Building Accurate Workload Models Using Markovian Arrival Processes

Giuliano Casale*
Imperial College London
Department of Computing
London, SW7 2AZ, UK
g.casale@imperial.ac.uk

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Modeling Techniques

General Terms

Algorithms, Performance, Theory

Keywords

Workload models, Markovian arrival process, fitting techniques

1. TUTORIAL OVERVIEW

The application of Markov chains and queueing theory to real performance evaluation problems often raises the question on how to best integrate in the model the observed characteristics of a workload. For example, what if job inter-arrival times to the system are statistically correlated? How can this be described compactly in a Markov model? What if the service time distribution is heavy-tailed? What if arrivals are periodic? Markovian arrival processes (MAPs) offer an elegant solution to these questions using the familiar framework of Markov theory [6]. MAPs have been developed with the aim of fitting in a compact Markov model workloads with statistical correlations and non-exponential distributions. This compact model can be then embedded in the infinitesimal generator of a performance model for a real system to represent events that occur with non-exponential inter-arrival times. For example, using matrix-geometric methods, one can easily study $MAP/MAP/1$ queues, where both inter-arrival times and service times are MAPs that describe temporal dependent processes [6]. MAPs also provide greater flexibility and realism compared to independent and identically distributed (i.i.d.) workload models such as Erlang, Coxian, or hyper-exponential; i.i.d. models describe the statistical distribution of inter-arrival times between events, but not their temporal order. Thus they cannot represent features such as burstiness, periodicities, or variability at multiple time scales that are frequent in computer workloads and network traffic [4, 7, 5]. MAPs have been explicitly developed to overcome this limitation of i.i.d. models.

Considered as a time-series modeling technique, MAPs may be seen as a class of hidden Markov models where observations depend on the state of an underlying continuous-time Markov chain (CTMC). Observations are continuous values that follow a phase-type distribution [6], which is the class of distributions describing the time to reach an absorbing state in a CTMC. The phase-type class is very flexible, e.g. it is a super-set of hypo-exponential, exponential, and hyper-exponential distributions that are popular models of workloads with different degrees of variability. As for other hidden Markov models, the expectation-maximization (EM) algorithm allows to parameterize the MAP to fit distribution and statistical order of observations in a measured time series [2].

When embedded in a system performance model, a MAP becomes a device to represent events that occur with non-exponential inter-arrival times. The nature of the event depends on the application, but usually in a queueing model it represents the arrival or the completion of a job. Events are defined as follows. State transitions in a MAP are classified as either *hidden* transitions, which do not lead to the occurrence of an event, or as *observable* transitions, which conversely represent the occurrence of the event modeled by the MAP. Both type of transitions can change the current state of the MAP and of the system performance model in which it is embedded. Hidden transitions may have an associated semantic as well, such server breakdown/repair, or may simply delay the occurrence of an observable event by moving to a state where observable transitions have different rates.

Finally, MAPs are also useful as a theoretical tool to investigate inter-arrival times of events that occur in a stochastic model. Given the CTMC underlying a performance model for a system, one can distinguish the state transitions into hidden and observable ones in order to extract a MAP from the CTMC. Such a MAP characterizes the inter-arrival times of the events chosen to be observable. In particular, the main advantage is that this MAP provides closed-form expressions to study distribution and correlations of the observable events. For instance, given a queueing network model, it is simple to mark transitions corresponding to arrival events to a queue. This provides then a simple way to study distribution, correlations, and average performance indexes embedded at arrival instants of jobs.

1.1 Mathematical Model

A MAP is defined by two square matrices D_0 and D_1 such that $Q = D_0 + D_1$ is an irreducible infinitesimal generator for the CTMC underlying the process, and $D_0(i,j)$ (resp. $D_1(i,j)$) is the rate of hidden (resp. observable) transitions from state i to state j. For example, a MAP(2) is a 2-state Markovian arrival process with

$$D_0 = \begin{bmatrix} -\sigma_1 & \lambda_{1,2} \\ \lambda_{2,1} & -\sigma_2 \end{bmatrix}, \quad D_1 = \begin{bmatrix} \mu_{1,1} & \mu_{1,2} \\ \mu_{2,1} & \mu_{2,2} \end{bmatrix}$$

where $\lambda_{i,j} \geq 0$, $\mu_{i,j} \geq 0$, for all i, j. The diagonal elements of D_0 are $\sigma_1 = \lambda_{1,2} + \mu_{1,1} + \mu_{1,2} > 0$ and $\sigma_2 = \lambda_{2,1} + \mu_{2,2} + \mu_{2,1} > 0$ such that the underlying CTMC Q has no absorbing states. The above process works as follow. First, the MAP is initialized in any of the two states according to a given probability distribution $\pi = [\pi_1, \pi_2]$, say state 1. After an exponentially distributed time with rate σ_1, it emits an observable event with probability

*The work of Giuliano Casale is supported by the Imperial College Junior Research Fellowship.

SIGMETRICS'11, June 7–11, 2011, San Jose, California, USA.
ACM 978-1-4503-0262-3/11/06.

$(\mu_{1,1} + \mu_{1,2})/\sigma_1$ and, if the event is generated, it instantaneously moves to state 2 with probability $\mu_{1,2}/(\mu_{1,1} + \mu_{1,2})$ or otherwise stays in state 1. If the event is not generated, the MAP moves into state 2 according to the hidden transition rate $\lambda_{1,2}$. Similar definitions hold for state 2. Note that after generating the event, the MAP does not need to be reinitialized and it evolves according to the above rules to generate additional events. MAPs with more than 2 states behave similarly, but probabilities are used also for the hidden transitions in order to select the destination state.

Let us first show that MAPs are a super-set of i.i.d. workload models. For instance, an exponential distribution with rate μ is readily specified as $D_0 = [-\mu]$, $D_1 = [\mu]$. A two stage hypo-exponential distribution with rates μ_1 and μ_2 may be represented as

$$D_0 = \begin{bmatrix} -\mu_1 & \mu_1 \\ 0 & -\mu_2 \end{bmatrix}, \quad D_1 = \begin{bmatrix} 0 & 0 \\ \mu_2 & 0 \end{bmatrix}$$

Similarly, a two-phase hyper-exponential distribution with probability $p > 0$ of selecting the first phase has MAP representation

$$D_0 = \begin{bmatrix} -\mu_1 & 0 \\ 0 & -\mu_2 \end{bmatrix}, \quad D_1 = \begin{bmatrix} p\mu_1 & (1-p)\mu_1 \\ p\mu_2 & (1-p)\mu_2 \end{bmatrix}$$

where after generating an event, the next event will be from state 1 with probability p and from state 2 with probability $1 - p$.

Notice that, in the above examples, the state in which the MAP moves after an observable event is either fixed or independent of the last state visited before generating the event. This makes the inter-arrival times between events statistically independent. The additional flexibility of MAPs compared to phase-type distribution models is that one can vary the rates in D_1 in order to introduce dependence between inter-arrival times. For example,

$$D_0 = \begin{bmatrix} -\mu_1 & 0 \\ 0 & -\mu_2 \end{bmatrix}, \quad D_1 = \begin{bmatrix} p_1\mu_1 & (1-p_1)\mu_1 \\ p_2\mu_2 & (1-p_2)\mu_2 \end{bmatrix}$$

is a MAP with hyper-exponentially distributed samples generated according to rates μ_1 and μ_2, but where the phase selection probabilities $p_1 > 0$ and $p_2 > 0$ may now be chosen to introduce temporal dependence. For instance, $p_1 = 0.99$ and $p_2 = 0.02$ define a MAP where successive events tend to be sampled from the same phase for an extended period of time, thus creating bursts of arrivals in the fast phase and long periods with few arrivals in the slow phase. Conversely, $p_1 = p_2$ makes again the MAP an i.i.d. hyper-exponential workload model.

2. FITTING MAP WORKLOAD MODELS

Fitting the inter-arrival times of observable events to a measured trace requires to capture the properties of a time series in terms of distribution and statistical correlations between samples. A fundamental property of MAPs is to enjoy closed-form analytical expressions for the joint probability density of event inter-arrival times. This provides the starting point for defining closed-form formulas for statistical descriptors and then related workload fitting methods. The joint density for a MAP is

$$\begin{aligned}\Pr[X_0 = t_0, X_1 = t_1, \ldots, X_n = t_n] \\ = \pi e^{D_0 t_0} D_1 e^{D_0 t_1} D_1 \cdots e^{D_0 t_n} D_1 1\end{aligned}$$

where π is an initialization vector for the CTMC underlying the MAP, 1 stands for a vector of ones, and $e^A = \sum_{k=0}^{\infty} A^k/k!$ is the matrix exponential function for matrix A. It can be shown that inter-arrival times are stationary if $\pi(-D_0)^{-1}D_1 = \pi$. Moreover, it is know that matrix $(-D_0)^{-1}D_1$ is a discrete-time Markov chain and gives in element (i, j) the conditional probability that a sample generated from a MAP initialized in phase i is followed by a sample obtained starting the MAP in phase j. Thus, matrix $(-D_0)^{-1}D_1$ needs to be assigned accurately in order to fit the temporal dependent patterns of a measured time-series. Note that this defines a nonlinear search problem in the rates of D_0 and D_1. Similar expressions exist for a variety of statistical descriptors for inter-arrival times, such as moments of the distribution, joint moments, autocorrelations, index of dispersion [1, 2, 3, 4, 6]. However, fitting these descriptors requires nonlinear techniques as well.

In spite of the non-linearity of the fitting problem, several methods have been developed in recent years to fit a Markovian arrival process. These include, among others, techniques such as analytical and optimization-based moment matching, expectation maximization methods, and compositional methods based on the superposition or synchronization of multiple MAPs [1, 2, 3, 4]. Here, we limit to illustrate the synchronization-based technique proposed in [3], called Kronecker product composition.

Given two MAPs $A = (A_0, A_1)$ and $B = (B_0, B_1)$, consider the following process defined by synchronization of A and B

$$C = A \otimes B = (D_0, D_1) = (-A_0 \otimes B_0, A_1 \otimes B_1)$$

where $\otimes$ denotes the Kronecker product operator. It is possible to show that C is a valid MAP if at least one between A_0 and B_0 is a diagonal matrix. Process C defines a synchronization between the transition rates of A and B according to the semantics of the Kronecker product operator. That is, C tracks the individual states of A and B: if A (resp. B) is in a state where events are observed at a rate λ (resp. μ), then events are observed in C with rate $\lambda\mu$.

Stemming from the properties of Kronecker products it is then possible to show that moments and correlations for the process C are simple functions of the corresponding moments for A and B, e.g., $E[X_C^k] = E[X_A^k]E[X_B^k]/k!$, where X_C is the random variable denoting the inter-arrival times of events in process C and X_A and X_B are similarly defined for A and B. Similar decomposition formulas apply to autocorrelations and joint moments that describe the statistical order of observations in the trace. Thus, the synchronization method defines a divide-and-conquer approach to MAP fitting where one is concerned only with defining moments and correlations for small processes A and B in order to create by composition a process C that fits a trace. A MATLAB toolbox that implements this divide-and-conquer fitting approach is given in [3].

3. 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] P. Buchholz. An EM-Algorithm for MAP Fitting from Real Traffic Data. Springer LNCS, vol. 2794, 218–236.

[3] G. Casale, E. Z.Zhang, E. Smirni. KPC-Toolbox: Best Recipes for Automatic Trace Fitting Using Markovian Arrival Processes . *Perf. Eval.*, 67(9):873–896, Sep 2010.

[4] A. Horváth and M. Telek. Markovian modeling of real data traffic. In *Performance Evaluation of Complex Systems, LNCS Tutorial Vol 2459*, 405–434, 2002.

[5] 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.

[6] M. F. Neuts. *Structured Stochastic Matrices of M/G/1 Type and Their Applications*. Marcel Dekker, New York, 1989.

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