

KPC-Toolbox: Simple Yet Effective Trace Fitting Using Markovian Arrival Processes

Giuliano Casale, Eddy Z. Zhang, Evgenia Smirni
 Department of Computer Science
 College of William and Mary
 23187-8795 Williamsburg, VA, USA
 {casale, eddy, esmirni}@cs.wm.edu

Abstract

We present the KPC-Toolbox, a collection of MATLAB scripts for fitting workload traces into Markovian Arrival Processes (MAPs) in an automatic way. We first present detailed sensitivity analysis that builds intuition on which trace descriptors are most important for queueing. This sensitivity analysis stresses the importance of matching higher-order correlations (i.e., joint moments) of the process inter-arrival times rather than higher order moments of the distribution and provides guidance on the relative importance of different descriptors on queueing. Given that the MAP parameterization space can be very large, we focus on first determining the order of the smallest MAP that can fit the trace well, using the Bayesian Information Criterion (BIC) for determining the best order-accuracy trade-off. Having determined the order of the target MAP, the KPC-Toolbox automatically derives a MAP that captures accurately the most essential features of the trace. Extensive experimentation illustrates the effectiveness of the KPC-Toolbox in fitting traces that are well-documented in the literature as very challenging to fit, showing that the KPC-Toolbox provides a simple and powerful solution to fitting accurately trace data into MAPs.

1 Introduction

Markovian Arrival Processes (MAP) are a class of Markov-modulated processes [15] used for fitting real workload traces with time-varying characteristics, e.g., for approximating workloads with short- or long-range dependent behavior [1, 10]. Traces of this type are commonly found in networks and systems, such as disk drives or e-commerce applications [17, 14]. Accurate models of these traces are indispensable in the capacity planning process, otherwise the robustness of capacity planning models may be undermined [4, 14].

The main advantage of MAPs is that they can be easily integrated within queueing systems or queueing networks, and then used in the computation of performance metrics such as mean response times or server utilizations [4]. However, it is often prohibitive to derive MAPs that can reproduce the characteristics of real workloads with temporal dependence. The main reason for this difficulty is the vast parameterization space of MAPs. Matching accurately traces with time-varying characteristics may require assigning the jumping rates between several tens of states, a task that must be supported by proper software tools which currently do not exist.

In this paper, we introduce the KPC-Toolbox, a set of MATLAB scripts for automatic fitting of real workload traces using MAPs. The KPC-Toolbox takes as input a trace of inter-arrival times, automatically searches for the best order of the MAP that can fit the trace accurately, and then derives a MAP which captures the most essential statistical features of the real workload. The underlying technology is the recently-proposed Kronecker Product Composition¹ (KPC) fitting method for MAPs [6]. KPC reduces fitting problems to determining the characteristics of small MAPs composed by no more than two phases. These MAPs can be easily fitted with closed-form formulas and are later composed into a large MAP by Kronecker products. A similar compositional approach, instead based on MAP superposition, has been widely-used in past MAP fitting literature [1]. The novelty of KPC is that the method is able for the first time to impose moments or correlations of *any order* to the resulting MAP, while superposition methods are mostly limited to first and second-order statistical descriptors (e.g., mean arrival intensity, variance-time curve) that can be largely insufficient for accurate queueing prediction [2]. We further stress the generality of KPC pointing out that, in addition to processes with time-varying characteris-

¹The theory of KPC is described in the technical report [6] that is available online.

tics, it can also fit traces into renewal processes if no autocorrelation exists in the trace. In such cases, the tool can be used for moment fitting and a phase-type (PH-type) renewal process is returned.

An open fundamental problem is determining which trace descriptors to fit in the target MAP. To give intuition on this problem, we study the performance of a $MAP/M/1$ queue in order to determine the best moments and correlations to be matched by KPC. We observe that higher-order statistical descriptors that are changed by a modification of the skewness (e.g., tail of the distribution, higher-order correlations) can result in dramatic performance changes. Instead, we find cases where performance can be quite insensitive to the coefficient-of-variation and to the autocorrelations values. We also give evidence that the higher-order correlations (i.e., joint moments [19]), rather than the tail or the higher-order moments of the distribution, are more likely to be responsible of performance differences of the $MAP/M/1$ system. Guided by the above observation, the KPC algorithm focuses on matching the higher-order correlations rather than the higher-order moments of the inter-arrival times.

A fundamental innovation of the KPC-Toolbox is to determine automatically the order of the MAPs used in fitting (i.e., the number of phases to be used in the underlying CTMC). Order selection in MAPs is an important issue because the MAP order can dramatically affect the running times of fitting. To the best of our knowledge, no criteria have been proposed in previous work for determining the size of a MAP. The KPC-Toolbox tackles this issue with an order-selection technique based on the Bayesian Information Criterion (BIC) [18], that is a widely-accepted method for determining the best order-accuracy tradeoff in fitting models such as ARIMA processes. We use a recursive characterization of the MAP autocorrelations given in [6] as input to BIC to evaluate the best MAP order to use.

The paper is organized as follows. Background on MAPs is given in Section 2. In Section 3 we present a sensitivity analysis on the $MAP/M/1$ queue that guides KPC to decide which moments and autocorrelations are the most important and should thus be matched. Section 4 introduces the KPC-Toolbox and the new BIC -based order selection method. We illustrate the effectiveness of the tool using the case studies in Section 5. Finally, Section 6 concludes the paper and outlines future work. The KPC-Toolbox is available for download at <http://www.cs.wm.edu/MAPQN/kpctoolbox.html>.

2 Markovian Arrival Processes

Here we provide a simple introduction to Markovian Arrival Processes and explain some fundamental difficulties of MAP fitting. We also summarize the Kronecker Product Composition (KPC) fitting method for MAPs [6]. For

a detailed theoretical treatment of properties of MAPs on which KPC is based upon, we direct the interested reader to [6]. Here, we present the tool implementation of the KPC as well as new results that aid the tool user in finding a best MAP fitting in an automatic way.

2.1 MAP Representation

A $MAP(N)$ [15] can be expressed as a continuous-time Markov chain (CTMC) with N phases and an absorbing state. Upon entering the absorbing state from phase s , an interarrival time sample ΔT_k is generated by considering the cumulative time passed from when the CTMC was started. Then, the CTMC is re-started from a phase s' with probability $p_{s,s'}$. Intuitively, the stochastic matrix $P = \{p_{s,s'}\}$, called the *embedded process* of the MAP, allows to create temporal dependence between consecutive values of ΔT_k , as it connects the sampling process of two consecutive inter-arrival times. This allows to create temporal dependent workloads and generalizes PH-type renewal processes [15], for which the only difference is that the probability $p_{s,s'} = \alpha_{s'}$ is independent of s .

The most frequently used MAP representation is the (D_0, D_1) description [13]. If the CTMC of the MAP has an infinitesimal generator Q of order N , the (D_0, D_1) representation is obtained by filtering the transitions of Q according to whether or not they lead to an absorption. That is, D_0 has the same diagonal as Q but its off-diagonal elements are the rates of transitions that do *not* jump to the absorbing state; D_1 includes only transitions to the absorbing state and is immediately computed as the matrix that assures $Q = D_0 + D_1$. In this representation, the embedded process is given by $P = (-D_0)^{-1}D_1$.

For fitting, the fundamental property of the (D_0, D_1) representation is the simplicity of evaluating MAP feasibility. A MAP is feasible if D_0 is invertible, Q is irreducible, and all entries of D_0 and D_1 are non-negative except for the diagonal of D_0 . These conditions can be easily checked. Other representations are possible, e.g. $((-D_0)^{-1}, P)$, but it is hard to discriminate without computing D_0 and D_1 if they produce a valid MAP. We do not consider these representations in the rest of the paper.

2.2 Moment Matching

Fitting a trace requires to capture the properties of a time series in terms of distribution and correlations between samples, which jointly summarize the observed patterns. Because of the difficulty in obtaining robust estimates of the probability density function, moment matching is largely used in fitting, e.g., [11, 19]. This also allows a compact representation of the most important characteristics of the trace.

The inter-arrival time distribution of a MAP is PH-type distributed, thus moments are obtained from standard theory of PH-type distributions as

$$E[X^k] = k! \bar{\pi}_e (-D_0)^{-k} \bar{e}, \quad k = 1, 2, \dots, \quad (1)$$

where $\vec{\pi}_e P = \vec{\pi}_e$. The vector $\vec{\pi}_e$ is independent of the last visited phase before absorption, which means that temporal dependence does not affect the inter-arrival time distribution. For this reason, fitting algorithms for PH-type distributions can be also applied to MAP distribution fitting [16].

A popular approach to matching the temporal dependence structure is to use second-order properties of the trace [1], such as the correlations

$$E[X_j X_{j+k}] = \vec{\pi}_e (-D_0)^{-1} P^k (-D_0)^{-1} \vec{e}, \quad k = 1, 2, \dots, \quad (2)$$

where X_j and X_{j+k} are inter-arrival times spaced by $k - 1$ arrivals. The autocorrelation function is a normalization of these values, i.e.,

$$\rho_k = \frac{\vec{\pi}_e (-D_0)^{-1} P^k (-D_0)^{-1} \vec{e} - E[X]^2}{2E[X^2] - E[X]^2}, \quad k = 0, 1, \dots$$

Equation (2) is more difficult to fit to real data than (1) because of P^k that accounts for the temporal dependence. In addition, imposing a certain distribution by (1) reduces the degrees of freedom for manipulating the matrices in (2), because D_0 and $\vec{\pi}_e$ have been already constrained to set some moments $E[X^k]$. This builds intuition on the main issue of inter-arrival time fitting in MAPs: we need to control the properties of moments and correlations by *jointly* manipulating products of matrices appearing in (1)-(2). Separate fitting of moments and correlations has been recently attempted [11] and can work successfully on small processes, but the underlying optimization suffers from numerical infeasibility on larger MAPs. The KPC method presented in the next section provides instead a simple divide-and-conquer approach of fitting both (1)-(2) in large MAPs.

2.3 Fitting Large MAPs: Kronecker Product Composition (KPC)

Kronecker Product Composition (KPC) is a new technique proposed in [6] for imposing moments and correlations of arbitrary order to a MAP. The idea is to use fundamental properties of the Kronecker (or tensor) product operator [3] to simplify the control of the matrix products and inversions appearing in (1)-(2). The aim of KPC is to provide a simple way to match a set of moments and autocorrelations by composing together small MAPs, typically MAPs with two phases that can be fitted easily with closed-form formulas [8, 9].

Given $MAP^a = \{D_0^a, D_1^a\}$ and $MAP^b = \{D_0^b, D_1^b\}$, we define the KPC of the two processes as the new MAP

$$MAP^a \otimes MAP^b = \{D_0, D_1\} = \{-D_0^a \otimes D_0^b, D_1^a \otimes D_1^b\}, \quad (3)$$

where \otimes denotes the Kronecker product operator. If MAP^a has K_a phases and MAP^b has K_b phases, then the process $MAP^a \otimes MAP^b$ has $K_a K_b$ phases. This suggests that the KPC operator should be used parsimoniously to preserve model compactness.

As an example of Kronecker product, if the original processes have D_0 matrices

$$D_0^a = \begin{bmatrix} -a_{1,1} & 0 \\ 0 & -a_{2,2} \end{bmatrix}, \quad D_0^b = \begin{bmatrix} -b_{1,1} & b_{1,2} \\ b_{2,1} & -b_{2,2} \end{bmatrix},$$

where $a_{i,j} \geq 0$ and $b_{i,j} \geq 0$ are nonnegative real numbers, then the composition yields $D_0 = -D_0^a \otimes D_0^b$ where

$$D_0 = \begin{bmatrix} -a_{1,1}b_{1,1} & a_{1,1}b_{1,2} & 0 & 0 \\ a_{1,1}b_{2,1} & -a_{1,1}b_{2,2} & 0 & 0 \\ 0 & 0 & -a_{2,2}b_{1,1} & a_{2,2}b_{1,2} \\ 0 & 0 & a_{2,2}b_{2,1} & -a_{2,2}b_{2,2} \end{bmatrix}$$

is a feasible D_0 matrix, having negative elements on the main diagonal only. The last feasibility condition is always enforced by KPC if at least one matrix between D_0^a and D_0^b is a diagonal matrix (e.g., D_0^a in the previous example); otherwise the zero entries in $-D_0^a \otimes D_0^b$ are replaced by negative values that make the MAP infeasible. Note that this condition does not place constraints on the generality of the KPC method since one matrix can always be arbitrary.

The logic behind the KPC definition (3) is that the Kronecker product satisfies the algebraic relations

$$(A \otimes B)(C \otimes D) = AC \otimes BD, \quad (A \otimes B)^{-k} = A^{-k} \otimes B^{-k},$$

which allow to decompose matrix products and inversions in terms of similar operations on smaller matrices. For instance, $(-D_0)^{-1} = (-D_0^a)^{-1} \otimes (-D_0^b)^{-1}$ and similarly it can be shown that $P = P^a \otimes P^b$ and $\vec{\pi}_e = \vec{\pi}_e^a \otimes \vec{\pi}_e^b$, where the indexes a and b refer to MAP^a and MAP^b . From these relations, it is easy to show that moments and correlations are decomposed as well. For example, the mean of the KPC $MAP^a \otimes MAP^b$ is decomposable as $E[X] = E_a[X]E_b[X]$, where $E_a[\cdot]$ and $E_b[\cdot]$ denote moments and correlations of the original processes MAP^a and MAP^b . The general decomposition formulas are as follows

$$E[X^k] = E_a[X^k]E_b[X^k]/k!, \quad (4)$$

$$E[X_j X_{j+k}] = E_a[X_j X_{j+k}]E_b[X_j X_{j+k}], \quad (5)$$

and generalize similarly to higher-order moments [6], e.g.,

$$E[X_j^u X_{j+k}^v X_{j+k+h}^z] = E_a[X_j^u X_{j+k}^v X_{j+k+h}^z] \times E_b[X_j^u X_{j+k}^v X_{j+k+h}^z] / (u!v!z!). \quad (6)$$

The KPC technique, shown above for the case of two MAPs, generalizes in a recursive fashion to KPC of several processes. For example, the mean of the composition $MAP^a \otimes MAP^b \otimes MAP^c$ is given by $E[X] = E_a[X]E_b[X]E_c[X]$.

Formulas (4)-(6) state that the MAP fitting of real traces can always be re-formulated in a divide-and-conquer fashion as follows: *assign moments and correlations of smaller*

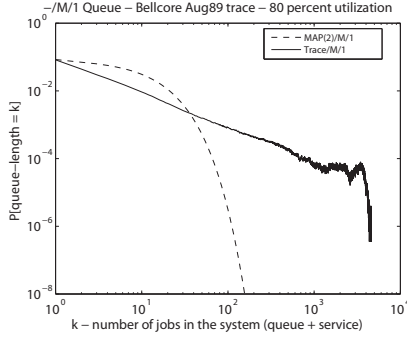


Figure 1. Inaccurate queue-length predictions of a MAP(2) fitted by exact matching of the trace’s most important moments and correlations.

MAPs such that their composition (3) yields a MAP with prescribed moments and correlations by (4)-(6). The KPC-Toolbox presented in Section 4 defines an automatic fitting scheme based on the KPC divide-and-conquer approach.

3 What is Important for Fitting?

Before describing the features of the KPC-Toolbox, we focus our attention on a challenging, but fundamental, question: what statistical descriptors are the most important in fitting by MAPs? A standard approach in current literature is to match the most important moments and correlation coefficients using the simplest available model, e.g., a two-phase MAP(2) process. For example, the first three moments $E[X]$, $E[X^2]$, $E[X^3]$, and the lag-1 correlation $E[X_j X_{j+1}]$, which determines the lag-1 autocorrelation ρ_1 , are sufficient to fully parameterize a MAP(2), see [8, 9] for fitting formulas. Matching these four parameters is often considered a viable approach to fit a trace; however, we argue that this frequently results in models with poor predictive capabilities. For example, Figure 1 shows the simulated queue-length probabilities for a *Trace/M/1* queue driven by the Bellcore Aug89 trace inter-arrival times² [12] and compares them with the probabilities of a *MAP(2)/M/1* queue, where the MAP(2) matches exactly the first three moments and the lag-1 autocorrelation of the Bellcore trace. The results clearly show the poor modeling accuracy of the MAP(2) fitted with this approach.

The experiment in Figure 1 motivates the investigation in this section: we study numerically the sensitivity of the *MAP/M/1* queue-length distribution with respect to the parameters used in MAP fitting. The aim is to derive qualitative recommendations for the best moments and correlations to be matched.

²The Bellcore Aug89 trace is widely used in the literature for accuracy evaluation of fitting techniques.

3.1 Evaluation Methodology

We perform a sensitivity analysis in two phases. We first evaluate the *MAP(2)/M/1* sensitivity and later confirm our observations using a larger MAP(4) process defined by the KPC of two MAP(2)s. Concerning the first analysis on the MAP(2), the process is studied as a function of its first three normalized moments (mean inter-arrival time *MEAN*, squared coefficient-of-variation *SCV*, and skewness *SKEW*) and the lag-1 autocorrelation coefficient ρ_1 of inter-arrival times. A fundamental difficulty in the analysis is that a variation of a single parameter may result in several modifications of the process, e.g., any variation of *SKEW* results in a simultaneous change of the skewness of the distribution, of the tail of the distribution, of higher-order moments, and of higher-order correlations (e.g., the bicorrelations $E[X_j X_{j+k} X_{j+k+z}]$ in (6)). This is a consequence of linear dependencies that relate moments and correlations in a MAP(*N*) [6]. Only first and second-order moments and correlations are unaffected by changes of *SKEW*. Our conjecture is that higher-order correlations, such as the bicorrelations (6), rather than the tail of the distribution or the skewness, are the main determinant of the quality of the fitting. We will provide evidence of this claim in Section 3.3; before, we will generically call “higher-order properties” the moments and correlations changed in a *MAP(2)* by a variation of *SKEW* for fixed *MEAN*, *SCV*, and ρ_1 .

The sensitivity analysis is performed as follows. Recall that for a *MAP/M/1* queue, the queue-length probabilities decay asymptotically as $P(n = k) \sim c_0 \eta^k$, where η is the decay rate (also called caudal characteristic [15]) and c_0 is a positive constant [15]. We investigate the sensitivity of the *MAP/M/1* results by determining the queue-length value x where $\eta^x < 10^{-8}$. Results for different values of the threshold are qualitatively similar. Intuitively, x represents a position of the queue-length distribution after which the probability values are in practice too small to affect performance. If the *MAP/M/1* queue is nearly-insensitive to a *MAP* parameter (e.g., a moment), we expect x to change slightly under variations of that parameter, i.e., the shape of $P(n = x)$ should not be altered significantly. Therefore, we check sensitivity to a fitting parameter by evaluating the sensitivity of “the threshold” x ; the analysis is done for different server utilization levels $\rho \in [0.1, 0.9]$ by varying the mean of the exponential service process.

3.2 MAP(2) Fitting

We first investigate the sensitivity of the *MAP(2)/M/1* results with respect to *SCV*, *SKEW*, and ρ_1 . We focus on the cases $SCV > 1$ and autocorrelations $\rho_k > 0$, which are the most frequently found in real traces.

Sensitivity to *SKEW*. We consider a MAP(2) with unitary mean $MEAN = 1$, $SCV = 10$, $\rho_1 = 0.446$ (which corresponds to an autocorrelation decay rate




Figure 2. Impact of $SKEW$ and SCV on the decay rate of the $MAP(2)/M/1$ queue-length probabilities for different utilization values. Lower skewness values imply heavier tail and lower higher-order correlations; this overall results in dramatic changes of the queue performance.

$\rho_k/\rho_{k-1} = 0.99$), and evaluate the cases $SKEW = 5$ and $SKEW = 100$. The value $SKEW = 100$ represents a distribution with light tail, since the large asymmetry places most of the probability mass around small values; conversely, $SKEW = 5$ has a fat tail. Similar considerations apply to the higher-order correlations: if we normalize the bicomrelations $E[X_j X_{j+1} X_{j+k}]$ it is easy to verify numerically that $SKEW = 5$ has a considerably larger temporal dependence than $SKEW = 100$.

Figure 2(a) shows the impact of $SKEW$ on the threshold x : for all utilization levels, the threshold is ~ 1000 times larger if $SKEW = 5$ instead of $SKEW = 100$, i.e., the performance degrades dramatically for $SKEW = 5$ with the tail of the queue-length probabilities being orders of magnitude longer than for $SKEW = 100$. It is striking to see that the impact of $SKEW$ is considerable also at utilizations as low as $\rho = 0.1$, where the variation of x is even maximal (about three orders of magnitude).

The result indicates that the higher-order properties affected by a change of $SKEW$ can have a remarkable influence on queueing predictions. We remark that if the same experiment is performed after setting all autocorrelations to zero, the variations of x under changes of $SKEW$ are not as dramatic. This is an indication that a long tail in the distribution or higher-order moments may be insufficient to capture alone the $MAP/M/1$ queueing performance, which is consistent with the claim that higher-order correlations are the main determinants of queueing performance.

Sensitivity to SCV . We have performed several experiments for different values of SCV , $SKEW$, and ρ_1 , and found that for large skewness (e.g., $SKEW = 100$), the queue-length probabilities are weakly sensitive to SCV . From now on, we focus on the most difficult case where the tail is not too small and specifically we consider MAPs with $SKEW = 5$ and $SKEW = 15$, where the 99th per-

centile of $SKEW = 5$ is in the range $[11.7, 16.2]$, while for $SKEW = 15$ it is in $[5.6, 9.6]$; in the previous section, for $SKEW = 100$ the 99th percentile is in $[4.4, 4.5]$.

The sensitivity to the SCV is evaluated by setting $SCV = 5$ or $SCV = 10$. Since in general a change of SCV implies also a change in the autocorrelations [6], for $SCV = 10$ we either keep the same ρ_1 of the model for $SCV = 5$ or the same index of dispersion $IDC = SCV(1 + 2\sum_{k=1}^{\infty} \rho_k)$. The results of the experiments are shown in Figure 2(b) and Figure 2(c). The value of $SKEW$ is again fundamental: if $SKEW = 5$, the impact of the SCV is mostly at low utilization, otherwise for $SKEW = 15$ only the high utilization is affected by a change of SCV . The choice of fixing ρ_1 or IDC impacts only at larger loads and suggests that the differences between the two figures cannot be attributed to the autocorrelations only. Explaining the large difference between Figure 2(b) and Figure 2(c) is difficult and the tail of the distribution does not give any clear intuition behind these effects. Instead, a more detailed analysis of temporal dependence reveals that the *distance* between bursty arrival periods is dramatically changed by SCV and $SKEW$, e.g. for $SKEW = 5$ the autocorrelation in counts [7] between consecutive bursty periods of length $T = MEAN$ is $\rho_1^c = 0.91$ for $SCV = 10$ and $\rho_1^c = 0.57$ for $SCV = 5$, while for $SKEW = 15$ the two cases are similar. This suggests that the busy period of the queue may be substantially affected by the SCV and $SKEW$ changes and this could reasonably explain the very different results in Figure 2(b)-2(c). Once again, this stresses the importance of matching correlations, which clearly shape the average distance between bursty periods.

Sensitivity to ρ_1 . Due to limited space, we do not report figures for this case and limit only to summarize our findings. We have evaluated different MAPs for $SCV = 10$,

$SKEW = 5$ or $SKEW = 15$, and $\rho_1 \in [0.0, 0.446]$. The results indicate the following properties: if $SKEW = 15$, then ρ_1 mostly impacts for utilization values larger than 70% with a gap of two orders of magnitude for x between the extreme cases $\rho_1 = 0.0$ ($x \approx 10^3$) and $\rho_1 = 0.446$ ($x \approx 10^5$). For utilization values smaller than 30% and $SKEW = 15$ the queue is almost insensitive to changes in ρ_1 . Conversely, for $SKEW = 5$ the threshold x varies up to two order of magnitude under changes of ρ_1 for all utilization values. This is consistent with our previous findings that the higher-order properties are critical for a good fitting, but also stresses that for medium-high utilizations a good match of the autocorrelations is *always* fundamental.

Summary. The experiments performed in this section indicate that it may be difficult to fit real traces by relying on first and second order properties of the trace only. The higher-order properties controlled by $SKEW$ affect dramatically the performance of a $MAP/M/1$ queue. Also SCV and autocorrelations remain very important, although we have found some insensitivity for certain combinations of utilization and $SKEW$ values; in light of this last observation, it also appears quite difficult to discriminate which is more important between SCV and ρ_1 , as their relative impact changes with the utilization and the $SKEW$ values. In the next section, we complete our analysis on higher-order properties by investigating which is the most important for the performance of the $MAP/M/1$ queue: higher-order moments or higher-order correlations.

3.3 General MAP Fitting

We use KPC to discriminate whether higher-order correlations are more important than higher-order moments in fitting traces with temporal dependence. We proceed as follows. We consider two $MAP(2)$ processes MAP_a and MAP_b such that the resulting KPC $MAP_a \otimes MAP_b$ has fat tail and temporal dependence, and we study the resulting threshold x as a function of the utilization ρ . Then, we repeat the experiment using two different $MAP(2)$ processes MAP_c and MAP_d such that the resulting $MAP_c \otimes MAP_d$ has same $MEAN$, SCV , $SKEW$, and approximately the same tail of the distribution and higher-order moments as $MAP_a \otimes MAP_b$. Table 1 gives the parameters of the MAPs used. The key difference is that $MAP_a \otimes MAP_b$ and $MAP_c \otimes MAP_d$ have considerably different higher-order correlations, i.e., while the inter-arrival time autocorrelations ρ_k are identical, the autocorrelations of counts ρ_k^c reveal that the temporal dependence of $MAP_c \otimes MAP_d$ ($\rho_1^c = 0.82$) is much larger than in $MAP_a \otimes MAP_b$ ($\rho_1^c = 0.32$); similar conclusions are obtained from the normalized bicorrelations.

Since the two distributions are virtually identical in terms of tail decay, we would expect reasonably similar performance if the distribution would be the main determinant of the $MAP/M/1$ queue performance. However, Figure

$MAP(2)$	SCV	$SKEW$	ρ_1
MAP_a	19	19.82	0.468
MAP_b	4	85.88	0.371
MAP_c	19	77.51	0.468
MAP_d	4	210.8	0.371
$MAP_a \otimes MAP_b$	49	57.77	0.482
$MAP_c \otimes MAP_d$	49	57.77	0.482

Table 1. MAP input parameters used in the KPC experiment. All MAPs have $MEAN = 1$.

3 shows that the variation of the threshold x in the two cases is extreme. Except for the high-utilization case where the results must converge because of the identical value of the index of dispersion IDC of the two MAPs, the results of $MAP_c \otimes MAP_d$ are orders of magnitude worse in terms of performance than $MAP_a \otimes MAP_b$. This confirms our intuition that fitting higher-order correlations is fundamental. Further, it appears consistent with the very different counting-process fitting approach in [10], which illustrates the importance of fitting higher-order properties of the counting process.

Summary. Matching higher-order correlations in the KPC-Toolbox has priority over matching higher-order moments or the tail of the distribution. In practice, it is reasonable to assume that matching the first two-three moments of the distribution of temporal dependent traces may be sufficient for capturing the impact of the distribution on queueing performance. All residual degrees of freedom of the model should be spent into matching correlations of different orders. Due to the cost of computing on a real trace a large number of correlations with more than three terms, the correlations $E[X_j X_{j+k}]$ and $E[X_j X_{j+k} X_{j+k+z}]$ are often the most practical descriptors to be matched. Note also that higher-order correlations of the type $E[X_j^2 X_{j+k}]$ are included as special cases of $E[X_j X_{j+k} X_{j+k+z}]$.

4 The KPC Toolbox

The KPC-Toolbox is a set of MATLAB scripts for fitting real traces into MAPs. The toolbox implements an algorithmic solution of the divide-and-conquer fitting problem defined by KPC, see [6]. Several design choices have been made to strike a good balance between simplicity of use and accuracy of fitting; in particular, the underlying fitting paradigm is based on the sensitivity analysis results found in Section 3. In the next subsections we outline the key ideas behind the design of the KPC-Toolbox.

4.1 Order Selection

The KPC-Toolbox approach to MAP order selection is based on a recursive characterization of MAP autocorrelations found in [6]. We consider the trace autocorrela-




Figure 3. Effects of higher-order temporal dependence on the MAP(4)/M/1 queue performance.

tions and, using the characterization in [6], we bound from above the goodness-of-fit of a MAP(N) model, for different choices of the order N . From these values, we select the best order using an information-theoretic approach. For simplicity, we consider autocorrelations, but the approach may be easily extended to evaluate moments or higher-order correlations using the formulas in [6]; this is left as a future extension of the KPC-Toolbox.

It is possible to show that MAP moments and correlations satisfy simple linear recurrence expressions [6]. For example, the autocorrelations ρ_k of a MAP with N phases satisfy

$$\rho_k = a_1 \rho_{k-1} + a_2 \rho_{k-2} + \dots + a_N \rho_{k-N}, \quad k \geq N, \quad (7)$$

where the a_k coefficients are computed from the eigenvalues of $P = (-D_0)^{-1} D_1$ [6]. Therefore, a *necessary* condition for a good matching of the *measured* trace autocorrelations $\hat{\rho}_k$ is that they can be fitted accurately by (7) for some choice of the coefficients a_k . This is not a *sufficient* condition, unless the a_k 's are constrained in sign and magnitude to be generated from a feasible MAP(N); unfortunately, feasibility expressions for the a_k 's are prohibitive to obtain for large MAPs because of the high-order of the nonlinear equations involved. For this reason, our order selection approach relies only on necessary conditions.

We select the best MAP order N^* by (7) as follows. For a trace we consider m autocorrelation coefficients $\hat{\rho}_k$ and give a preliminary evaluation of MAP(N) fitting by defining a linear system of equations (7), i.e.,

$$\begin{cases} \hat{\rho}_{k+1} = a_1 \hat{\rho}_k + a_2 \hat{\rho}_{k-1} + \dots + a_N \hat{\rho}_{k-N+1}, \\ \hat{\rho}_{k+2} = a_1 \hat{\rho}_{k+1} + a_2 \hat{\rho}_k + \dots + a_N \hat{\rho}_{k-N+2}, \\ \vdots \\ \hat{\rho}_{k+m} = a_1 \hat{\rho}_{k+m-1} + a_2 \hat{\rho}_{k+m-2} + \dots + a_N \hat{\rho}_{k-N+m}. \end{cases} \quad (8)$$

The linear system is solved efficiently by linear regression to compute the residual sum of squares $RSS = \sum_k (\hat{\rho}_k - \rho_k)^2$. Linear regression of hundreds of equations (7) can be solved with small time and space requirements, usually in a few seconds. In particular, the KPC-Toolbox uses 500 logarithmically-spaced autocorrelations $\hat{\rho}_k$, $k \in [1, 10^5]$, to parameterize the linear system (8). If the trace has low RSS relatively to the normalizing constant $\sum_k \hat{\rho}_k^2$, then its correlation structure should be relatively easy to fit with a MAP, and the results of the KPC-Toolbox are expected to provide a good fitting model because the temporal dependence structure can be matched well by a MAP. Otherwise, large RSS values indicate that MAPs are unlikely to provide a good fitting of the considered trace and, if no order N has good RSS value, other models with higher-order or different from MAPs should be considered.

After this first evaluation, we use the computed RSS values for different choices of the model size $N \in \{2, 4, 8, 16, 32, 64\}$ to select the best trade-off between accuracy and model size as follows. We adopt the Bayesian Information Criterion (BIC) [18] as quantitative method to estimate the best trade-off. The BIC is defined as [18]

$$BIC(N) = m \log \left(\frac{RSS}{m} \right) + N \log m,$$

where, in our application, m is the number of autocorrelations $\hat{\rho}_k$ used in the regression. According to its definition [18], the BIC should be intended as a cost function, i.e., lower values of BIC denote better trade-offs. Thus, the best order is immediately selected as $N^* = \{N : \min_N BIC(N)\}$, and it is used by the KPC-Toolbox as the best choice for the MAP order. Indeed, similar criteria may be defined by replacing BIC with similar cost-accuracy objective functions, e.g., the Akaike Information Criteria (AIC). However, the BIC is known to be better than AIC as the number of available observations becomes asymptotically large [18] and is therefore used by KPC-Toolbox as the best-available information-criterion method.

We conclude by remarking that, in the special case of traces that are not autocorrelated, the BIC order selection method does not apply, but the KPC toolbox can be used to fit a PH-type renewal process by manually specifying the number of moments desired for the fitting. Matching the first five moments is often sufficient for very accurate distribution fitting, see [6] for case studies.

4.2 KPC-Toolbox Fitting Algorithm

Starting from the observations in the previous section, we have implemented the KPC-Toolbox fitting algorithm as follows. The KPC-Toolbox first performs the BIC order selection and determines the optimal number $J = \log_2 N^*$ of MAP(2)s to be composed by KPC. Then, the tool searches for an actual set of J MAP(2)s that can match accurately the first three moments, the autocorrelations, and

the bicorrelations of the trace. The number of autocorrelations and bicorrelations to be fitted can be specified by the user, otherwise the KPC-Toolbox uses by default 500 logarithmically-spaced autocorrelation coefficients and 100 bicorrelation values obtained from a grid of 10×10 logarithmically-spaced points. The range of sampling is set by default to lags in $[1, 10^5]$ for the autocorrelations and to lags in $[1, 10^5] \times [1, 10^5]$ for the bicorrelations.

The KPC-Toolbox fitting algorithm is organized around three stages: *SCV* and autocorrelation KPC fitting (*stage 1*); mean, skewness and bicorrelation KPC fitting (*stage 2*); generation of the final MAP (*stage 3*). In the first stage, the KPC-Toolbox fits the measured autocorrelations $\hat{\rho}_k$ by searching the values of *SCV* and lag-1 autocorrelation for each MAP(2) that minimize the residual sum-of-squares $RSS = \sum_k (\hat{\rho}_k - \rho_k)^2$, where $\hat{\rho}_k$ is the measured autocorrelation and ρ_k is the autocorrelation of the final KPC process. It can be shown by (4)-(5) that ρ_k can be computed directly from the values of *SCV* and lag-1 autocorrelation of each composing MAP(2) without the need of imposing at this stage mean and skewness for the MAP(2)s, see [6] for a closed-form formula. The optimization program used in stage 1 also includes a constraint imposing that the maximum error between the measured *SCV* and the KPC process *SCV* must be less than 10%.

Mean and skewness of the J MAP(2)s are determined during stage 2 of the fitting algorithm, where we match with a least-squares approach the bicorrelations, mean and skewness of the trace. The only significant difference with respect to the previous stage is that we also impose constraints on the feasible mean and skewness values for the composing MAP(2)s. This is because, upon fixing the *SCV* and lag-1 autocorrelation values found in stage 1, not all combinations of mean and skewness evaluated in stage 2 would result in feasible MAP(2)s; this issue is tackled by adding to the optimization program the feasibility constraints found in [9]. Note also that throughout the entire stage 2 each output couple (*MEAN*, *SKEW*) has a one-to-one mapping with a specific couple (*SCV*, ρ_1) determined by the previous stage. At the end of this stage, we move to determining a (D_0, D_1) representation of the results.

In the final stage 3, we use the values of the first three moments and lag-1 autocorrelation coefficient found in the first two stages to determine the (D_0, D_1) representation of each MAP(2) using the closed-form formulas in [8, 9]. The resulting MAP(2)s are finally composed by KPC into a MAP(N^*) that is returned to the user. More precisely, the three stages described above are run several times to find the best MAP to be returned to the user. The selection algorithm operates as follows. The KPC-Toolbox first runs several times stage 1 while keeping stored the solutions that produced the best ten *RSS* value of the stage 1 optimization. Then, stage 2 is run several times for each of the solu-

tions selected in stage 1. The KPC-Toolbox returns as best MAP the stage 2 solution associated to the overall lowest *RSS* for the bicorrelations. According to this selection approach, the final MAP should have a high-quality fitting of the autocorrelations (stage 1) and a nearly-optimal fitting of the bicorrelations (stage 2) and is thus consistent with the sensitivity analysis conclusions presented in Section 3.

5 Numerical Experiments

We conclude with experiments on the accuracy of the KPC-Toolbox in fitting both real and synthetic traces. Consistently with the sensitivity analysis, we focus our evaluation on the queueing prediction accuracy for a $-/M/1/FCFS$ queue and we fit two traces:

- *Seagate Web trace*: This trace is composed by 3.6×10^6 interarrival times of requests at the disk drive of a Web server, see [17] for a description of this trace and related analyses of its temporal dependence structure.
- *Bellcore Aug89 trace*: this is a benchmark case for evaluating the quality of long-range dependent trace fitting approximations [1, 10, 11]. The traffic consists of 1 million inter-arrival time samples collected in 1989 at the Bellcore Morrison Research and Engineering facility.

We remark that the two traces are used also in the manually-performed fitting experiments in [6]. However, the results presented here are generated *automatically*, and not manually, using the KPC-Toolbox. In this section, we compare the results for the Seagate and the Bellcore traces with the manually-generated results in [6] in order to prove that the automatic fitting approach compares very well with time-consuming manual fitting.

5.1 Order Selection Results

For each of the three traces, the KPC-Toolbox runs the *BIC* order selection for different orders, i.e., $\{4, 8, 16, 32, 64\}$. The results of the *BIC* order selection are shown in Figure 4. For the Bellcore Aug89 trace, the *BIC* selection in Figure 4 indicates that order 16 is the best choice, order 32 is a close candidate, while the other orders are significantly worse in terms of *BIC*. Similarly, for the Seagate Web trace, the *BIC* criterion indicate a best order of 16 phases. These results are consistent with our manual fitting experience, since we were never able to fit MAPs with less than 16 phases that could reproduce effectively the Seagate and Bellcore traces queueing performance. In particular, to the best of our knowledge, the best available fitting for the Bellcore Aug89 trace have either 16 or 32 states [1, 10], which further confirms the reliability of the *BIC* results in Figure 4.

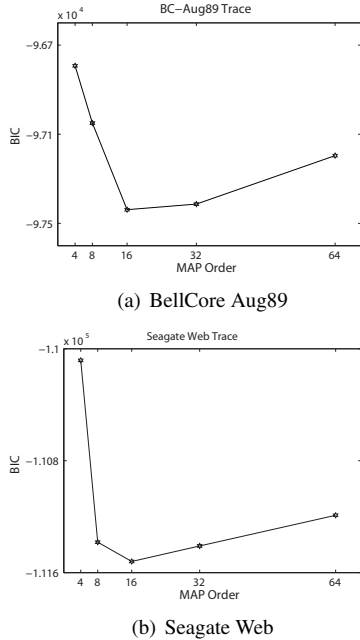


Figure 4. Order selection for the BellCore Aug89 trace and the Seagate Web trace.

5.2 Queuing Results

In figures 5 and 6, we present the complementary cumulative distribution function (CCDF) of the queue length probabilities for a $MAP/M/1$ queue and the empirical CCDF obtained by simulating the $Trace/M/1$ queue. The CCDFs of the $MAP/M/1$ queue are obtained by solving the underlying quasi-birth death process using MAMSolver (available at <http://www.cs.wm.edu/MAMSolver/>). The service rate of the exponential server is adjusted to tune the load of the server at different utilization levels. We plot the queueing distributions at utilizations 20%, 50%, and 80%, standing for the low, medium and high load, respectively.

The Bellcore Aug89 trace has been identified as a difficult trace to fit and thus extensively used in the literature of MAP fitting to evaluate the effectiveness of fitting [5]. At 20% utilization level, the fitted MAP(16) underestimates the queueing probabilities. When the utilization increases, the queueing probability prediction improves. At 50% utilization level, the fitted MAP captures the small and medium queue lengths probabilities better than the 20% utilization level. At 80% utilization, the fitted MAP(16) almost overlaps the trace up to the queue length equal to 8×10^3 . Further, as shown in comparison with the manual fitting, the queueing prediction accuracy for the KPC is comparable with the best in the literature [6] with the additional benefit that KPC automatically generates the MAP while other methods require exhaustive manual tuning to obtain a good

MAP. Similarly with the Seagate Web trace, as utilization increases, the fitted MAP(16) captures more accurately the queueing behavior of the Trace/M/1 queue. For utilization levels 50% and 80%, the fitted MAP captures the beginning probabilities of the queue length very well. Also in this case, the result is comparable with the manually fitted MAP determined in [6] and also plotted in the figure.

6 Conclusion

In this paper, we have presented the KPC-Toolbox, a set of MATLAB scripts for fitting workload traces into MAPs. One of the greatest challenges in MAP fitting is to (a) decide the order of the MAP that is necessary to fit the trace data and (b) determine the relative importance of the various stochastic descriptors to be of the trace that should be matched by the MAP. The KPC-Toolbox meets the above challenges with a novel approach that uses the *BIC* criterion to determine the best order-accuracy trade-off for a MAP and by using optimization to explore a vast parameter space of alternatives such that the most important stochastic properties of the trace are captured by the resulting MAP. Detailed queueing analysis that confirms the importance of matching higher-order correlations (i.e., joint moments) rather than higher order moments is used to guide the optimization.

The tool implements the theoretical results of [6], and guided by *BIC* and the various derived criteria for fitting, it produces good MAP fittings of challenging traces in an *automatic* way. Experimental results on real traces from both the systems and networking domains show the effectiveness of deriving a MAP that captures well the workload characteristics. Although the robustness of the toolbox has been extensively examined and has been shown successful with a large set of traces that are not presented here due to lack of space, we intend to extend the theory and the tool implementation so that it can also apply to traces with arbitrary autocorrelation functions (e.g., periodicity, complex eigenvalues). Ongoing work also includes the trace preprocessing for noise reduction to improve the stability of fitting. The KPC-Toolbox is available for download at <http://www.cs.wm.edu/MAPQN/kpctoolbox.html>.

Acknowledgment

This work was supported by the National Science Foundation under grants ITR-0428330 and CNS-0720699, and by Seagate Research.

References

- [1] A. T. Andersen and B. F. Nielsen. A Markovian approach for modeling packet traffic with long-range dependence. *IEEE JSAC*, 16(5):719–732, 1998.
- [2] A. T. Andersen and B. F. Nielsen. On the use of second-order descriptors to predict queueing behavior of MAPs. *Naval Res. Logistics*, 49(4):391–409, 2002.

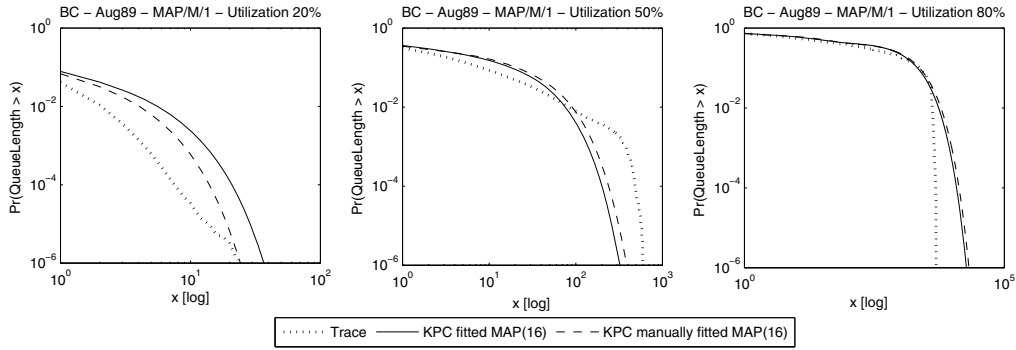


Figure 5. Queuing results for comparison between the Bellcore Aug89 trace and the fitted MAP

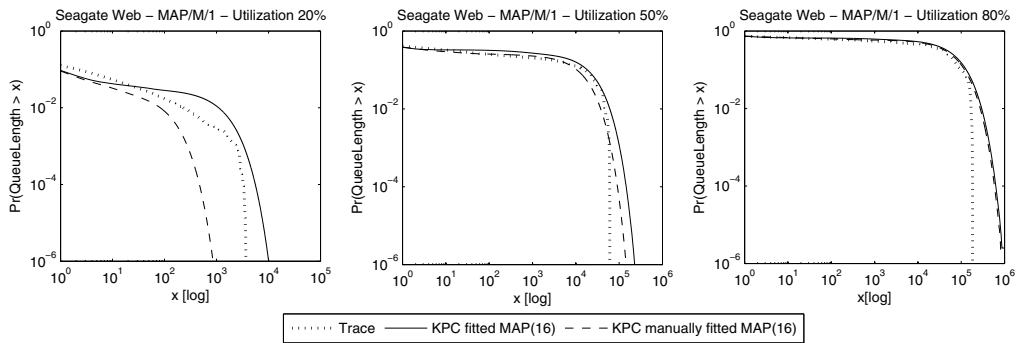


Figure 6. Queuing results for comparison between the Seagate Web trace and the fitted MAP

- [3] J. Brewer. Kronecker products and matrix calculus in system theory. *IEEE T. Circuits and Sys.*, 25(9), 1978.
- [4] G. Casale, N. Mi, and E. Smirni. Bound analysis of closed queueing networks with workload burstiness. In *Proc. of ACM SIGMETRICS 2008*, pages 13–24, ACM Press, 2008.
- [5] G. Casale, E. Zhang, and E. Smirni. Characterization and synthesis of Markovian workload models. In *Proc. of the Globecom 2007 Workshop on Enabling the Future Service-Oriented Internet*, pages 1–5, IEEE Press, Nov 2007.
- [6] G. Casale, E. Zhang, E. Smirni. Interarrival times characterization and fitting for Markovian traffic analysis. College of William and Mary, TR WM-CS-2008-02. Available at <http://www.wm.edu/computerscience/techreport/2008/WM-CS-2008-02.pdf>.
- [7] D. Cox and P. Lewis. *The Statistical Analysis of Series of Events*. John Wiley and Sons, New York, 1966.
- [8] J. Diamond and A. Alfa. On approximating higher-order MAPs with MAPs of order two. *Queueing Systems*, 34:269–288, 2000.
- [9] A. Heindl, K. Mitchell, and A. van de Liefvoort. Correlation bounds for second-order MAPs with application to queueing network decomposition. *Perf. Eval.*, 63(6):553–577, 2006.
- [10] A. Horváth and M. Telek. Markovian modeling of real data traffic: Heuristic phase type and MAP fitting of heavy tailed and fractal like samples. In *Performance Evaluation of Complex Systems: Techniques and Tools, IFIP Performance 2002, LNCS Tutorial Series Vol 2459*, pages 405–434, 2002.
- [11] G. Horváth, P. Buchholz, and M. Telek. A MAP fitting approach with independent approximation of the inter-arrival time distribution and the lag correlation. In *Proc. of 2nd Conf. on Quantitative Evaluation of Systems (QEST)*, pages 124–133, 2005.
- [12] W. E. Leland, M. S. Taqqu, W. Willinger, D. V. Wilson. On the self-similar nature of ethernet traffic. *IEEE/ACM TON*, 2(1):1–15, 1994.
- [13] D. M. Lucantoni. New results on the single server queue with a batch Markovian arrival process. *Stochastic Models*, 7:1–46, 1991.
- [14] N. Mi, Q. Zhang, A. Riska, E. Smirni, and E. Riedel. Performance impacts of autocorrelated flows in multi-tiered systems. *Perf. Eval.*, 64(9-12):1082–1101, 2007.
- [15] M. F. Neuts. *Structured Stochastic Matrices of M/G/1 Type and Their Applications*. Marcel Dekker, New York, 1989.
- [16] A. Riska, V. Diev, and E. Smirni. An EM-based technique for approximating long-tailed data sets with PH distributions. *Perf. Eval.*, 55(1-2):147–164, Jan. 2004.
- [17] A. Riska and E. Riedel. Long-range dependence at the disk drive level. In *Proc. of 3rd Conf. on Quantitative Evaluation of Systems (QEST)*, pages 41–50, IEEE Press, 2006.
- [18] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.
- [19] M. Telek and G. Horváth. A minimal representation of Markov arrival processes and a moments matching method. *Perf. Eval.*, 64(9–12):1153–1168, 2007.