y(at)$, for all $t > 0$ and $a > 0$, where H ($0 < H < 1$) is the Hurst parameter. Clearly, the process $Y(t)$ can not be stationary. However, if $Y(t)$ has stationary increments then again $X(k) = X^{(1)}(k)$ is equivalent, in the sense of finite-dimensional distributions, to $m^{1-H}X^{(m)}(k)$. This illustrates that a traffic model developed for fitting self-similar behavior must preferably enable the matching of the distribution on several time scales.

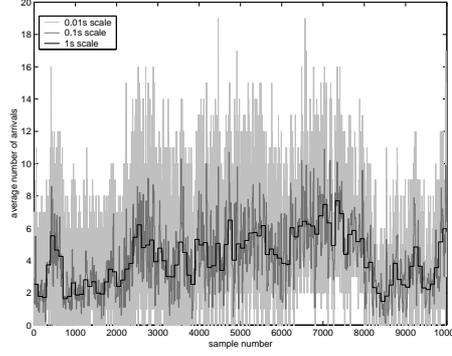


Figure 1: LRD processes exhibit fluctuations over a wide range of time scales (Example: trace pOct).

Long-range dependence is associated with stationary processes. Consider now that $X(k)$ is second-order stationary with variance σ^2 and autocorrelation function $r(k)$. Note that, in this case, $X^{(m)}(k)$ is also second-order stationary. The process $X(k)$ has long-range dependence (LRD) if its autocorrelation function is non-summable, that is, $\sum_n r(n) = \infty$. Intuitively, this means that the process exhibits similar fluctuations over a wide range of time scales. Taking the case of the pOct Bellcore trace, it can be seen in Figure 1 that the fluctuations over the 0.01, 0.1 and 1s time scales are indeed similar.

Equivalently, one can say that a stationary process is LRD if its spectrum diverges at the origin, that is $f(v) \sim c_f |v|^{-\alpha}, v \rightarrow 0$. Here, α is a dimensionless scaling exponent, that takes values in $(0, 1)$; c_f takes positive real values and has dimensions of variance. On the other hand, a short range dependent (SRD) process is simply a stationary process which is not LRD. Such a process has $\alpha = 0$ at large scales, corresponding to white noise at scales beyond the so-called characteristic scale or correlation horizon. The Hurst parameter H is related with α by $H = (\alpha + 1)/2$.

There are several estimators of LRD. In this study we use the semi-parametric estimator developed in [22]. Here, one looks for alignment in the so-called Logscale Diagram (LD), which is a log-log plot of the variance estimates of discrete wavelet transform coefficients, against scale, complete with confidence intervals about these estimates at each scale. It can be thought of as a spectral estimator where large scale corresponds to low frequency. Traffic is said to be LRD if, within the limits of the confidence intervals, the log of the variance estimates fall on a straight line, in a range of scales from some initial value j_1 up to the largest one present in data and the slope of the straight line, which is an estimate of the scaling exponent α , lies in $(0, 1)$.

There is a close relationship between long-range dependent and self-similar processes. In fact, if $Y(t)$ is self-similar with stationary increments and finite variance then $X(k)$ is long-range dependent, as long as $\frac{1}{2} < H < 1$. The process $X(k)$ is said to be exactly second-order self-similar ($\frac{1}{2} < H < 1$) if

$$r(n) = 1/2 \left[(n+1)^{2H} - 2n^{2H} + (n-1)^{2H} \right] \quad (2)$$

for all $n \geq 1$, or is asymptotically self-similar if

$$r(n) \sim n^{-(2-2H)} L(n) \quad (3)$$

as $n \rightarrow \infty$, where $L(n)$ is a slowly varying function at infinity. In both cases the autocovariance decays hyperbolically, which indicates LRD. Any asymptotically second-order self-similar process is LRD, and vice-versa.

3 Markov Modulated Poisson Processes

The discrete-time Markov Modulated Poisson Process (dMMPP) is the discrete-time version of the popular (continuous-time) MMPP and may be regarded as a Markov random walk where the increments in each instant have a Poisson distribution whose parameter is a function of the state of the modulator Markov chain. More precisely, the (homogeneous) Markov chain $(Y, J) = \{(Y_k, J_k), k = 0, 1, \dots\}$ with state space $\mathcal{I}N_0 \times S$ is a dMMPP if and only if for $k = 0, 1, \dots$,

$$P(Y_{k+1} = m, J_{k+1} = j | Y_k = n, J_k = i) = \begin{cases} 0 & m < n \\ p_{ij} e^{-\lambda_i} \frac{\lambda_i^{m-n}}{(m-n)!} & m \geq n \end{cases} \quad (4)$$

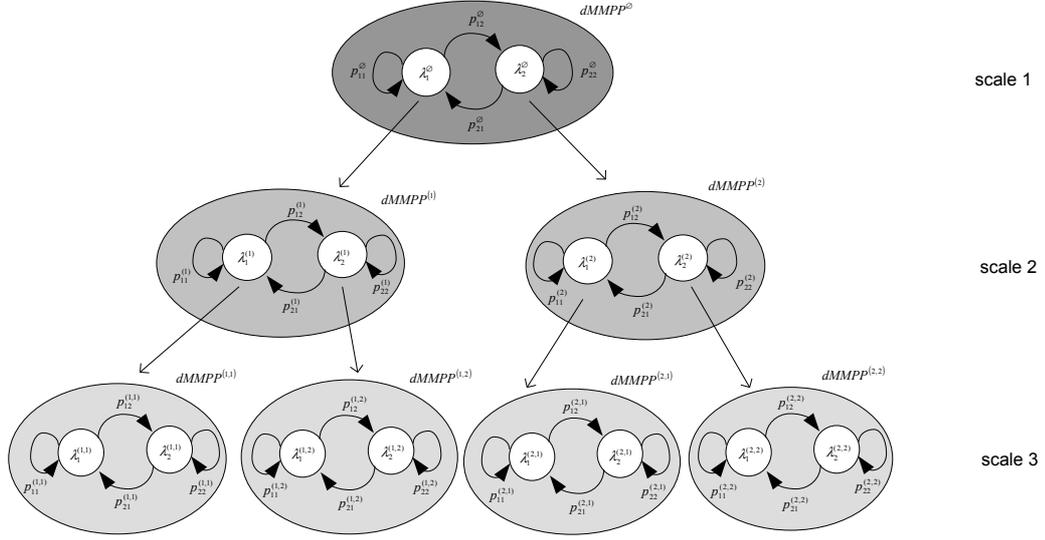


Figure 2: Illustration of the dMMPP construction.

for all $m, n \in \mathbb{N}_0$ and $i, j \in S$, with $\lambda_i, i \in S$, being nonnegative real constants and $\mathbf{P} = (p_{ij})$ being a stochastic matrix. Note that the distribution of $Y_{k+1} - Y_k$ given $J_k = j$ is Poisson with mean λ_j , so that λ_j represents the mean increment of the process Y when the modulating Markov chain is in state j .

Whenever (8) holds, we say that (Y, J) is a dMMPP with set of modulating states S and parameter (matrices) \mathbf{P} and $\mathbf{\Lambda}$, and write

$$(Y, J) \sim \text{dMMPP}_S(\mathbf{P}, \mathbf{\Lambda}) \quad (5)$$

where $\mathbf{\Lambda} = (\lambda_{ij}) = (\lambda_i \delta_{ij})$. The matrix \mathbf{P} is the transition probability matrix of the modulating Markov chain J , whereas $\mathbf{\Lambda}$ is the matrix of Poisson arrival rates. If S has cardinality r , we say that (Y, J) is a dMMPP of order r (dMMPP $_r$). When, in particular, $S = \{1, 2, \dots, r\}$ for some $r \in \mathbb{N}$, then

$$\mathbf{P} = \begin{bmatrix} p_{11} & p_{12} & \dots & p_{1r} \\ p_{21} & p_{22} & \dots & p_{2r} \\ \dots & \dots & \dots & \dots \\ p_{r1} & p_{r2} & \dots & p_{rr} \end{bmatrix} \quad \text{and} \quad \mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \lambda_r \end{bmatrix} \quad (6)$$

and we write simply that $(Y, J) \sim \text{dMMPP}_r(\mathbf{P}, \mathbf{\Lambda})$. The stationary distribution of J is denoted by $\pi = [\pi_1 \pi_2, \dots, \pi_r]$.

4 Proposed model

The goal of this work is to propose a dMMPP model that is able to incorporate traffic characteristics of different time scales. Specifically, we work with the PMF of the packet counts at each time scale. This is accomplished through a construction procedure that successively decomposes dMMPP states into new dMMPPs, thus refining the traffic process by incorporating the characteristics offered by finer time scales. We start at the largest time scale, by inferring a dMMPP that matches the PMF of this time scale. As part of the parameter fitting procedure, each time interval of the data sequence is assigned to a dMMPP state; in this way, a PMF can be associated with each dMMPP state. At the next finer time scale, each dMMPP state is decomposed into a new dMMPP that matches the contribution of this time scale to the PMF of the state it descends from. In this way, a child dMMPP provides a description of its parent state PMF. This refinement process is iterated until a pre-defined number of time scales are integrated. Finally, a dMMPP incorporating this hierarchical structure is derived.

We consider that the number of time scales, L , is fixed *a priori*. Time scales will be numbered in an increasing way, from $l = 1$ (corresponding to the largest time scale) to $l = L$ (corresponding to the smallest time scale). The construction process can be described through a tree where, except for the root node, each tree node corresponds to a dMMPP state and each tree level to a time scale. A dMMPP state will be represented by a vector indicating the path in the tree from its higher level ancestor (i.e. the state it descends from at the

largest scale, $l = 1$) to itself. Thus, a state at time scale l will be represented by $\vec{s} = (s_1, s_2, \dots, s_l)$, $s_i \in \mathcal{I}$. Each dMMPP will be represented by the state that generated it (i.e. its parent state). We let $\text{dMMPP}^{\vec{s}}$ denote the dMMPP generated by state \vec{s} and $\{1, 2, \dots, N_{\vec{s}}\}$ the set of corresponding states, where $N_{\vec{s}}$ is its number of states. The root node of the tree corresponds to a virtual state, denoted by $\vec{s} = \emptyset$, that is used to represent the dMMPP of the largest time scale, $l = 1$. This dMMPP will be called the root dMMPP.

Thus, the dMMPP states in the tree are characterized by

$$\vec{s} = (s_1, s_2, \dots, s_l), l \in \mathcal{I} \quad (7)$$

with $s_{i+1} \in \{1, 2, \dots, N_{\vec{s}_i}\}$, $i = 0, 1, \dots, l-1$; here, \vec{s}_j denotes the sub-vector of \vec{s} given by (s_1, s_2, \dots, s_j) , with $j < |\vec{s}|$, and $\vec{s}_0 = \emptyset$, where $|\vec{s}|$ denotes the length of vector \vec{s} . Note that, using this notation, a vector \vec{s} can either represent state \vec{s} or the dMMPP generated by \vec{s} . Also, the time scale of $\text{dMMPP}^{\vec{s}}$ is $|\vec{s}| + 1$.

Figure 2 illustrates the decomposition process for the simple case of three time scales and two-state dMMPPs.

5 Inference Procedure

The inference procedure is represented schematically in the flowchart of Figure 3, where the following main steps can be identified:

(i) calculation of the data sequences (corresponding to the average number of arrivals per time interval) for each time scale, starting with the smallest one and going through an aggregation process up to the largest one.

(ii) inference of the dMMPP at the largest time scale, $l = 1$, that matches the empirical PMF at this time scale.

(iii) for all other time scales, in increasing order, $l = 2, \dots, L-1$, and for each parent dMMPP state, identification of the time intervals assigned to the state, calculation of the corresponding PMF and inference of the dMMPP that matches the contribution of the time scale to the PMF of the state;

(iv) finally, calculation of matrices \mathbf{A} and \mathbf{P} of the dMMPP incorporating the previous hierarchical structure. Note that the dimensions of all dMMPPs are computed as part of the fitting procedure. We will now describe in detail the various steps of the inference method.

5.1 Calculation of the data aggregates

Having defined the time interval at the smallest time scale, Δt , the number of time scales, L , and the level of aggregation, a , the aggregation process starts by computing the data sequence corresponding to the average number of arrivals in intervals of length Δt , i.e., in the smallest time scale, which will be denoted by $D^{(L)}(k)$, $k = 1, 2, \dots, N$. Then, it calculates the data sequences of the remaining time scales, $D^{(l)}(k)$, $l = L-1, \dots, 1$, corresponding to the average number of arrivals in intervals of length $\Delta t a^{(L-l)}$. This is given by

$$D^{(l)}(k) = \begin{cases} \Psi \left(\frac{1}{a} \sum_{i=0}^{a-1} D^{(l+1)}(k + ia^{L-l-1}) \right), & \frac{k-1}{a^{L-l}} \in \mathcal{I}_0 \\ D^{(l)}(k-1), & \frac{k-1}{a^{L-l}} \notin \mathcal{I}_0 \end{cases} \quad (8)$$

where $\Psi(x)$ represents round toward the integer nearest x . Note that the block length of equation (1) is related with a and l by $m = a^{L-l}$. Note also that all data sequences have the same length N and that $D^{(l)}(k)$ is formed by sub-sequences of a^{L-l} successive equal values; these sub-sequences will be called l -sequences.

5.2 Inference of the tree dMMPPs

All tree dMMPPs are inferred in order to fit an empirical PMF. For the largest time scale, it is the PMF of the most aggregated data sequence, $D^{(1)}(k)$. For all other time scales, one dMMPP is inferred for each state of the immediately higher time scale. For each dMMPP and time scale, the matched PMF represents the contribution of the time scale to the PMF of its parent state. The parameter fitting procedure of each tree dMMPP comprises several steps, highlighted in the flowchart of Figure 3 and explained in more detail in the next sub-sections.

An important step of the fitting procedure is the identification of the time intervals assigned to each dMMPP state. Let $E^{\vec{s}}$ denote the set of time intervals associated with state \vec{s} , i.e., with $\text{dMMPP}^{\vec{s}}$. Using this notation, the set associated with dMMPP^{\emptyset} will be $E^{\emptyset} = \{1, 2, \dots, N\}$, where N is the number of time intervals, i.e.,

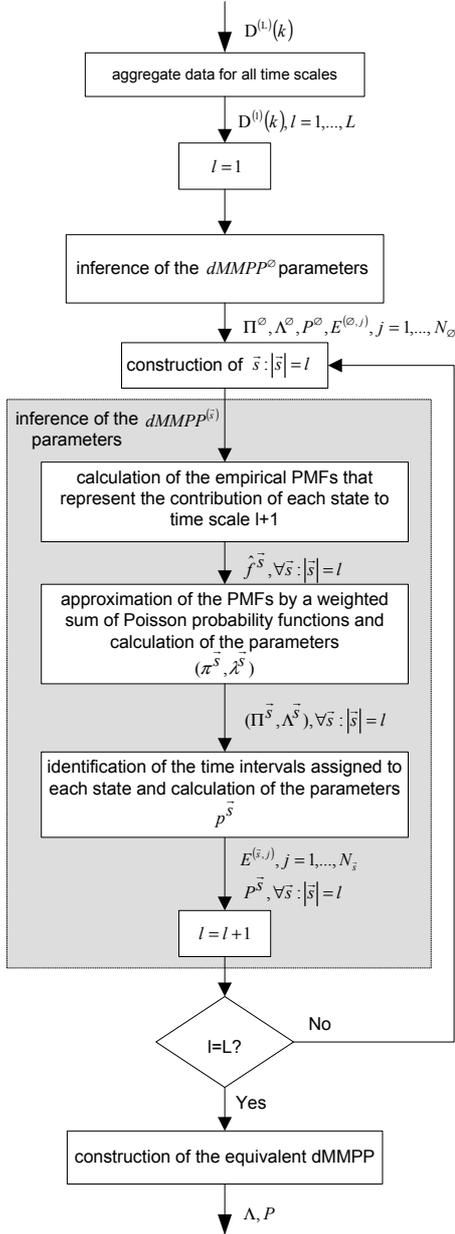


Figure 3: Flow diagram of the inference procedure.

E^0 contains all time intervals. Starting from E^0 , the sets $E^{\vec{s}}$ are successively partitioned at each time scale in a hierarchical fashion. Thus, if states \vec{s} and \vec{t} are such that $|\vec{s}| = |\vec{t}| = l$ and $\vec{s} \neq \vec{t}$, then $E^{\vec{s}} \cap E^{\vec{t}} = \emptyset$ and $\bigcup_{\vec{s}: |\vec{s}|=l} E^{\vec{s}} = E^0$. Moreover, if state \vec{s} is a parent of state \vec{t} , that is $\vec{t} = (\vec{s}, j)$, then $E^{\vec{t}} \subseteq E^{\vec{s}}$ and $\bigcup_{j=1, \dots, N_{\vec{s}}} E^{(\vec{s}, j)} = E^{\vec{s}}$.

5.2.1 Calculation of the PMFs

Each dMMPP will be inferred from a PMF that represents a contribution to a particular time scale. Except for the root dMMPP, the contribution of a dMMPP at time scale l generated from state \vec{s} corresponds to the deconvolution of the empirical PMFs, calculated over the set of time intervals $E^{\vec{s}}$, at this time scale $l = |\vec{s}| + 1$ and previous time scale $l - 1 = |\vec{s}|$, i.e., $\hat{f}_p^{\vec{s}}(x) = [\hat{p}^{\vec{s}, |\vec{s}|+1} \otimes^{-1} \hat{p}^{\vec{s}, |\vec{s}|}](x)$, where $\hat{p}^{\vec{s}, l}$ represents the PMF obtained from the data sequence $D^l(k)$, $k \in E^{\vec{s}}$. Note that the two empirical PMFs are obtained from the same set of time intervals but aggregated at different levels.

However, this may result in probability mass at negative arrival rates for the dMMPP $^{\vec{s}}$, which will occur whenever $\min \{x : \hat{p}^{\vec{s}, |\vec{s}|}(x) > 0\} < \min \{x : \hat{p}^{\vec{s}, |\vec{s}|+1}(x) > 0\}$. To correct this, the dMMPP $^{\vec{s}}$ will be fitted to

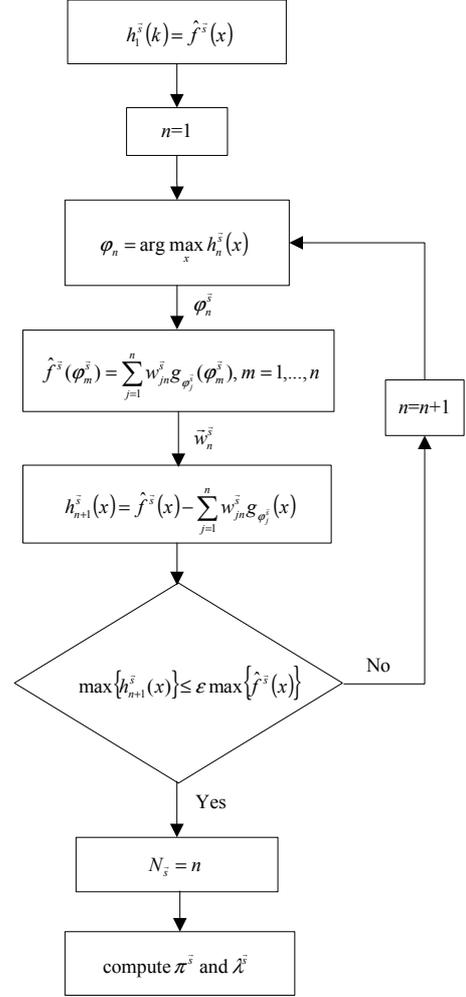


Figure 4: Algorithm for calculating the number of states and the Poisson arrival rates of dMMPP $^{\vec{s}}$.

$$\hat{f}^{\bar{s}}(x) = \hat{f}_p^{\bar{s}}(x + e^{\bar{s}}) \quad (9)$$

where $e^{\bar{s}} = \min\left(0, \min\left\{x : \hat{f}_p^{\bar{s}}(x) > 0\right\}\right)$, which assures $\hat{f}^{\bar{s}}(x) = 0, x < 0$. These additional factors are removed in the next step of the inference procedure.

5.2.2 Inference of the parameters

The first step in the inference of the dMMPP $^{\bar{s}}$ parameters, is the approximation of $\hat{f}^{\bar{s}}$, by a weighted sum of Poisson probability functions. This is based on an algorithm that progressively subtracts a Poisson probability function from $\hat{f}^{\bar{s}}$. The most important steps of this algorithm are depicted in the flowchart of Figure 4 and will be explained in the next paragraphs.

Let the n^{th} Poisson probability function, with mean $\varphi_n^{\bar{s}}$, be represented by $g_{\varphi_n^{\bar{s}}}(x)$ and define $h_n^{\bar{s}}(x)$ as the difference between $\hat{f}^{\bar{s}}(x)$ and the weighted sum of Poisson probability functions at the n^{th} iteration. Initially, we set $h_1^{\bar{s}}(x) = \hat{f}^{\bar{s}}(x)$ and, in each step, we first detect the maximum of $h_n^{\bar{s}}(x)$. The corresponding x -value, $\varphi_n = \arg \max_x h_n^{\bar{s}}(x)$, will be considered the n^{th} Poisson rate of the dMMPP $^{\bar{s}}$. We then calculate the weights of each Poisson probability function, $\vec{w}_n^{\bar{s}} = [w_{1n}^{\bar{s}}, w_{2n}^{\bar{s}}, \dots, w_{nn}^{\bar{s}}]$, through the following set of linear equations:

$$\hat{f}^{\bar{s}}(\varphi_m^{\bar{s}}) = \sum_{j=1}^n w_{jn}^{\bar{s}} g_{\varphi_j^{\bar{s}}}(\varphi_m^{\bar{s}}) \quad (10)$$

for $m = 1, \dots, n$. This assures that the fitting between $\hat{f}^{\bar{s}}(x)$ and the weighted sum of Poisson probability functions is exact at $\varphi_m^{\bar{s}}$ points, for $m = 1, 2, \dots, n$. The final step in each iteration is the calculation of the new difference function

$$h_{n+1}^{\bar{s}}(x) = \hat{f}^{\bar{s}}(x) - \sum_{j=1}^n w_{jn}^{\bar{s}} g_{\varphi_j^{\bar{s}}}(x). \quad (11)$$

The algorithm stops when the maximum of $h_n^{\bar{s}}(x)$ is lower than a pre-defined percentage of the maximum of $\hat{f}^{\bar{s}}(x)$. At this point, the number of states of the dMMPP $^{\bar{s}}$, $N_{\bar{s}}$, is made equal to n .

After $N_{\bar{s}}$ has been determined, the parameters of the dMMPP $^{\bar{s}}$ $\{(\pi_j^{\bar{s}}, \lambda_j^{\bar{s}}), j = 1, 2, \dots, N_{\bar{s}}\}$ are set equal to

$$\pi_j^{\bar{s}} = w_{jN_{\bar{s}}}^{\bar{s}} \quad \text{and} \quad \lambda_j^{\bar{s}} = \varphi_j^{\bar{s}}. \quad (12)$$

Note that the number of states of each dMMPP depends on the level of accuracy employed in the approximation of $f^{\bar{s}}$ by the weighted sum of Poisson probability functions.

The next step of the parameter inference procedure is to associate, at each time scale, one of the dMMPP $^{\bar{s}}$ states with each time interval. Recall that the set of time intervals associated with dMMPP $^{\bar{s}}$ is $E^{\bar{s}}$ and that the data sequences aggregated at time scale l have a^{L-l} successive equal values called l-sequences. The goal here is to partition $E^{\bar{s}}$ into subsets $E^{(\bar{s},j)}$, $j = 1, \dots, N_{\bar{s}}$. The state assignment process considers only the first time interval of each l-sequence, defined by $i = a^{L-(|\bar{s}|+1)}(k-1) + 1, k \in \mathbb{N}, i \in E^{\bar{s}}$. The state that is assigned to l-sequence i is calculated randomly according to the probability vector $\vec{\theta}^{\bar{s}}(i) = \left\{\theta_1^{\bar{s}}(i), \dots, \theta_{N_{\bar{s}}}^{\bar{s}}(i)\right\}$, with

$$\theta_n^{\bar{s}}(i) = \frac{g_{\lambda_n^{\bar{s}}}(D^{(|\bar{s}|+1)}(i))}{\sum_{j=1}^{N_{\bar{s}}} g_{\lambda_j^{\bar{s}}}(D^{(|\bar{s}|+1)}(i))} \quad (13)$$

where $n = 1, \dots, N_{\bar{s}}$. Recall that $\lambda_j^{\bar{s}}$ represents the Poisson arrival rate of the j^{th} state of dMMPP $^{\bar{s}}$ and $g_{\lambda}(y)$ represents a Poisson probability distribution function with mean λ . The elements of this vector represent the probability that the state j had originated the number of arrivals $D^{(l)}(k)$ at time interval k from time scale l .

After this step, we infer the dMMPP $^{\bar{s}}$ transition probabilities, $p_{od}^{\bar{s}}$, with $o, d = 1, \dots, N_{\bar{s}}$, counting the number of transitions between each pair of states. If $n_{od}^{\bar{s}}$ represents the number of transitions from state o to state d of the dMMPP $^{\bar{s}}$, then we let

$$p_{od}^{\bar{s}} = \frac{n_{od}^{\bar{s}}}{\sum_{m=1}^{N_{\bar{s}}} n_{om}^{\bar{s}}}, o, d = 1, \dots, N_{\bar{s}} \quad (14)$$

The transition probability and the Poisson arrival rate matrices of the dMMPP \vec{s} are then given by

$$\mathbf{P}^{\vec{s}} = \begin{bmatrix} p_{11}^{\vec{s}} & p_{12}^{\vec{s}} & \cdots & p_{1N_{\vec{s}}}^{\vec{s}} \\ p_{21}^{\vec{s}} & p_{22}^{\vec{s}} & \cdots & p_{2N_{\vec{s}}}^{\vec{s}} \\ \cdots & \cdots & \cdots & \cdots \\ p_{N_{\vec{s}}1}^{\vec{s}} & p_{N_{\vec{s}}2}^{\vec{s}} & \cdots & p_{N_{\vec{s}}N_{\vec{s}}}^{\vec{s}} \end{bmatrix} \quad \text{and} \quad \mathbf{\Lambda}^{\vec{s}} = \begin{bmatrix} \lambda_1^{\vec{s}} & 0 & \cdots & 0 \\ 0 & \lambda_2^{\vec{s}} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \lambda_{N_{\vec{s}}}^{\vec{s}} \end{bmatrix} + e^{\vec{s}}\mathbf{I} \quad (15)$$

The diagonal matrix of the steady-state probabilities will be designated by $\mathbf{\Pi}^{\vec{s}}$.

5.3 Construction of the equivalent dMMPP model

In this Section we construct a dMMPP equivalent to the tree structure of dMMPPs derived in previous sections. The goal is to incorporate in the model the level of detail given by the finest time scale. Thus, the equivalent dMMPP will have a number of states equal to the number of states in the finest time scale of the tree structure, L . These can be identified by paths on the tree structure of the form $\vec{s} = (s_1, s_2, \dots, s_L)$. Note that each state \vec{s} results from its associated states on the corresponding path, $\vec{s}_{[i+1]} = (s_1, s_2, \dots, s_{i+1})$, $i = 0, 1, \dots, L-1$ of the dMMPP $\vec{s}_{[i]}$. Thus, the states of the equivalent dMMPP will have Poisson rates which are the sum of the Poisson rates of its associated states in the tree structure, i.e.,

$$\lambda_{\vec{s}} = \sum_{j=0}^{L-1} \lambda_{s_{j+1}}^{\vec{s}_{[j]}}. \quad (16)$$

The transition between each pair of states is determined by the shortest path in the tree structure, going through the root dMMPP, that joins the two states. Any pair of states descend from one or more common dMMPPs. The first one, at the time scale with higher l , will be denoted by $\vec{s} \wedge \vec{t} = (s_1, s_2, \dots, s_k)$ where $k = \max \{i : s_j = t_j, j = 1, 2, \dots, i\}$.

We first consider the case of $\vec{s} \neq \vec{t}$. The probability of transition from \vec{s} to \vec{t} , $p_{\vec{s},\vec{t}}$, is given by the product of three factors. The first factor accounts for the time scales where \vec{s} and \vec{t} have the same associated states and is given by

$$\phi_{\vec{s},\vec{t}} = \begin{cases} \prod_{j=0}^{|\vec{s} \wedge \vec{t}|-1} p_{s_{j+1}, t_{j+1}}^{\vec{s}_{[j]}}, & |\vec{s} \wedge \vec{t}| \neq 0 \\ 1, & |\vec{s} \wedge \vec{t}| = 0 \end{cases} \quad (17)$$

The second factor accounts for the transition in the time scale where \vec{s} and \vec{t} are associated to different states of the same dMMPP, which corresponds to $p_{s_{|\vec{s} \wedge \vec{t}+1}, t_{|\vec{s} \wedge \vec{t}+1}}^{\vec{s} \wedge \vec{t}}$. The third factor accounts for the steady-state probabilities of states associated to \vec{t} in the time scales that are not common to \vec{s} and is given by

$$\psi_{\vec{s},\vec{t}} = \prod_{j=|\vec{s} \wedge \vec{t}+1}^{L-1} \pi_{t_{j+1}}^{\vec{t}_{[j]}} \quad (18)$$

where an empty product is equal to one.

Finally, for $\vec{s} \neq \vec{t}$,

$$p_{\vec{s},\vec{t}} = \phi_{\vec{s},\vec{t}} p_{s_{|\vec{s} \wedge \vec{t}+1}, t_{|\vec{s} \wedge \vec{t}+1}}^{\vec{s} \wedge \vec{t}} \psi_{\vec{s},\vec{t}}. \quad (19)$$

In case $\vec{s} = \vec{t}$, the transition probability is simply

$$p_{\vec{s},\vec{t}} = \phi_{\vec{s},\vec{t}}. \quad (20)$$

6 Overview of the traffic traces

Two different traces of aggregated IP traffic were selected to test the accuracy of the proposed fitting procedure: (i) the well known and publicly available Bellcore pOct LAN trace [1] and (ii) a trace measured at the backbone of a Portuguese ISP ADSL network, characterizing the downstream Internet access traffic of approximately 65 simultaneous users. A third trace was also considered, corresponding to the downstream traffic from 10 users of the file sharing application Kazaa, a protocol running over TCP. This trace was measured at the premises of

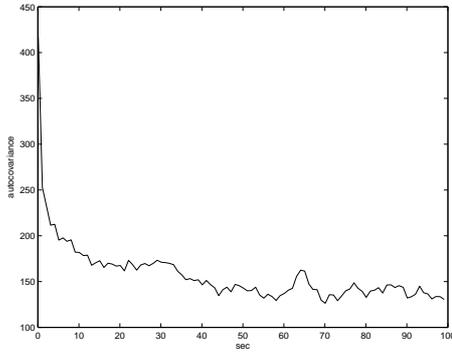


Figure 5: Autocovariance of packet counts, trace pOct.

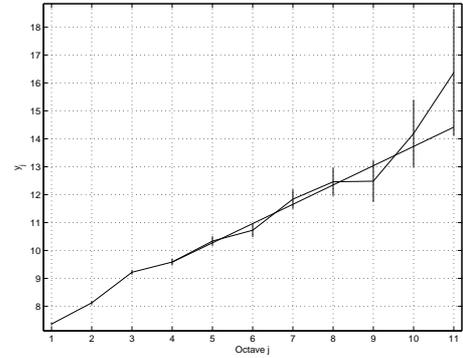


Figure 6: Second order Logscale Diagram, trace pOct.

the same Portuguese ISP and its inclusion is due to the fact that an increasing percentage of the overall Internet traffic belongs to peer-to-peer protocols of the same type as Kazaa. For all our measurements, the traffic analyzer was a 1.2 GHz AMD Athlon PC, with 1.5 Gbytes of RAM and running WinDump, and recorded the arrival instant and the IP header of each packet. The main characteristics of all selected traces are described in Table 1.

All traces exhibit self-similar characteristics: taking trace pOct, for example, the analysis of its autocovariance function (Figure 5) lead us to suspect that it exhibits LRD behavior, due to the slow decay for large time lags. This is confirmed by the scaling analysis, since the y_j values in the logscale diagram are aligned between a medium octave (7) and octave 14, the highest one present in data (Figure 6). A similar analysis was made for the other traces, also revealing the same LRD behavior.

7 Numerical Results

We assess the suitability of the proposed MMPP fitting procedure using several criteria: (i) comparing the Hurst parameters of the original and synthesized (according to the parameters inferred for the resulting dMMPP) data traces; (ii) comparing the PMFs of the packet counts in different time scales, calculated also from the original and synthesized traces and (iii) comparing the queuing behavior, in terms of packet loss ratio (PLR) and average waiting time in queue (AWT), through a trace-driven simulation using those traces as inputs. All simulations were carried out using a fixed packet length equal to the mean packet length of the trace. For all traces, the sampling interval of the counting process was chosen to be 0.1s and three different time scales were considered: 0.1s, 0.2s and 0.4s. Larger aggregation levels were also considered, with good fitting results. For each trace, the estimation procedure took less than 2 minutes, using a MATLAB implementation running in the PC described above, which shows that the procedure is computationally very efficient.

In order to verify that the proposed fitting approach captures the traffic LRD behavior, we compare in Table 1 the Hurst parameters estimated for the original and dMMPP fitted traffic, for each one of the three selected data traces. The LRD estimator that was used is the Logscale Diagram, introduced in Section 2, and Table 1 also includes the range of time scales where the wavelet coefficients follow a straight line, written in parenthesis near to the corresponding Hurst parameter value. As we can see, there is a very good agreement between the Hurst parameter values of the original and fitted traffic, so LRD behavior seems to be well captured by this fitting approach.

The next evaluation criteria is based on the comparison between the PMFs of the original and dMMPP

Trace name	Capture period	Trace size (pkts)	Mean rate (byte/s)	Mean pkt size (bytes)
pOct	Bellcore trace	0.5 million	322790	568
ISP	10.26pm to 10.49pm, October 18 th 2002	0.5 million	583470	797
Kaaza	10.26pm to 11.31pm, October 18 th 2002	0.5 million	131140	1029

Table 1: Main characteristics of measured traces.

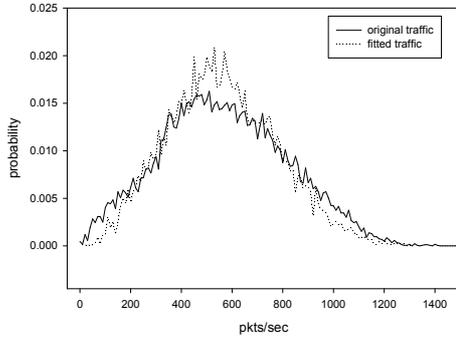


Figure 7: PMF at the smallest time scale, trace pOct

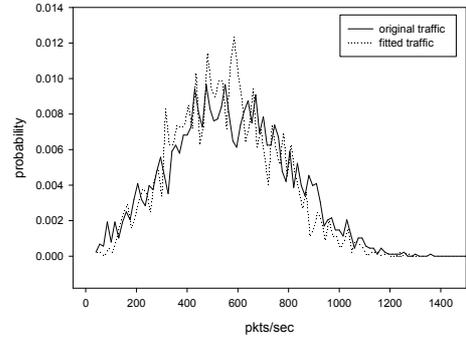


Figure 8: PMF at the intermediate time scale, trace pOct

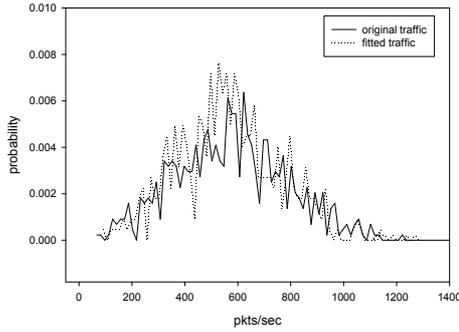


Figure 9: PMF at the largest time scale, trace pOct

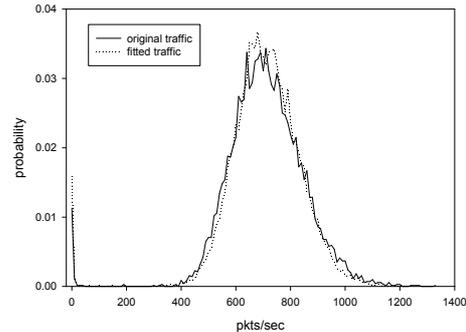


Figure 10: PMF at the smallest time scale, trace ISP

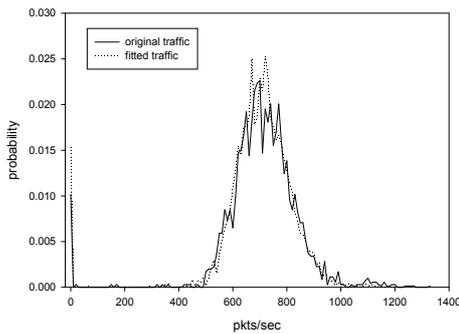


Figure 11: PMF at the intermediate time scale, trace ISP

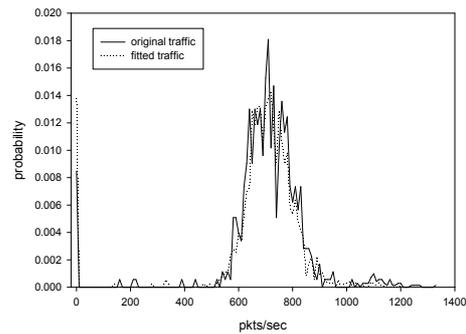


Figure 12: PMF at the largest time scale, trace ISP

fitted traces, for different time scales. Starting with trace pOct, we see from figures 7, 8 and 9 that there is a good agreement between the PMFs of the original and dMMPP fitted traces, for the smallest, intermediate and largest time scales. This is achieved with a dMMPP having 81 states. For trace ISP, the resulting dMMPP has 74 states and the comparison between the PMFs of the original and fitted traces, shown in figures 10, 11 and 12 for the smallest, intermediate and largest time scales, also reveals a good agreement. Finally, for trace Kazaa the resulting dMMPP has 38 states and the PMFs of the original and fitted traces, shown in figures 13, 14 and 15 for the smallest, intermediate and largest time scales, also reveal a good agreement. Note that, as stated before, the number of states is directly related to the level of accuracy used in the fitting task that approximates

Trace	original	fitted
pOct	0.846 (4,11)	0.859 (4,11)
ISP	0.954 (4,10)	0.956 (4,10)
Kazaa	0.917 (8,12)	0.897 (6,12)

Table 2: Comparison between Hurst parameter values

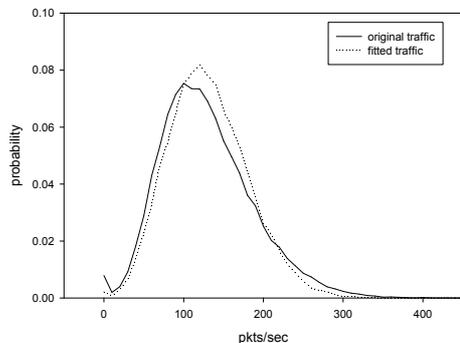


Figure 13: PMF at the smallest time scale, trace

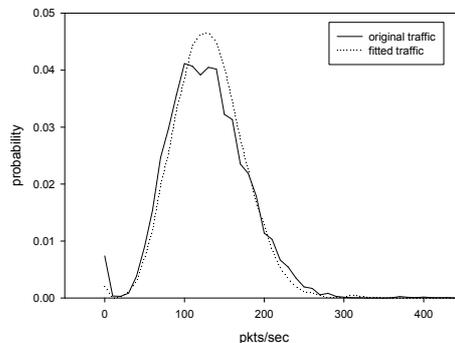


Figure 14: PMF at the intermediate time scale, trace

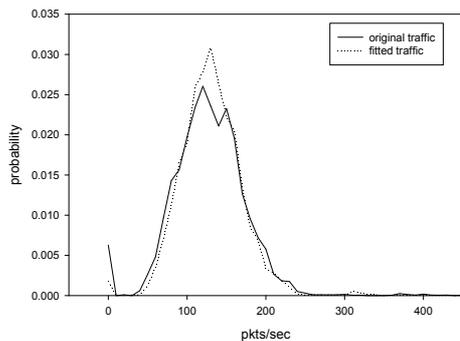


Figure 15: PMF at the largest time scale, trace

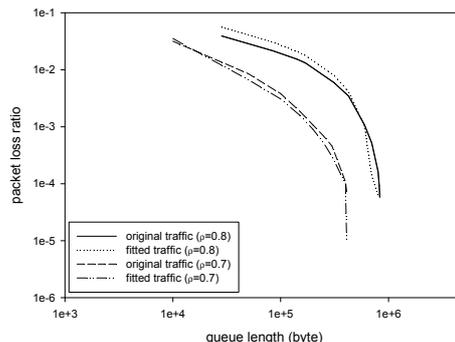


Figure 16: Packet loss ratio, trace pOct.

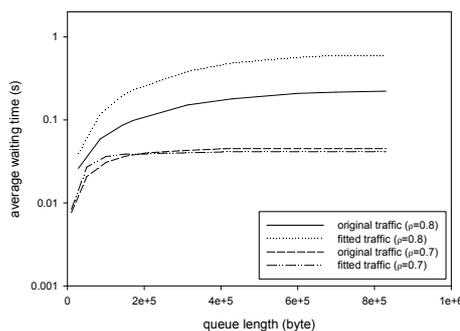


Figure 17: Average waiting time in queue, trace pOct.

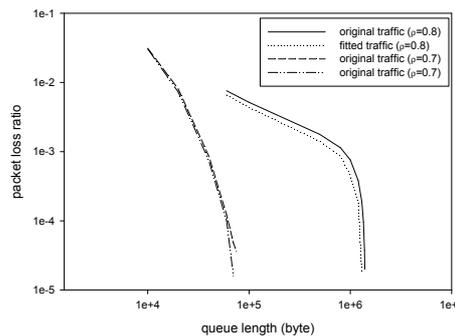


Figure 18: Packet loss ratio, trace Kazaa.

the empirical PMF at each time scale by a weighted sum of Poisson probability functions. So, there is always a tradeoff between the number of states of the resulting dMMPP and the intended level of accuracy.

We now verify if the close match obtained in the Hurst parameter values and in the PMFs at each time scale is enough to guarantee a similar queuing behavior between the original traffic and the fitted models. For each selected trace we compare the PLR and AWT values obtained through trace-driven simulation of the original and dMMPP fitted traces. Two different sets of utilization ratios were used in the simulations: for traces pOct and Kazaa, we used $\rho = 0.7$ and $\rho = 0.8$ and for trace ISP the selected values were $\rho = 0.8$ and $\rho = 0.9$. This is due to the lower burstiness of the ISP traffic, which leads to lower packet losses for the same link utilization. From figures 16 and 17 it is possible to see that, for trace pOct, PLR behavior is very well approximated by the equivalent dMMPP for both utilization ratios, while the agreement of the AWT curves is less accurate specially for higher utilization ratios. For trace Kazaa, the results are depicted in figures 18 and 19 and for trace ISP the results are illustrated in figures 20 and 21. For both traces, the agreement between the PLR curves corresponding to the original and fitted traces is good. However, as the utilization ratio increases the deviation slightly increases, because the sensitivity of the metrics to differences in the traces under comparison is higher. Regarding AWT, the agreement between the curves corresponding to the original and fitted traces is also good, specially for higher utilization ratios.

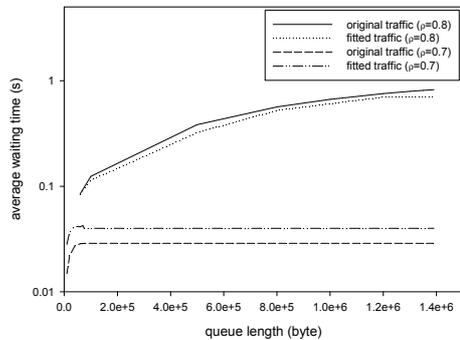


Figure 19: Average waiting time in queue, trace Kazaa.

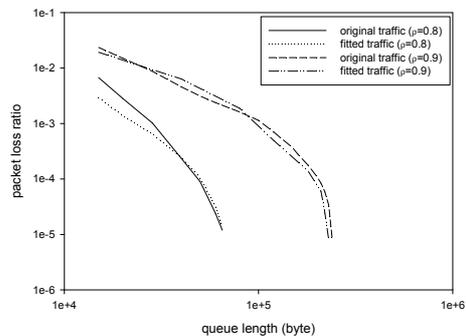


Figure 20: Packet loss ratio, trace ISP.

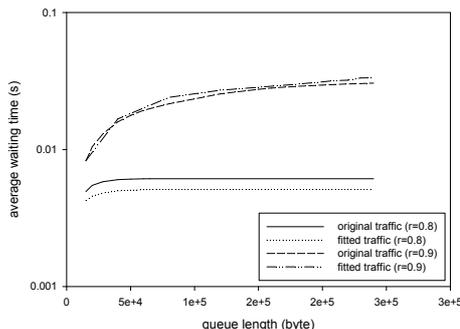


Figure 21: Average waiting time in queue, trace ISP.

As a final remark, we can say that the proposed fitting approach provides a close match of the Hurst parameters and probability mass functions at each time scale, and this agreement reveals itself sufficient to drive a good queuing performance in terms of packet loss ratio and average waiting time in queue. The computational complexity of the fitting method is also very small.

8 Conclusions

We proposed a MMPP, and its associated parameter fitting procedure, which is able to capture self-similarity over a range of time scales. This is accomplished through a hierarchical construction procedure that, starting from a MMPP that matches the distribution of packet counts at the coarsest time scale, successively decomposes each MMPP state into new MMPPs, that incorporate a more detailed description of the distribution at finer time scales. The traffic process is then represented by a MMPP equivalent to the constructed hierarchical structure. The accuracy of the fitting procedure was evaluated by comparing the Hurst parameter, the probability mass function at each time scale and the queuing behavior (as assessed by the loss probability and average waiting time), corresponding to the measured and to synthetic traces generated from the inferred models. Several measured traffic traces exhibiting self-similar behavior were considered: the well-known pOct Bellcore trace, a trace of aggregated IP WAN traffic, and a trace corresponding to the the popular file sharing application Kazaa. Our results show that the proposed model and parameter fitting procedure are very effective in matching the main characteristics of the measured traces over the different time scales present in data.

References

- [1] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the self-similar nature of Ethernet traffic (extended version),” *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, Feb. 1994.
- [2] J. Beran, R. Sherman, M. Taqqu, and W. Willinger, “Long-range dependence in variable-bit rate video traffic,” *IEEE Transactions on Communications*, vol. 43, no. 2/3/4, pp. 1566–1579, 1995.
- [3] V. Paxson and S. Floyd, “Wide-area traffic: The failure of Poisson modeling,” *IEEE/ACM Transactions*

- on *Networking*, vol. 3, no. 3, pp. 226–244, June 1995.
- [4] B. Ryu and A. Elwalid, “The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities,” *ACM Computer Communication Review*, vol. 26, pp. 3–14, Oct. 1996.
- [5] M. Crovella and A. Bestavros, “Self-similarity in World Wide Web traffic: Evidence and possible causes,” *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, Dec. 1997.
- [6] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, “Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level,” *IEEE/ACM Transactions on Networking*, vol. 5, no. 1, pp. 71–86, Feb. 1997.
- [7] W. Willinger, V. Paxson, and M. Taqqu, *Self-similarity and Heavy Tails: Structural Modeling of Network Traffic*, A Practical Guide to Heavy Tails: Statistical Techniques and Applications. Birkhauser, 1998.
- [8] D.P.Heyman and T.V.Lakshman, “What are the implications of long range dependence for VBR video traffic engineering?,” *IEEE/ACM Transactions on Networking*, vol. 4, no. 3, pp. 301–317, June 1996.
- [9] A. Neidhardt and J. Wang, “The concept of relevant time scales and its application to queuing analysis of self-similar traffic,” in *Proceedings of SIGMETRICS’1998/PERFORMANCE’1998*, 1998, pp. 222–232.
- [10] M. Grossglauser and J. C. Bolot, “On the relevance of long-range dependence in network traffic,” *IEEE/ACM Transactions on Networking*, vol. 7, no. 5, pp. 629–640, Oct. 1999.
- [11] A. Nogueira and R. Valadas, “Analyzing the relevant time scales in a network of queues,” in *SPIE’s International Symposium ITCOM 2001*, Aug. 2001.
- [12] K. Meier-Hellstern, “A fitting algorithm for Markov-modulated Poisson process having two arrival rates,” *European Journal of Operational Research*, vol. 29, 1987.
- [13] P. Skelly, M. Schwartz, and S. Dixit, “A histogram-based model for video traffic behaviour in an ATM multiplexer,” *IEEE/ACM Transactions on Networking*, pp. 446–458, Aug. 1993.
- [14] R. Grünenfelder and S. Robert, “Which arrival law parameters are decisive for queueing system performance,” in *ITC 14*, 1994.
- [15] S. Kang and D. Sung, “Two-state MMPP modelling of ATM superposed traffic streams based on the characterisation of correlated interarrival times,” *IEEE GLOBECOM’95*, pp. 1422–1426, Nov. 1995.
- [16] S. Li and C. Hwang, “On the convergence of traffic measurement and queuing analysis: A statistical-match and queuing (SMAQ) tool,” *IEEE/ACM Transactions on Networking*, pp. 95–110, Feb. 1997.
- [17] A. Andersen and B. Nielsen, “A Markovian approach for modeling packet traffic with long-range dependence,” *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 5, pp. 719–732, June 1998.
- [18] C. Nunes and A. Pacheco, “Parametric estimation in MMPP(2) using time discretization,” *Proceedings of the 2nd International Symposium on Semi-Markov Models: Theory and Applications*, Dec. 1998.
- [19] P. Salvador and R. Valadas, “A fitting procedure for Markov modulated Poisson processes with an adaptive number of states,” in *Proceedings of the 9th IFIP Working Conference on Performance Modelling and Evaluation of ATM & IP Networks*, June 2001.
- [20] T. Yoshihara, S. Kasahara, and Y. Takahashi, “Practical time-scale fitting of self-similar traffic with Markov-modulated Poisson process,” *Telecommunication Systems*, vol. 17, no. 1-2, pp. 185–211, 2001.
- [21] P. Salvador, A. Pacheco, and R. Valadas, “Multiscale fitting procedure using Markov modulated Poisson processes,” *Telecommunications Systems*, vol. 23, no. 1-2, pp. 123–148, June 2003.
- [22] D. Veitch and P. Abry, “A wavelet based joint estimator for the parameters of LRD,” *IEEE Transactions on Information Theory*, vol. 45, no. 3, Apr. 1999.