KPC-Toolbox: Best Recipes Toward Automatization of Workload Fitting *

Eddy Zheng Zhang
College of William & Mary
Computer Science Dept.
Williamsburg, 23187, VA, US
eddy@cs.wm.edu

Giuliano Casale
College of William & Mary
Computer Science Dept.
Williamsburg, 23187, VA, US
casale@cs.wm.edu

Evgenia Smirni
College of William & Mary
Computer Science Dept.
Williamsburg, 23187, VA, US
esmirni@cs.wm.edu

ABSTRACT

We present the KPC-Toolbox, a set of MATLAB scripts for fitting workload traces into Markovian Arrival Processes (MAPs) in an automatic way. 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). Having determined the order of the target MAP, the KPC-Toolbox automatically derives a MAP that captures accurately the moments and temporal dependence of the trace. We present experiments showing the effectiveness of the KPC-Toolbox in fitting traces that are well-documented in the literature as very challenging ones to fit.

1. INTRODUCTION

Markovian Arrival Processes (MAPs) are a class of Markov-modulated processes used for fitting real workload traces with time-varying characteristics, e.g., traces with short- or long-range dependent behavior. Traces of this type are commonly found in networks and systems, such as disk drives or e-commerce applications, and need to be modeled accurately to support capacity planning decisions with reliable quantitative predictions [10, 11]. The advantage of MAPs with respect to other fitting models 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 by closed-form formulas MAPs that can reproduce the characteristics of real workloads with temporal dependence. The main 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 extended abstract, 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, e.g., of interarrival times, automatically searches for the smallest order of the MAP that can fit the trace accurately, and then derives the MAP which captures the most essential statistical features of the real workload. The underlying technology is the recently-proposed Kronecker Product Composition (KPC) fitting method for MAPs [5]. The basic idea of KPC is to reduce the moment and temporal dependence fitting problem to assigning the characteristics of smaller MAPs composed by no more than two states. These MAPs can be easily fitted with closed-form formulas and are later composed by Kronecker products into a larger MAP that fits accurately the workload trace. Superposition of small MAPs has been widely-used in past MAP fitting literature, but the novelty of KPC is that the method is able for the first time to impose moments or autocorrelations *of any order* to the fitted MAP, while existing 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 characteristics, it can also fit traces into renewal processes if no autocorrelation exists in the trace. In such cases, the tool emphasis is on moment fitting. We summarize in Section 2.2 the main ideas of the KPC technique and point to [5, 6] for additional details.

Another fundamental innovation of the KPC-Toolbox is the automatic selection of the order of the MAPs used in fitting, i.e., the number of states of the underlying Markov process. Order selection is a challenging issue, because the selected MAP order can dramatically affect the running times of fitting but also the computational costs of solving MAP-based queueing models. The KPC-Toolbox tackles this problem by an order-selection technique based on the Bayesian Information Criterion (BIC) [12] in combination with linear regression models. To our best knowledge, this is the first time that a formal order selection algorithm for MAP fitting is proposed. We describe in Section 3.1 the effectiveness of BIC in MAP fitting using an example of a real trace.

This paper is organized as follows. In Section 2, we give background on MAPs and describe the BIC selection algorithm and the Kronecker Product Composition technique. We provide two case studies in Section 3 that show the effectiveness of the KPC-Toolbox. In Section 4, we conclude the paper. The KPC-Toolbox can be downloaded from *http://www.cs.wm.edu/MAPQN/kpctoolbox.html.*

2. MAP FITTING

A MAP(n) can be expressed as a Markov process that jumps between n states, spending an exponentially-distributed time in each state visit. Some jumps trigger arrival events (*completion transitions*), while others do not produce any effect except for the state change (*background transitions*). According to this definition, a MAP(n) can be represented

*This work was partially supported by the National Science Foundation under grants ITR-0428330 and CNS-0720699.

by a pair of square matrices D_0 and D_1 of order n, where the off-diagonal entries in D_0 are the jumping rates of background transitions, the entries of D_1 are the jumping rates of completion transitions with arrival events, and the diagonal entries of D_0 are imposed so that $Q = D_0 + D_1$ is the infinitesimal generator of the underlying Markov process. We focus on the interval stationary process that describes the interarrival time between two adjacent arrival events. The interarrival times of a MAP are phase-type distributed with k-th moment

$$E[X^k] = k!\vec{\pi}_e(-D_0)^{-k}\vec{e}, \quad k \geq 0, \tag{1}$$

where $\vec{\pi}_e = \vec{\pi}_e P$, $P = (-D_0)^{-1}D_1$, and $\vec{e} = (1, 1, \ldots, 1)^T$. This formula implies that the squared coefficient of variation of the interarrival times is $CV^2 = 2E[X]^{-2}\vec{\pi}_e(-D_0)^{-2}\vec{e}-1$.

Temporal relations between successive interarrival times may be described by autocorrelations or more general joint moments. The lag-k autocorrelation coefficient of interarrival times is computed in MAPs as

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

Let X_i be the i-th interarrival time from an arbitrary starting epoch $i_0 = 0$, and consider a sequence $X_{i_1}, X_{i_2}, ..., X_{i_L}$, where $0 \leq i_1 < i_2 < ... < i_L$. The joint moment of these L ordered interarrivals is given by

$$H(\vec{i}, \vec{k}) = E[X_{i_1}^{k_1} X_{i_2}^{k_2} \cdots X_{i_L}^{k_L}], \tag{3}$$

where $\vec{i} = (i_1, i_2, \ldots, i_L)$ and $\vec{k} = (k_1, k_2, \ldots, k_L)$. The joint moments $H(\vec{i}, \vec{k})$ are known to completely characterize a MAP [2, 13]. They are computed analytically as [13]

$$H(\vec{i}, \vec{k}) = \vec{\pi}_e \left(\prod_{l=1}^{L} k_l!(-D_0)^{-k_l} P^{i_l - i_{l-1}} \right) \vec{e}, \tag{4}$$

where $i_0 = 0$. Note that (4) reduces to (1) in the case $L = 1$.

The problem of MAP fitting is to determine the size of the MAP, that is, the dimension n of D_0 and D_1, and to assign the jumping rates in D_0 and D_1 in order to obtain the desired moments and correlations of the process that best describes the trace. According to these requirements, the KPC Toolbox fitting algorithm consists of two parts: the selection of the optimal order n and the fitting of the resulting n-state MAP based on the KPC technique. These steps are outlined in the following subsections.

2.1 BIC Order Selection

For a given trace, several MAP models with different number of states can fit observations accurately, but in general one wants to minimize the cardinality of these states. This results in simplification of fitting and also reduces the computational costs of evaluating MAP-based queueing models. In order to determine the optimal number of states n that may fit a given trace, i.e., the order n of the MAP, we use an information-theoretic criterion known as the Bayesian Information Criterion (BIC) [12] which we apply starting from the following characterization of MAP autocorrelations.

LEMMA 1. *In a MAP(n), any $n+1$ consecutive autocorrelations are linearly dependent according to the relation*

$$\rho_k = -\sum_{j=1\ldots n} a_j \rho_{k-j}, \; \rho_0 = \left(1 - 1/\mathrm{CV}^2\right)/2, \; k \geq n, \tag{5}$$

where a_j, $j = 1, \ldots, n$, are the coefficients of the characteristic polynomial of P and $\sum_{j=1}^{n} a_j = 0$.

A proof of Lemma 1 and additional details can be found in [5]. Our fundamental observation for order selection is that (5) depends on the order n of the MAP and constraints the class of feasible autocorrelation coefficients for a MAP process. Since the autocorrelations between interarrival times are critical drivers of the performance of queueing models, one usually wants to fit measured autocorrelations accurately and thus we can use (5) to determine for which order n the measured autocorrelation is the most likely to be fitted accurately. This likelihood is estimated using the BIC applied to the residuals of the following linear regression model. For a trace with $n + m$ measured autocorrelations $\hat{\rho}_i$, we build a linear regression model

$$\begin{cases} \hat{\rho}_{k+1} = c_1\hat{\rho}_k + c_2\hat{\rho}_{k-1} + \ldots + c_n\hat{\rho}_{k+1-n} + \epsilon_1, \\ \hat{\rho}_{k+2} = c_1\hat{\rho}_{k+1} + c_2\hat{\rho}_k + \ldots + c_n\hat{\rho}_{k+2-n} + \epsilon_2, \\ \vdots \qquad \vdots \qquad \qquad \vdots \\ \hat{\rho}_{k+m} = c_1\hat{\rho}_{k+m-1} + c_2\hat{\rho}_{k+m-2} + \ldots + c_n\hat{\rho}_{k+m-n} + \epsilon_m, \end{cases}$$

where $\hat{\rho}_i$ represents the i-th estimated autocorrelation from the trace of interarrival time samples, and ϵ_i is the observation error of $\hat{\rho}_i$. We use the Bayesian Information Criterion (BIC) [12] to determine the best trade-off between order n and autocorrelation fitting accuracy as estimated by the residuals of the linear regression model. The BIC is known to be better than other cost-accuracy objective functions like the Akaike Information Criterion (AIC), as the number of available observations becomes asymptotically large. BIC is therefore used by KPC-Toolbox as the best-available information-criterion method. We here focus on autocorrelations, but the approach may be easily extended to moments or higher-order correlations using recurrence formulas in [5] similar to (5).

We evaluate the residual sum of squares $RSS = \sum_k (\hat{\rho}_k - \rho_k)^2$ for different choices of n in the linear regression model. Then, we assess the goodness of MAP size using the Bayesian Information Criterion [12]

$$BIC \equiv BIC(n) = m \log\left(\frac{RSS}{m}\right) + n \log m.$$

The BIC function is intended as a cost function [12], i.e., lower values of BIC indicate a better size-accuracy trade-off. As the model size increases, the BIC typically decreases since more states improve the accuracy of fitting. However, when the order becomes too large the RSS values do not decrease significantly compared to smaller orders and thus the cost term $n \log m$ makes the BIC value to increase with n. Thus, we select the best MAP order according to the lowest BIC value as $\{n : BIC(n) = \min_N BIC(N)\}$, which corresponds to the best trade-off between accuracy and MAP size with respect to the estimate of the linear regression model. We remark that the computational costs of the order selection are often small due to the efficiency of linear regression and that our RSS-based estimate of the quality of autocorrelation fitting is only a preliminary estimate, since here we do not compute the MAP jumping rates and the actual fitting of autocorrelations is performed later by KPC. We also point out that large RSS values indicate that the respective MAPs are unlikely to provide a good fitting of the trace and, in these cases, the KPC-Toolbox casts a warning suggesting that the other fitting models should be considered.

2.2 Kronecker Product Composition

Kronecker Product Composition (KPC) is a technique for defining large MAPs in terms of two-state models. Given J MAPs $\{D_0^j, D_1^j\}$, we define their KPC as the new MAP

$$\{D_0^{\mathrm{kpc}}, D_1^{\mathrm{kpc}}\} = \{(-1)^{J-1} D_0^1 \otimes \cdots \otimes D_0^J, D_1^1 \otimes \cdots \otimes D_1^J\},$$

where $\otimes$ denotes the Kronecker product operator [3]. It can be shown by the properties of the Kronecker product that $P^{\mathrm{kpc}} = -(D_0^{\mathrm{kpc}})^{-1} D_1^{\mathrm{kpc}} = P^1 \otimes \cdots \otimes P^J$ and $\vec{\pi}_e^{kpc} = \vec{\pi}_e^1 \otimes \cdots \otimes \vec{\pi}_e^J$, thus our composition generates a matrix P^{kpc} with simple structure. When all composing MAPs have two states, the parameter J is related to the KPC model order n by $J = \log_2 n$. In KPC, in order to generate a valid MAP, we require that at least $J-1$ of the composing processes have diagonal D_0. The fundamental property of KPC is that the moments and correlations of the composed MAP $\{D_0^{\mathrm{kpc}}, D_1^{\mathrm{kpc}}\}$ are in simple relation with those of the original MAPs. This allows to fit a large MAP generated by KPC by assigning the properties of the smaller MAPs used in the composition, i.e., KPC provides a divide-and-conquer approach to MAP fitting. In fact, the moments of the KPC process satisfy

$$E[X_{kpc}^k] = (k!)^{1-J} E[X_1^k] E[X_2^k] \cdots E[X_J^k].$$

The joint moments of KPC process satisfy

$$H_{kpc}(\vec{i}, \vec{k}\,) = H_1(\vec{i}, \vec{k}\,) H_2(\vec{i}, \vec{k}\,) \cdots H_J(\vec{i}, \vec{k}\,)/(k_1! k_2! \cdots k_L!),$$

with $\vec{i} = (i_1, i_2, ..., i_L)$ and $\vec{k} = (k_1, k_2, ..., k_L)$. Similar formulas exist for the autocorrelation coefficients which allow to compute them in terms of the autocorrelation of the small MAPs that are composed together by KPC [5].

2.3 MAP fitting by KPC

We perform trace fitting based on the KPC moment and autocorrelation formulas. We first assume that the desired MAP order n has been obtained by *BIC* order selection and we determine the corresponding value $J = \log_2 n$ that is the number of MAP(2)s to be composed by KPC. We restrict our attention to two-phase MAPs for composition, because they can be fitted analytically without difficulty [7]. We then proceed to trace fitting according to the following three steps.

Step 1 - Autocorrelation and CV^2 Fitting. We fit the autocorrelations and the squared coefficient of variation CV^2 of interarrival times by a nonlinear optimization program. This program is essentially a least-square optimization constrained to the properties of KPC. The results are two sets $CV^2(j)$ and $\gamma_2(j)$ (the decay rate of the autocorrelation function [5]), for $j = 1, \ldots, J$, which specify the optimal CV^2 and autocorrelations for each of the J MAP(2)s used in the KPC.

Step 2 - Moment and Higher-Order Correlation Fitting. Once that the optimal values of $CV^2(j)$ and $\gamma_2(j)$ are obtained after one or more runs of the previous step, we search for the missing parameters required to define valid MAP(2)s, namely the means E[X](j) and third moments $E[X^3](j)$ for all $j = 1, \ldots, J$. We spend the degrees of freedom in assigning $E[X](j)$ and $E[X^3](j)$ by solving in this step a new nonlinear optimization program that selects the MAP(2)s that result in the best fitting of a set of joint moments $\widehat{H}(\vec{i}, \vec{k})$. In particular, we focus on the joint moments $E[X_k X_{k+j} X_{k+j+h}]$, i.e., the bicorrelation coefficients, that can be computed efficiently from the trace and describe higher-order properties of the measurements not explicitly captured by the autocorrelation coefficients ρ_k [6].

Step 3 - MAP(n) Generation. Finally, given the target optimal values for the $E[X](j)$, $CV^2(j)$, $E[X^3](j)$, and $\gamma_2(j)$ we generate J MAP(2)s as follows. We use standard fitting algorithms, such as the one in [7], to determine D_0 and D_1 from moments and autocorrelations for each of the J MAP(2)s. Whenever this MAP(2) fitting results in an infeasible process (i.e., complex or negative rates in D_1 or in the off-diagonal elements of D_0), we perform a least square fitting to best match the target $E[X^3](j)$ and $\gamma_2(j)$ while keeping $E[X](j)$ and $CV^2(j)$ fixed. Once that J feasible MAP(2)s are obtained in this way, the final process is immediately computed by KPC.

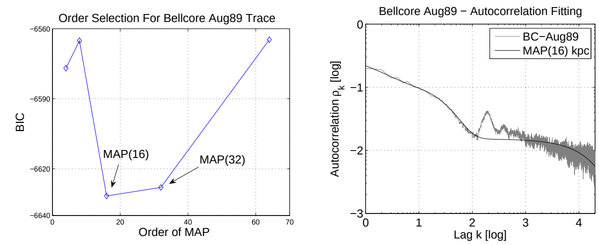

Figure 1: Order selection results

3. CASE STUDY

We illustrate the fitting procedure implemented in the KPC Toolbox using two case studies. In the first case study, we show the effectiveness of MAP order selection on a packet traffic trace. In the second case study, we give a quantitative idea of the quality of the proposed KPC-based fitting by modeling a real workload trace with temporal dependence and a MAP queueing system to reproduce the performance behavior of the original trace as obtained by simulation. In particular, we illustrate the results of fitting using the following traces:

- Bellcore Aug89 trace: This is a benchmark case for evaluating the quality of long-range trace fitting [1, 8, 9]. The traffic consists of 1 million interarrival time samples collected in 1989 at the Bellcore Morriston Research and Engineering facility. We use this trace to illustrate the *BIC* order selection.

- Seagate Web trace: This trace is collected from HTTP Web traffic with 3.6 million interarrival times of requests at the storage system of a Web server, see [11] for a description of this trace and related analyses of its temporal dependence structure. We use this trace to illustrate the KPC-based fitting.

In the next subsections, we discuss the two case studies on order selection and KPC-based MAP fitting.

3.1 Order Selection Results

We select the MAP order n for the Bellcore Aug89 trace using the *BIC* criterion. We evaluate MAP orders ranging in $n = \{4, 8, 16, 32, 64\}$. Figure 1 plots in the left diagram the *BIC* value versus the MAP order. The *BIC* indicates

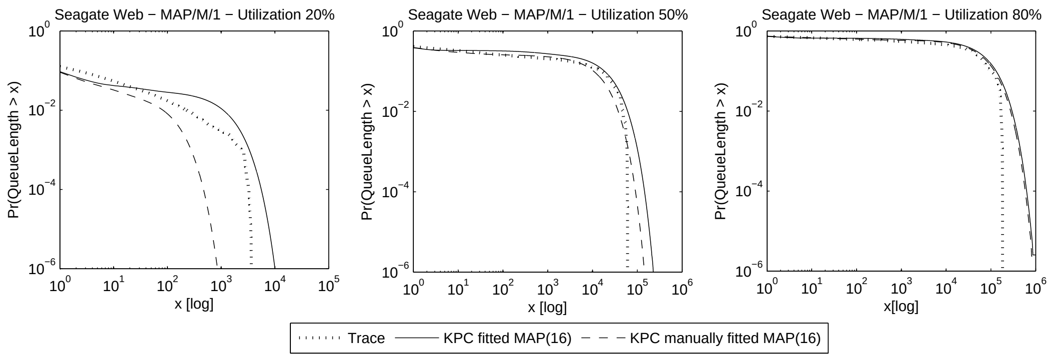

Figure 2: Queuing results for comparison between the Seagate Web trace and the fitted MAP

that the best orders for fitting the Bellcore Aug89 trace are $n = 16$ and $n = 32$. This is significant, because all the best MAPs fitted in the literature for the Bellcore Aug89 traces have exactly 16 or 32 states [5]. We select the order $n = 16$ since the BIC values for $n = 16$ and $n = 32$ are comparable and a MAP(16) has a smaller state space than a MAP(32). We then fit the MAP(16) by KPC of four MAP(2)s. In the right graph in Figure 1, we compare the autocorrelation of the real trace and of the fitted MAP(16). The fitted curve follows well the trace autocorrelations confirming the good quality of the order selection. We were unable to generate results of similar quality using MAPs of different order such as a MAP(8).

3.2 Queuing Prediction Results

We illustrate in Figure 2 the queuing prediction ability of the MAP fitted on the Seagate Web trace. The order selection on this trace returns $n = 16$ as the optimal MAP size. In the figure, we present the complementary cumulative distribution function (CCDF) of the queue length probabilities for two $MAP/M/1$ queues and the empirical CCDF obtained by simulating the $Trace/M/1$ queue. A set of results for one $MAP/M/1$ queue is obtained using as MAP the process fitted by KPC automatically, while the other $MAP/M/1$ queue uses the manually-fitted MAP obtained in [5]. The CCDFs of the $MAP/M/1$ queue are obtained by solving the underlying quasi-birth death process by 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.

At 20% utilization, the fitted MAP(16) overestimates the queuing probabilities. However it is still better than the manually fitted MAP(16). As utilization increases, the accuracy of queuing prediction improves. At 50% utilization, the fitted MAP captures the small and medium queue lengths better than the 20% utilization level. At 80% utilization, the fitted MAP(16) predicts well the queueing probabilities up to $x = 2 \cdot 10^5$.

4. CONCLUSION

We have presented the KPC-Toolbox, a set of MATLAB scripts for fitting workload traces into MAPs. The challenges of MAP fitting are to decide the order of the MAP and to determine the set of MAP jumping rates that yield a process that best fits the trace. The KPC-Toolbox meets the above challenges with a novel approach that uses the BIC criterion to estimate the optimal MAP order and by using optimization to explore a vast parameter space of alternative processes and find the MAP that best captures the essential descriptors of the trace. We have reported experimental results on real traces from both the systems and networking domains that show the effectiveness of the KPC-Toolbox. The KPC-Toolbox is available for downloading at *http://www.cs.wm.edu/MAPQN/kpctoolbox.html*.

5. 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.

[3] J. Brewer. Kronecker products and matrix calculus in system theory. *IEEE T. Circ. 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*, 13–24, 2008.

[5] G. Casale, E. Z. Zhang, and E. Smirni. Interarrival times characterization and fitting for Markovian traffic analysis. *Under Submission*, 2007.

[6] G. Casale, E. Z. Zhang, and E. Smirni. KPC-Toolbox: Simple yet effective trace fitting using Markovian Arrival Processes. To appear in *Proc. of the 5th Intl. Conference on the Quantitative Evaluation of Systems (QEST)*, St. Malo, France, IEEE Press, September 2008.

[7] J. Diamond, A. Alfa. On approximating higher order MAPs with MAPs of order two. *Queueing Systems*, 34:269–288, 2000.

[8] 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*, Performance 2002 Tutorial Lectures, 405-434, LNCS 2459-Springer, 2002.

[9] 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 QEST, IEEE Press*, 124–133, 2005.

[10] 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.

[11] A. Riska and E. Riedel. Long-range dependence at the disk drive level. In *Proc. of QEST*, 41–50. IEEE Press, 2006.

[12] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.

[13] 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.