ETAQA Truncation Models for the MAP/MAP/1 Departure Process

Armin Heindl*
Computer Networks and Communication Systems
University Erlangen-Nuremberg
91058 Erlangen, Germany
Armin.Heindl@informatik.uni-erlangen.de

Qi Zhang Evgenia Smirni†
Department of Computer Science
College of William and Mary
Williamsburg, VA 23187-8795
{qizhang,esmirni}@cs.wm.edu

Abstract

We propose a family of finite approximations for the departure process of a MAP/MAP/1 queue. The departure process approximations are derived via an exact aggregate solution technique (called ETAQA) applied to Quasi-Birth-Death processes (QBDs) and require only the computation of the frequently sparse fundamental-period matrix $\mathbf{G}$*. The approximations are indexed by a parameter* n*, which determines the size of the output model as* $n + 1$ *QBD levels. The marginal distribution of the true departure process and the lag correlations of the interdeparture times up to lag* $n - 1$ *are preserved exactly. Via experimentation we show the applicability of the proposed approximation in traffic-based decomposition of queueing networks and investigate how correlation propagates through tandem queues.*

1. Introduction

Queueing networks are often used as models of complex computer and communication systems. In such systems, both the arrival and service processes may exhibit correlations – be they strong or weak, short-term or long-term –, and such correlations have been shown empirically and theoretically to have a significant impact on queueing behavior (e.g., [11, 17, 4]). Note that the presence of correlations complicates the solution of queueing networks, as classic analytic techniques cease to apply. One can resort to simulation for the solution of such systems. Approximate analysis techniques are an alternative to simulation and are based on analyzing each queue in isolation, which is also known as traffic-based decomposition (see [7] for an overview). Approximating the output traffic of a single queueing system such that its "relevant" properties are captured is crucial for the effectiveness of traffic-based decomposition. For example, traffic-based decomposition can be used to model e-commerce servers that are organized as three-tiered systems [1] (with web, application, and database servers forming three tandem queues).

In this paper, we employ Markovian arrival processes (MAPs) to approximate the departure process of a MAP/MAP/1 queue. MAPs are widely used in traffic engineering and workload modeling to match correlated arrival/service processes. Our output approximation captures the marginal distribution of the true departure process of a MAP/MAP/1 queue exactly and preserves the first coefficients of correlation of the lagged interdeparture intervals up to a given order, depending on the selectable size of the output model. Thus, in accordance with many other authors [6, 21, 12, 9], we focus on the marginal distribution and the autocorrelation structure as the "relevant" properties in traffic-based decomposition.

Independent of traffic-based decomposition of queueing networks, one may be interested in the impact of different arrival and/or service processes on the departure process of a MAP/MAP/1 queue, e.g., for the purpose of traffic shaping. This can be studied with our output models – and is studied rudimentarily in the experiments we give in this paper. Furthermore, our models can assist in examining the propagation of arrival/service autocorrelation, e.g., through a tandem network.

1.1. Related work

It is well known that the departure process of a MAP/MAP/1 queue is generally nonrenewal (the M/M/1 queue being an exception). More precisely, the departure process can be described exactly as a MAP with an infinite number of states, where the Quasi-Birth-Death process (QBD) of the MAP/MAP/1 queue is filtered according to transitions causing a departure or not. Note that the

* Armin Heindl has been supported by DFG under grant Ho1257/17, while he has been on leave from the Technische Universität Berlin, Institut für Technische Informatik und Mikroelektronik.

† Evgenia Smirni has been partially supported by the National Science Foundation under grants CCR-0098278 and ACI-0090221.

Markov chain of the QBD is infinite, therefore one may resort to truncation of the infinite MAP representation in order to find an analytically tractable finite-state approximation of this true departure process. Different truncation models have been proposed and our output model also falls within this group. In the following, we briefly present truncation approaches from the literature, which also capture the interdeparture distribution and the first lag coefficients of correlation of the departure process.

Kumaran et al. [9] suggest a truncation model for the departure process of an ME/ME/1 queue. An ME/ME/1 queue (ME stands for the matrix-exponential distribution) is (slightly) more general than a PH/PH/1 queue (PH stands for phase-type distribution), but entails no correlations – neither for service nor for interarrival times. However, the truncation model preserves the autocorrelations in the departure process up to an index defined by the model size (i.e., by the truncation level of the QBD). Compellingly, the matrices of the output model (namely a correlated matrix-exponential process, whose notation strongly resembles that of a MAP) are directly given in terms of the vectors/matrices that specify the interarrival and the service distribution. Consequently, this model comes essentially at no cost, but the employed renewal assumptions for service and arrivals prevent its generalization to correlated processes.

The truncation models proposed by Green [5, 6] work for MAP/PH/1 queues (i.e., when there is no correlation among service times). The approximations of the departure process reflect the busy-period behavior of the queue, which however requires the computation of the departure-epoch queue length distribution, which is an expensive computation. Regarding the preserved properties of the output models, like marginal distribution and correlation coefficients, thorough proofs can be found in [5].

The family of truncation models proposed in [21] also claims to preserve the marginal distribution and parts of the correlation structure. The departure process of a MAP/MAP/1 queue may not only be truncated, but also lower levels (as they arise from the QBD structure of the queue) may be condensed based on flow arguments. Still, parts of the stationary queue length distribution have to be computed to perform the approximation.

Our truncation models also apply to the MAP/MAP/1 queue and therein are closest to [21]. We preserve the same characteristics of the true departure process as the techniques mentioned above. But our approximation is more general than the ones reported in [9, 6], as it does capture correlations in both the arrival and service processes. Furthermore, the presented approximation is significantly more efficient than those reported in [21, 6].

1.2. The ETAQA approach

Recently, an aggregation method for M/G/1-type processes, called ETAQA, has been proposed [18, 20]. This method directly applies to the solution of QBDs, as QBDs are a special case of M/G/1-type processes. ETAQA exactly preserves the level distribution, i.e., the stationary probability distribution of the MAP/MAP/1 queue in our setting, for the non-aggregated levels. This aggregated and thus finite-dimensional representation may be filtered to derive an output approximation of the MAP/MAP/1 queue of the same dimension. We show that this output MAP has the desired properties regarding the marginal distribution and autocorrelation structure. Moreover, the block elements of the MAP notation are given directly from the arrival and service representations and the fundamental-period matrix $\mathbf{G}$. In this paper we focus on MAP/MAP/1 queues only, but the proposed approximation is easily generalized to queues with a QBD representation (e.g., a Σ_i MAP$_i$/MAP/m queue with multiple inputs and multiple servers). Although here we restrict our discussions to continuous-time queues, the truncation technique can be directly adapted to discrete-time queues.

The remainder of the paper is organized as follows: Section 2 briefly recalls the definitions of MAPs and QBDs and summarizes the exact aggregate solution of QBDs. In Section 3, we construct our family of output approximations and prove some of their properties. Numerical examples in Section 4 demonstrate the applicability of these models in traffic-based decomposition. Section 5 discusses the computational complexity of our method. Section 6 recapitulates our contributions and outlines future work.

2. Background

Here, we recall the definition and properties of Markovian arrivals processes (MAPs) and cite a theorem on the aggregate solution of Quasi-Birth-Death processes (QBDs) upon which our departure process approximation builds.

2.1. Markovian Arrival Processes (MAPs)

Informally, MAPs are ergodic Continuous-Time Markov Chains (CTMCs), in which transitions are distinguished by whether they cause an event or not. The rates are correspondingly grouped into two matrices $\mathbf{D}_1$ and $\mathbf{D}_0$, such that

$\mathbf{D}_1$ is a nonnegative ($m_{\text{MAP}} \times m_{\text{MAP}}$)-rate matrix, where m_{MAP} is the order of the MAP.

$\mathbf{D}_0$ is a matrix of the same dimension as $\mathbf{D}_1$, with negative diagonal and nonnegative off-diagonal elements.

The matrix $\mathbf{Q}_{\text{MAP}} = \mathbf{D}_0 + \mathbf{D}_1$ is the irreducible infinitesimal generator of the CTMC, where $\mathbf{D}_0$ governs transitions that do not correspond to events, while $\mathbf{D}_1$ governs those transitions that do correspond to events. Let $\boldsymbol{\pi}_{\text{MAP}}$ be the stationary probability vector of the CTMC generator (i.e., $\boldsymbol{\pi}_{\text{MAP}}\mathbf{Q}_{\text{MAP}} = \mathbf{0}, \boldsymbol{\pi}_{\text{MAP}}\mathbf{e} = 1$, where $\mathbf{0}$ and $\mathbf{e}$ denote the vectors of zeros or ones of the appropriate dimension). Special cases of MAPs include Markov-modulated Poisson processes (MMPPs) and PH-type renewal processes. From $\mathbf{D}_0$ and $\mathbf{D}_1$ a variety of performance measures can be computed. The arrival rate and the squared coefficient of variation (of the PH-type marginal distribution) of the MAP with interevent time X are given by

$$\begin{aligned} \lambda_{\text{MAP}} &= \boldsymbol{\pi}_{\text{MAP}}\mathbf{D}_1\mathbf{e} \quad , & (1)\\ c^2_{\text{MAP}} &= \frac{E[X^2]}{(E[X])^2} - 1 = 2\lambda_{\text{MAP}}\boldsymbol{\pi}_{\text{MAP}}(-\mathbf{D}_0)^{-1}\mathbf{e} - 1 \; . & (2) \end{aligned}$$

The autocorrelation function of a stationary MAP, i.e., the lag-k coefficients of correlation ($k > 0$), is computed as [14]

$$\begin{aligned} \text{corr}[X_0, X_k] &= \frac{E[(X_0 - E[X])(X_k - E[X])]}{\text{Var}[X]} & (3)\\ &= \frac{\lambda_{\text{MAP}}\boldsymbol{\pi}_{\text{MAP}}((-\mathbf{D}_0)^{-1}\mathbf{D}_1)^k(-\mathbf{D}_0)^{-1}\mathbf{e} - 1}{2\lambda_{\text{MAP}}\boldsymbol{\pi}_{\text{MAP}}(-\mathbf{D}_0)^{-1}\mathbf{e} - 1} \; , \end{aligned}$$

where X_0 and X_k denote two interevent times k lags apart.

2.2. Quasi-Birth-Death processes (QBDs)

The solution of a MAP/MAP/1 queue translates to the solution of a QBD. The infinitesimal generator $\mathbf{Q}_\infty$ of a QBD is block tridiagonal. Generally, the state space $\mathcal{S}$ of such a CTMC is partitioned into levels: the boundary states $\mathcal{S}^{(0)} = \{s_1^{(0)}, \ldots, s_{m_0}^{(0)}\}$, and the sets of states $\mathcal{S}^{(j)} = \{s_1^{(j)}, \ldots, s_m^{(j)}\}$, for $j \geq 1$ and $m_0, m \geq 1$. Let $\boldsymbol{\pi}^{(0)}$ and $\boldsymbol{\pi}^{(j)}$ for $j \geq 1$ be the stationary probability vectors (of dimensions m_0 and m, respectively) for states in $\mathcal{S}^{(0)}$ and $\mathcal{S}^{(j)}$, respectively. Intuitively, $\mathcal{S}^{(0)}$ represents the state configuration when the queue is empty. The states that account for the state of the system when the queue is nonempty (with j customers) correspond to sets $\mathcal{S}^{(j)}$, for $j \geq 1$, and the interaction of successive sets has a "repetitive" structure.

For the general case of QBDs, the infinitesimal generator $\mathbf{Q}_\infty$ has the following form:

$$\mathbf{Q}_\infty = \begin{bmatrix} \widehat{\mathbf{L}} & \widehat{\mathbf{F}} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \widehat{\mathbf{B}} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{L} & \mathbf{F} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad . \tag{4}$$

We use the letters "L", "F" and "B" according to whether they describe "local", "forward" and "backward" transition rates, respectively, in relation to a set of states $\mathcal{S}^{(j)}$ for $j \geq 0$, and a "$\widehat{\ }$" for matrices related to $\mathcal{S}^{(0)}$. Key to the computation of the stationary probability vector

$$\boldsymbol{\pi}_\infty = \begin{bmatrix} \boldsymbol{\pi}^{(0)} & \boldsymbol{\pi}^{(1)} & \ldots \end{bmatrix} \quad , \tag{5}$$

defined by $\boldsymbol{\pi}_\infty\mathbf{Q}_\infty = \mathbf{0}$ and $\boldsymbol{\pi}_\infty\mathbf{e} = 1$, is the fact that a geometric relation holds among the vector $\boldsymbol{\pi}^{(j)}$ for $j \geq 1$,

$$\boldsymbol{\pi}^{(j)} = \boldsymbol{\pi}^{(1)} \cdot \mathbf{R}^{j-1}, \quad j \geq 1 \; . \tag{6}$$

Matrix $\mathbf{R}$, the geometric coefficient, is the solution of the matrix equation

$$\mathbf{F} + \mathbf{R} \cdot \mathbf{L} + \mathbf{R}^2 \cdot \mathbf{B} = \mathbf{0}$$

and has an important probabilistic interpretation: the entry (k, l) of $\mathbf{R}$ is the expected time spent in state l of $\mathcal{S}^{(i)}$, before the first visit into $\mathcal{S}^{(i-1)}$, expressed in time unit $\Delta_k^{(i-1)}$, given the starting state in $\mathcal{S}^{(i-1)}$ is k. $\Delta_k^{(i-1)}$ is the mean sojourn time in state k of $\mathcal{S}^{(i-1)}$ for $i \geq 2$ [15, pages 30-35]. $\mathbf{R}$ can be obtained using iterative numerical algorithms the most well-known one being the logarithmic reduction algorithm proposed by Latouche and Ramaswami [10]. $\boldsymbol{\pi}^{(0)}$ and $\boldsymbol{\pi}^{(1)}$ are obtained by solving the following system of $m_0 + m$ equations:

$$[\boldsymbol{\pi}^{(0)} \; \boldsymbol{\pi}^{(1)}] \cdot \left[\begin{array}{cc|c} \widehat{\mathbf{L}} & \widehat{\mathbf{F}} & \mathbf{e} \\ \widehat{\mathbf{B}} & \mathbf{L} + \mathbf{R} \cdot \mathbf{B} & (\mathbf{I} - \mathbf{R})^{-1} \cdot \mathbf{e} \end{array} \right] = [\mathbf{0} \,|\, 1], \tag{7}$$

where the last column in the matrix corresponds to normalization, which replaces any one of the other equations.

Another matrix that can greatly facilitate the computation of the probability distribution of all states in QBDs is $\mathbf{G}$ [10], solution of:

$$\mathbf{B} + \mathbf{L} \cdot \mathbf{G} + \mathbf{F} \cdot \mathbf{G}^2 = \mathbf{0} \quad .$$

Matrix $\mathbf{G}$ also has an important probabilistic interpretation: an entry (l, k) in $\mathbf{G}$ expresses the conditional probability of the process first entering $\mathcal{S}^{(i-1)}$ through state k, given that it starts from state l of $\mathcal{S}^{(i)}$ [16, page 81]. The following fundamental relation between matrices $\mathbf{R}$ and $\mathbf{G}$ holds [10, pages 137-8]:

$$\mathbf{R} \cdot \mathbf{B} = \mathbf{F} \cdot \mathbf{G} \quad . \tag{8}$$

Obviously, for QBDs, knowing $\mathbf{G}$ (or $\mathbf{R}$) implies a direct computation of $\mathbf{R}$ (or $\mathbf{G}$). Computing $\mathbf{G}$ is usually easier than computing $\mathbf{R}$.

In [18], ETAQA, a new methodology for the exact analysis of M/G/1-type Markov processes was proposed. Originally, ETAQA truncates these infinite Markov processes on block level $n = 2$ in such a way that the stationary level distributions $\boldsymbol{\pi}^{(0)}$ and $\boldsymbol{\pi}^{(1)}$ are preserved. The main theorem for the solution of M/G/1-type processes can be directly applied for the solution of QBDs, since QBDs are a special case of M/G/1-type processes.

Theorem 2.1 [ETAQA for QBDs] Given a QBD with infinitesimal generator $\mathbf{Q}_\infty$ (4) and with stationary probability vector $\boldsymbol{\pi}_\infty$ (5), the system of linear equations

$$\tilde{\boldsymbol{\pi}} \cdot \tilde{\mathbf{Q}} = \mathbf{0} \quad ,$$

where $\tilde{\mathbf{Q}} \in \mathbb{R}^{(m_0+2m)\times(m_0+2m)}$ is defined by

$$\tilde{\mathbf{Q}} = \begin{bmatrix} \widehat{\mathbf{L}} & \widehat{\mathbf{F}} & \mathbf{0} \\ \widehat{\mathbf{B}} & \mathbf{L} & \mathbf{F} \\ \mathbf{0} & \mathbf{B}-\mathbf{FG} & \mathbf{F}+\mathbf{L}+\mathbf{FG} \end{bmatrix} \quad ,$$

admits a unique solution $\tilde{\boldsymbol{\pi}} = [\boldsymbol{\pi}^{(0)}\ \boldsymbol{\pi}^{(1)}\ \boldsymbol{\pi}^{(*)}]$, where $\boldsymbol{\pi}^{(*)} = \sum_{i=2}^{\infty} \boldsymbol{\pi}^{(i)}$, given that we discard one column (any) and we replace it with a column of 1's due to the normalization condition $\tilde{\boldsymbol{\pi}}\mathbf{e} = 1$.

Proof: The proof follows directly from the main theorem in [18]. Figure 1(b) illustrates the high-level idea that ETAQA is based on, i.e., the aggregation of the repeating portion of the chain from set $\mathcal{S}^{(2)}$ and onward. □

In particular for MAP/MAP/1 queues, which are the focus of this paper, aggregation can occur for *any* level $n \geq 1$ (i.e., including $n = 1$, see Figure 1(c)). This is because the block matrices $\widehat{\mathbf{B}}$ and $\widehat{\mathbf{F}}$ are identical to $\mathbf{B}$ and $\mathbf{F}$, respectively. Using Kronecker notation, the block matrices are defined by:

$$\begin{aligned} \widehat{\mathbf{L}} &= \mathbf{D}_0^{(A)} \otimes \mathbf{I}_S \\ \mathbf{L} &= \mathbf{D}_0^{(A)} \oplus \mathbf{D}_0^{(S)} = \mathbf{D}_0^{(A)} \otimes \mathbf{I}_S + \mathbf{I}_A \otimes \mathbf{D}_0^{(S)} \\ \mathbf{B} &= \mathbf{I}_A \otimes \mathbf{D}_1^{(S)} \ \ (= \widehat{\mathbf{B}}\,) \\ \mathbf{F} &= \mathbf{D}_1^{(A)} \otimes \mathbf{I}_S \ \ (= \widehat{\mathbf{F}}\,) \quad , \end{aligned}$$

where $\mathbf{D}_0^{(A)}$ and $\mathbf{D}_1^{(A)}$ describe the MAP of the arrival process of order m_A and $\mathbf{D}_0^{(S)}$ and $\mathbf{D}_1^{(S)}$ describe the MAP of the service process of order m_S. In the remainder of this paper, we will keep the conventions $\widehat{\mathbf{B}} = \mathbf{B}$ and $\widehat{\mathbf{F}} = \mathbf{F}$, which essentially imply that *all* matrices $\mathbf{B}$, $\mathbf{F}$, $\mathbf{L}$ and $\widehat{\mathbf{L}}$ are square $(m \times m)$-matrices, and that vectors $\boldsymbol{\pi}^{(j)}$ for $j \geq 0$ have the same dimension m, where $m = m_A m_S$.

Thus, the resulting QBD can be solved according to the following lemma (where $n \geq 1$).

Lemma 2.1 Given an ergodic MAP/MAP/1 CTMC with infinitesimal generator $\mathbf{Q}_\infty$ (see (4), where $\widehat{\mathbf{B}} = \mathbf{B}$ and $\widehat{\mathbf{F}} = \mathbf{F}$) and with stationary probability vector $\boldsymbol{\pi}_\infty$ (5), the system of linear equations (parameterized with n)

$$\boldsymbol{\pi}_n \cdot \mathbf{Q}_n = \mathbf{0} \quad ,$$

where $\mathbf{Q}_n \in \mathbb{R}^{(n+1)m\times(n+1)m}$ is defined by

$$\mathbf{Q}_n = \begin{bmatrix} \widehat{\mathbf{L}} & \mathbf{F} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{B} & \mathbf{L} & \mathbf{F} & \ddots & \vdots \\ \mathbf{0} & \ddots & \ddots & \ddots & \mathbf{0} \\ \vdots & \ddots & \mathbf{B} & \mathbf{L} & \mathbf{F} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{B}-\mathbf{FG} & \mathbf{L}+\mathbf{F}+\mathbf{FG} \end{bmatrix} \quad , \tag{9}$$

admits a unique solution

$$\boldsymbol{\pi}_n = \left[\ \boldsymbol{\pi}^{(0)}\ \ \boldsymbol{\pi}^{(1)}\ \ \ldots\ \ \boldsymbol{\pi}^{(n-1)}\ \ \boldsymbol{\pi}^{(n,*)}\ \right] \quad ,$$

where $\boldsymbol{\pi}^{(n,*)} = \sum_{i=n}^{\infty} \boldsymbol{\pi}^{(i)}$, given that we discard one column (any) and replace it with a column of 1's due to the normalization condition, i.e., $\boldsymbol{\pi}_n \mathbf{e} = 1$.

Note again that the case $n = 1$ with two block levels only (namely 0 and 1) is explicitly included.

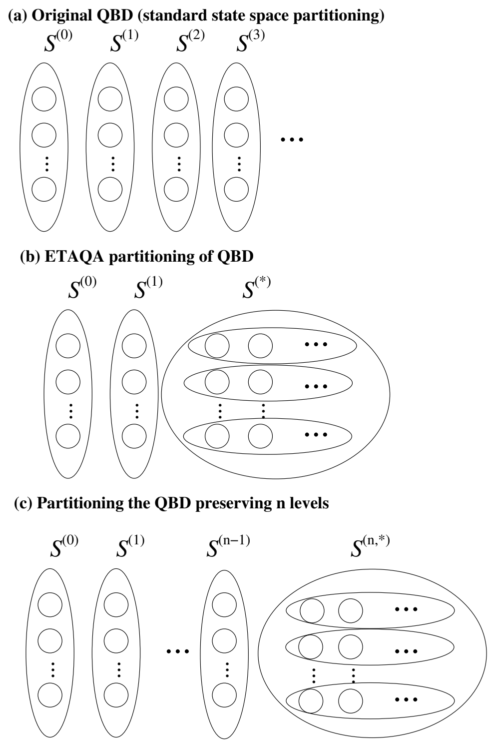

Figure 1. State-space partitioning of a QBD solved via the traditional matrix-analytic method (a), via ETAQA (b), and ETAQA that preserves the first n QBD levels, $n \geq 1$ (c).

3. Output process of the MAP/MAP/1 queue and its approximation

We apply filtration to the QBD of the MAP/MAP/1 queue and its aggregated representation to obtain the true departure process of this queueing system and a flexible approximation thereof. We also prove that the marginal distribution of the departure process is preserved.

3.1. Exact departure process

Starting from the infinitesimal generator $\mathbf{Q}_\infty$ (4) (with $\widehat{\mathbf{B}} = \mathbf{B}$ and $\widehat{\mathbf{F}} = \mathbf{F}$), we give the exact departure process of a MAP/MAP/1 queue as a MAP of infinite order. By so-called filtration (see [5]), i.e., by collecting in matrix $\mathbf{D}_1$ "backward" transitions of $\mathbf{Q}_\infty$ that correspond to departures, we arrive at the following MAP representation:

$$\mathbf{D}_{0,\infty}^{(D)} = \begin{bmatrix} \widehat{\mathbf{L}} & \mathbf{F} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{L} & \mathbf{F} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{L} & \mathbf{F} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} , \tag{10}$$

$$\mathbf{D}_{1,\infty}^{(D)} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{B} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{B} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{B} & \mathbf{0} & \mathbf{0} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} . \tag{11}$$

Despite its infinite order (which of course is impractical for further processing in network decomposition), several performance measures can be computed from this representation. We will dwell on this in the next subsection, when we compare properties for both the exact departure process and its truncated approximation model.

3.2. Approximations for the output process

Obviously, one way to obtain a tractable MAP approximation of the MAP/MAP/1 departure process is to truncate the infinite representation (10)/(11). This, however, can be done in many different ways (e.g., see the discussions in Section 1.1 or in [2]). In any case, the last block row (the level may be chosen arbitrarily) will have to be adjusted so that a properly defined MAP arises. Careless truncations will fudge desirable properties of the departures process.

The ETAQA truncation $\mathbf{Q}_n$ (9) of the QBD of the MAP/MAP/1 queue points at novel possibilities to achieve an appropriate truncation (at arbitrary level n, $n \geq 1$) for the departure process. At the same time, Lemma 2.1 guarantees that the stationary distribution of the true departure process (and of the original QBD) is preserved up to block index $n-1$. Of course, similar invariance properties for the marginal distribution and the autocorrelation function for both MAPs would be of eminent importance in our setting. We will see that the following output truncation model (with $n+1$ block levels) enjoys such favorable features and, moreover, admits an intuitive stochastic interpretation:

$$\mathbf{D}_{0,n}^{(D)} = \begin{bmatrix} \widehat{\mathbf{L}} & \mathbf{F} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{L} & \mathbf{F} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} & \mathbf{L} & \mathbf{F} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} & \mathbf{L}+\mathbf{F} \end{bmatrix} , \tag{12}$$

$$\mathbf{D}_{1,n}^{(D)} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \mathbf{B} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B}-\mathbf{FG} & \mathbf{FG} \end{bmatrix} . \tag{13}$$

Index n ($n \geq 1$) indicates that the bottom block row may be chosen arbitrarily (with the exception of level 0). The order of this truncated MAP is $(n+1)m = (n+1)m_S m_A$.

To make evident the stochastic underpinning, we apply (8) and (6) to point out that

$$\begin{aligned} \boldsymbol{\pi}^{(n,*)}(\mathbf{B}-\mathbf{FG}) &= \boldsymbol{\pi}^{(n,*)}(\mathbf{I}-\mathbf{R})\mathbf{B} = \boldsymbol{\pi}^{(n)}\mathbf{B} , \\ \boldsymbol{\pi}^{(n,*)}\mathbf{FG} &= \boldsymbol{\pi}^{(n,*)}\mathbf{RB} = (\Sigma_{i=n+1}^{\infty}\boldsymbol{\pi}^{(i)})\mathbf{B} . \end{aligned}$$

Note that $\boldsymbol{\pi}_\infty$ (see (5)) and $\boldsymbol{\pi}_n$ (see Lemma (2.1)) are the stationary distributions of both the respective output MAP and the corresponding (infinite or truncated) QBD. Thus, the block matrix $\mathbf{FG}$ in (13) captures the flow backward *within* the aggregate state encompassing original levels n to ∞, whereas $\mathbf{B}-\mathbf{FG}$ corresponds to the flow that actually leads from the aggregate state to level $n-1$.

This stochastic interpretation supports the validity of other qualitative and quantitative properties, in particular that the interdeparture time distribution and a finite number of lag-k correlation coefficients ($0 < k < n$) are preserved. Below, we prove that the arrival rate remains unchanged.

Using (1), we may compute the arrival rate of a MAP regardless of its dimension. Therefore, the corresponding identity for both the infinite representation (10)/(11) and the truncation model (12)/(13) follows from

$$\begin{aligned} \lambda_\infty &= \boldsymbol{\pi}_\infty\mathbf{D}_{1,\infty}^{(D)}\mathbf{e} = \begin{bmatrix} \boldsymbol{\pi}^{(1)}\mathbf{B} & \boldsymbol{\pi}^{(2)}\mathbf{B} & \ldots \end{bmatrix}\mathbf{e} \\ &= (\Sigma_{i=1}^{\infty}\boldsymbol{\pi}^{(i)})\mathbf{B}\,\mathbf{e} , \\ \lambda_n &= \boldsymbol{\pi}_n\mathbf{D}_{1,n}^{(D)}\mathbf{e} = \big[\ \boldsymbol{\pi}^{(1)}\mathbf{B} \quad \ldots \\ &\qquad \boldsymbol{\pi}^{(n-1)}\mathbf{B} \quad \boldsymbol{\pi}^{(n,*)}(\mathbf{B}-\mathbf{FG}) \quad \boldsymbol{\pi}^{(n,*)}\mathbf{FG}\ \big]\mathbf{e} \\ &= (\Sigma_{i=1}^{\infty}\boldsymbol{\pi}^{(i)})\mathbf{B}\,\mathbf{e} . \end{aligned}$$

4. Experimental results

In this section, we present a set of experimental results that show the effectiveness of the family of approximations of the previous section. For all experiments, we use dual tandem queues (see Figure 2). Inter alia, the purpose of the experiments is to illustrate that a level-n approximation of the departure process captures the exact lag coefficients up to $n-1$ for $n \geq 1$. For all experiments we show the autocorrelation function (ACF) of the arrival process to the tandem queue (i.e., at point "A", see Figure 2) and the ACF of the departure process of the first queue for different approximation levels n (i.e., at point "B" in Figure 2). In traffic-based decomposition, the approximation of the departure process becomes the arrival process to the second queue. Therefore, we also illustrate the average queue length and its distribution in the second server for different levels n. In an effort to show how correlation propagates in the system, we also show the ACF of the departure process from the second server, i.e., at point "C" in Figure 2. All analytic results are obtained via MAMSolver, a matrix-analytic methods tool [19]. To assess the quality of the approximation, simulation results are also presented.

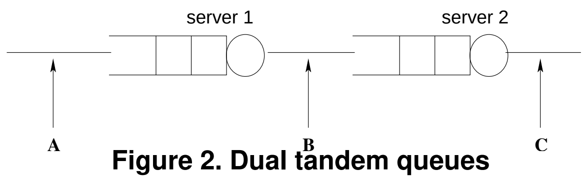

Figure 2. Dual tandem queues

4.1. Example MMPP(2)/M/1 → /Erlang-2/1

The first experiment uses the dual tandem queue MMPP(2) / M / 1 → / Erlang-2 / 1, parameterized as in [8]. In this experiment, the service process in the first queue is exponential, i.e., it has no correlation, while the arrival process does have correlation. In MAP notation, the arrival MMPP of order 2 is as follows:

$$\mathbf{D}_0^{(A)} = \begin{bmatrix} -6.9375 & 0.9375 \\ 0.0625 & -0.1958 \end{bmatrix}, \mathbf{D}_1^{(A)} = \begin{bmatrix} 6.0 & 0.0 \\ 0.0 & 0.1333 \end{bmatrix}. \tag{14}$$

This MMPP(2) has an arrival rate of 0.5 and a squared coefficient of variation (SCV) equal to 4.1. The mean service rate of server 1 is equal to 1, and the mean service rate of server 2 is equal to 1.25, resulting in a system that is moderately loaded. Analytic and simulation results of this network are given in Figure 3. The autocorrelation function (ACF) of the arrival process to the system (i.e., the ACF of the MMPP(2)) as a function of the lag is illustrated in Figure 3(a). Figure 3(b) plots the ACF of the departure process from server 1 (which is also the arrival process to server 2) for several truncation levels (as given by different values of parameter n). Simulation results are also reported in the figure. Although we did experiments for *several* values of n, we only plot representative ones here lest we overload the graph. As expected, the approximation with $n = 2$ is a poor one, as the ACF is not captured except for lag $k = 1$. The case with $n = 10$ captures *exactly* the lag correlations up to lag $k = 9$, and diverts beyond that point. Consistently, experiments with $n = 25$ and $n = 50$ capture the departure ACF exactly up to lag $k = 24$ and lag $k = 49$. The inset graph in Figure 3(b) provides a better look of how close the ACFs of various departure approximations match simulation for lags greater than 40. Note the difference in the scale of the graph – the y-axis is up to 0.0045, which practically means that for that large lags, the correlation is of almost no practical significance.[1] Note that queueing at the first node intensifies the ACF of the departure process in the tail, as compared with the ACF of the arrival process to server 1.

Figure 3(c) shows the average queue length in server 2 as a function of the truncation level and shows that in this particular case the approximation with $n = 10$ approaches simulation closely, while $n \geq 25$ gives nearly exact results (rel. error of 0.1%), i.e., the average queue length is equal to 0.95094, 0.95159 and 0.95160 for n=25, 50 and 100, respectively, and it is 0.95190 for simulation. The inset in the figure shows the queue length distribution. Up to queue length equal to 5 we use linear scale for the y-axis, and results for all values of n are nearly indistinguishable. Beyond 5, we use logarithmic scale for the y-axis and this allows us to better distinguish the tail of the queue. The graph shows that $n = 2$ is not as close to the simulation values, but $n = 25$ and $n = 50$ result in values indistinguishable to simulation.

Figure 3(d) shows the ACF of the departure process of server 2. We plot the simulation curve but also analytic curves with approximation parameter equal to $n = 10$ for server 1 and $n = 2, 3$ for server 2. The notation $n = x, n = y$ on the graph legend means that the approximation level for server 1 is equal to x and for server 2 equal to y. We also plot curves with approximation parameter $n = 50$ for server 1 and $n = 2, 3$ for server 2. The graph shows that if the first server is modeled with $n = 50$, analytic results on the second server even with small n are in excellent agreement with simulation.

We also note that at point "B" (see Figure 2) the marginal distribution (rate 0.5, SCV 2.67) is preserved for any approximation. While we conserve the flow also at point "C", the level-n approximation of the internal traffic at "B" also distorts the marginal distribution of the output approximation at the second server (compare the simulated SCV of 2.4561 to 2.4525 for $n = 10$ and 2.4560 for $n = 50$).

1 Note that the simulation graph for high lags is not smooth. This happens despite the fact that we simulate 100 million jobs.

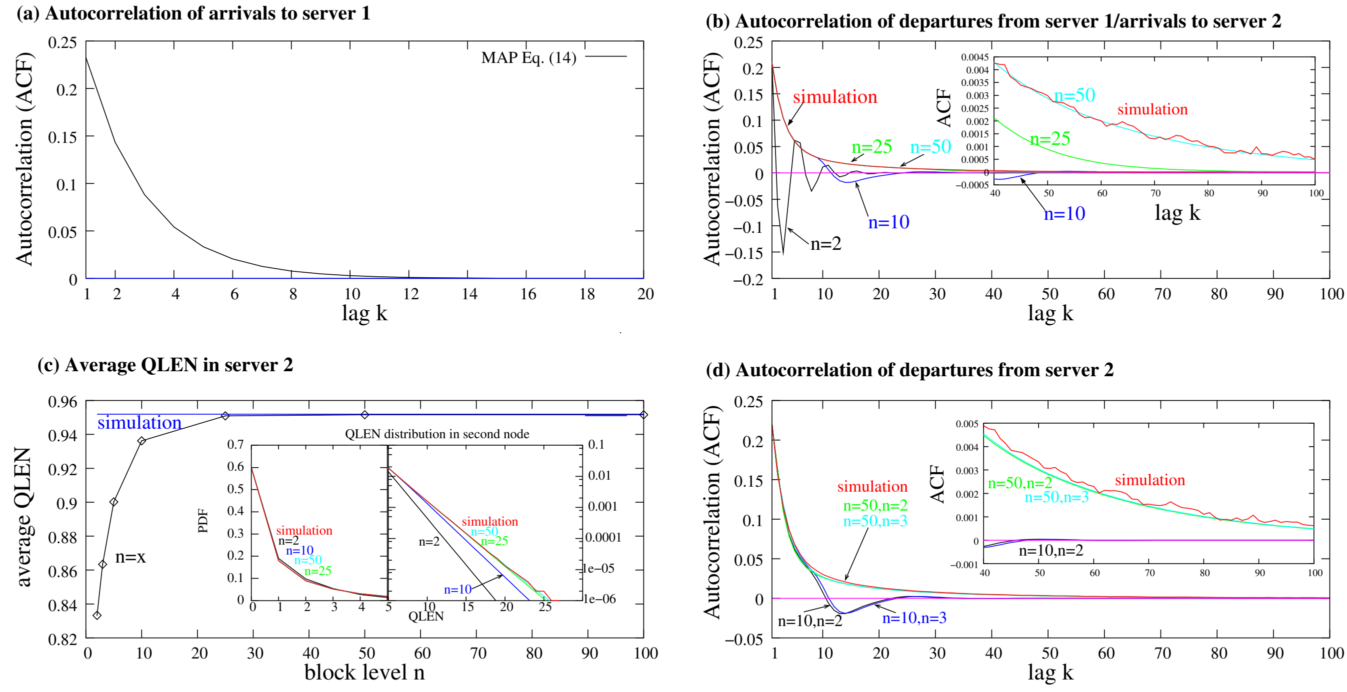

Figure 3. ACF of arrivals to server 1 (a), ACF of departures from server 1/arrivals to server 2 (b), mean queue length and queue length distribution at server 2 for different approximation levels (c), and ACF of departures from server 2 (d).

4.2. Example MAP(3)/H_2/1 $\rightarrow$ /Erlang-2/1

Here, we study another dual tandem queue of the scheme MAP(3) / Hyperexponential-2 / 1 $\rightarrow$ / Erlang-2 / 1. The arrival MAP(3) is taken from [3, page 229] (but scaled to a mean rate of 0.5):

$$\mathbf{D}_0^{(A)} = \begin{bmatrix} -0.542409519 & 0.0037279 & 0.0 \\ 0.004349217 & -0.02298872 & 0.000621317 \\ 0.0 & 0.001242633 & -2.269670072 \end{bmatrix} \tag{15}$$

$$\mathbf{D}_1^{(A)} = \begin{bmatrix} 0.020503453 & 0.0 & 0.518178166 \\ 0.0 & 0.017396869 & 0.000621317 \\ 2.259107688 & 0.004970534 & 0.004349217 \end{bmatrix} \tag{16}$$

This MAP has an SCV of about 18.86 and resulted from an EM-algorithm for fitting MAPs to data traces. The (uncorrelated) service times are given by two-stage hyperexponential (H_2) and two-stage Erlang distributions at the first and second server, respectively, with mean rates equal to 1. The rates of the H_2-distribution (with equal weights) are 10 and 0.52632, leading to an SCV equal to 2.62.

Figure 4(a) illustrates the autocorrelation of the arrival process to server 1. Note the jagged shape of the ACF. Figure 4(b) presents the ACF of the departure process from server 1 for different approximation levels n. Again, approximations with $n = x$ capture the lag coefficients up to $k = x - 1$. Surprisingly, the approximation with $n = 2$ appears to do much better than it did with the first example (see Figure 3(b)). Output approximation values of level n greater than 2 are (almost) indistinguishable from simulation, as there is severe overplotting, but can be distinguished in the inset. Note that the y-axis range goes up to 0.0035 and that approximation levels as high as 25 and 50 capture the jagged shape of the ACF. Even better than in Figure 3(b), the rather unpredictable behavior of the ACF of departure approximations is observed when k reaches n (especially for $n = 25$).

The average queue length is illustrated in Figure 4(c) and shows that levels greater or equal to 25 are needed to capture the average queue length (while Figure 4(b) suggested that smaller n might be sufficient for a reasonable approximation). The inset graph shows how the queue length distribution is underestimated in its tail – due to truncation, which tends to underestimate the positive correlation of the departure process. Note that the output approximations preserve the marginal distributions (rate 0.5, SCV 18.15). Stronger autocorrelations, higher variability of arrival process and service times and a lower service rate at the second node all contribute to a larger average queue length than in the first experiment.

Figure 4(d) shows the ACF of the departure process of

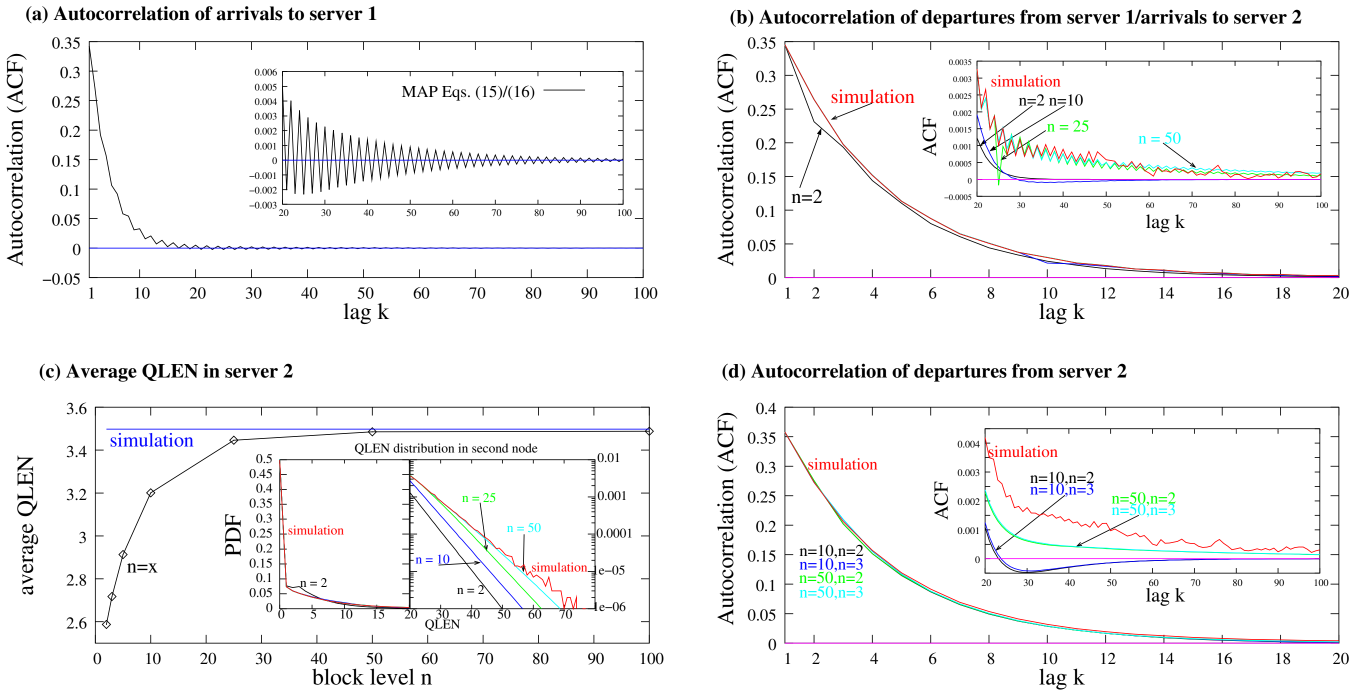

Figure 4. ACF of arrivals to server 1 (a), ACF of departures from server 1/arrivals to server 2 (b), mean queue length and queue length distribution at server 2 for different approximation levels (c), and ACF of departures from server 2 (d).

server 2. As in the first example, we experimented with different combinations of approximation levels at both nodes. The inset graph of Figure 4(d) shows that to closely capture correlation, higher approximation levels need to be chosen for the first truncation than for the second one. Note that the ACF values of the y-axis are *very* small, indicating that levels as low as 2 and 3 for the departure approximation of server 2 are still effective.

As opposed to the first example, the simulated SCV (equal to 17.04) of the departure process of the second queue (point "C") is now slightly overestimated by a level-n approximation ($n = 10$: SCV 17.11, $n = 50$: SCV 17.09), where n refers to the truncation level for the first queue.

4.3. Example MAP(3)/MAP(2)/1 → /Erlang-2/1

To evaluate the importance of correlation in the service process, we use the same scheme as before, but modify the H_2-service in server 1 so that this marginal distribution is preserved, but the consecutive service times become correlated. This is achieved by the following MAP, which describes alternating exponential service times:

$$\mathbf{D}_0^{(S_1)} = \begin{bmatrix} -10 & 0 \\ 0 & -0.52632 \end{bmatrix}, \mathbf{D}_1^{(S_1)} = \begin{bmatrix} 0 & 10 \\ 0.52632 & 0 \end{bmatrix} \quad (17)$$

Figure 5(a) depicts the ACF of the service process in server 1. Note that this is a case of strong correlation that oscillates between -0.3 and 0.3. Figure 5(b) shows that correlation in the service is also reflected in the ACF of the output process of the first node (compare the jags in the simulation curve in Figure 5(b) with those in Figure 4(b)). The inset of Figure 5(b) illustrates how this jagged behavior is reached by the approximation, as the level n increases. Due to the nature of these approximations, adding a block level does not necessarily improve the approximation of the mean queue length, as shown in Figure 5(c). Finally, Figure 5(d) gives the ACF of the departure process from server 2 and illustrates that the Erlang-2 distribution in server 2 seemingly takes out the intense jagged behavior of the arrival process to this server. Again, simulation results are given together with analytic results for different combinations of approximation levels for the departure processes at both servers. The inset shows that the jagged behavior of the ACF is still there, albeit significantly less prominent. To capture this, approximation levels larger than 3 are needed for server 2.

Due to the correlated service process, the SCVs decay more slowly than in the second example, namely to 18.3 at point "B" and 17.3 at point "C" (according to simulation). Again, traffic-based decomposition (slightly) overestimates the SCV of the departure process at the second

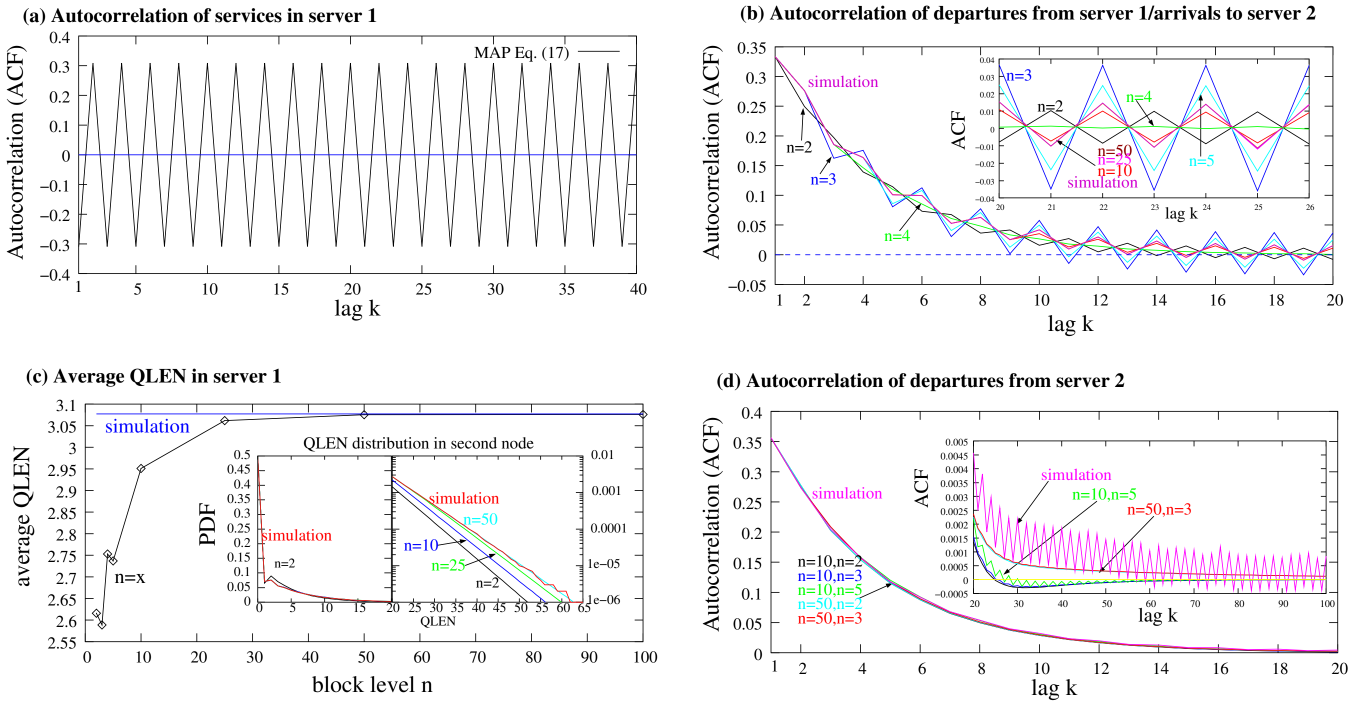

Figure 5. ACF of service process at server 1 (a), ACF of departures from server 1/arrivals to server 2 (b), mean queue length and queue length distribution at server 2 for different approximation levels (c), and ACF of departures from server 2 (d).

queue (17.316 for $n = 10$ and 17.305 for $n = 50$, where n is the truncation level for the internal traffic).

5. Complexity comparisons

As pointed out in Section 1.1, our approach compares favorably with other truncation methods, because it provides both generality and efficiency at the same time.

Here, we briefly comment on complexity issues in comparison with the approach in [21] for MAP/MAP/1 queues (this discussion also applies to the less general approach in [5]). Recall that $m = m_A m_S$ is the dimension of the QBD block levels. Our approach requires matrix $\mathbf{G}$ of dimension m, which is often sparse and can be efficiently computed by matrix-analytic techniques [10, 13]. The approaches in [21, 5] need to compute the matrix $\mathbf{R}$ (dimension m), which is at least as expensive as the computation of $\mathbf{G}$ [10].

The method of [21] additionally computes the stationary queue length distribution (or at least parts of it). When we desire only the output model (with $n+1$ levels, $n \geq 1$) of a MAP/MAP/1 queue and no other quantity, the *additional* (per-queue) effort for [21] consists in

1. solving a linear system of equations of size $2m$ (see (7), or equivalently a system of size m, where an additional matrix inversion for a matrix of order m is required),
2. $n-1$ vector-matrix multiplications of order m (related to (6)), and
3. a matrix inversion for a matrix of order m (or equivalently solving another linear system of equations of the same dimension, related to computing a vector like $\boldsymbol{\pi}^{(n,*)} = \boldsymbol{\pi}^{(n)}(\mathbf{I} - \mathbf{R})^{-1}$).

Here, we neglected some scalar-matrix multiplications and matrix subtractions. In iterative traffic-based decomposition for queueing networks with loops, frequently only the output model is required. When one is also interested in performance measures in terms of moments of the queue length distribution or waiting time – besides the proper output model – the computational cost of point 1 must be left out of the additional effort. If moreover, the queue length distribution is desired, the per-queue difference in computational cost between the approach in [21] and the one presented in this paper shrinks to point 3 (plus/minus some scalar-matrix multiplications and matrix subtractions). The savings with our truncation models are still significant with respect to the overall cost of $O(m^3)$, especially when a large number of departure processes needs to be approximated.

6. Concluding remarks

We have exploited finite aggregate representations of the QBD process, which arises from a MAP/MAP/1 queue, to derive a family of tractable departure process approximations of this queueing system. Due to its invariance properties, this truncation model lends itself for studying characteristics of the true departure process or for application in traffic-based decomposition of queueing networks – as demonstrated in numerical experiments. Future research will scrutinize the interdependence of truncation levels and correlation propagation and investigate the impact of various parameters, like utilization.

Other truncation methods also preserve the marginal distribution of the true departure process and the first $n-1$ lag coefficients of correlation (for an output approximation in $n+1$ QBD levels, $n \geq 1$), but our approach turns out to be more efficient and/or applies to a wider class of queueing systems (e.g., including correlated services). In fact, since our approximation (and all related proofs) rely on the QBD representation only, it works – with the same properties – for any queueing system with such a structure, e.g., a Σ_i MAP$_i$/MAP/m queue. As a shared drawback of all the discussed truncation techniques, we point out that the dimension of the output model is a multiple of the order product for arrival and service processes. In traffic-based decomposition of queueing networks, the repeated application of this truncation may soon lead to a state-space explosion, when not combined with other (available) fixed-size output models (e.g., [8, 12]).

References

[1] C. Amza, E. Cecchet, A. Chanda, A. Cox, S. Elnikety, R. Gil, J. Marguerite, K. Rajamani, and W. Zwaenepoel. Specification and implementation of dynamic content benchmarks. In *Proc. 5th IEEE Workshop on Workload Characterization (WWC-5)*, pages 3–13, 2002.

[2] N. G. Bean, D. A. Green, and P. G. Taylor. Approximations to the output process of MAP/PH/1/ queues. In *Proc. 2nd Int. Workshop on Matrix-Analytic Methods in Stochastic Models*, pages 151–159. Notable Publications, 1998.

[3] P. Buchholz. An EM-algorithm for MAP fitting from real traffic data. In P. Kemper and W. Sanders, editors, *Proc. 13th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, volume 2794 of *LNCS*, pages 218–236, 2003.

[4] A. Erramilli, O. Narayan, and W. Willinger. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Trans. Netw.*, 4(2):209–223, 1996.

[5] D. Green. *Departure Processes from MAP/PH/1 Queues*. PhD thesis, Department of Applied Mathematics, University of Adelaide, 1999.

[6] D. Green. Lag correlations of approximating departure processes of MAP/PH/1 queues. In *Proc. 3rd Int. Conf. on Matrix-Analytic Methods in Stochastic Models*, pages 135–151. Notable Publications, 2000.

[7] A. Heindl. *Traffic-Based Decomposition of General Queueing Networks with Correlated Input Processes*. Shaker Verlag, Aachen, Germany, 2001. PhD Thesis, TU Berlin.

[8] A. Heindl, K. Mitchell, and A. van de Liefvoort. The correlation region of second-order MAPs with application to queueing network decomposition. In P. Kemper and W. Sanders, editors, *Proc. 13th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, volume 2794 of *LNCS*, pages 237–254, 2003.

[9] J. Kumaran, K. Mitchell, and A. van de Liefvoort. Characterization of the departure process from an ME/ME/1 queue. *RAIRO Recherche Operationelle / Operations Research*, 2004. accepted for publication.

[10] G. Latouche and V. Ramaswami. *Introduction to Matrix-Analytic Methods in Stochastic Modeling*. Series on statistics and applied probability. ASA-SIAM, 1999.

[11] M. Livny, B. Melamed, and A. K. Tsiolis. The impact of autocorrelation on queueing systems. *Management Science*, 39:322–339, 1993.

[12] K. Mitchell and A. van de Liefvoort. Approximation models of feed-forward G/G/1/N queueing networks with correlated arrivals. *Performance Evaluation*, 51:137–152, 2003.

[13] V. A. Naoumov, U. Krieger, and D. Wagner. Analysis of a multi-server delay-loss system with a general Markovian arrival process. In Chakravarthy and Alfa, editors, *Matrix-Analytic Methods in Stochastic Models*, pages 43–66, New York, 1997. Marcel Dekker.

[14] M. Neuts. *Algorithmic Probability: A Collection of Problems*. Chapman and Hall, 1995.

[15] M. F. Neuts. *Matrix-geometric Solutions in Stochastic Models*. Johns Hopkins University Press, Baltimore, MD, 1981.

[16] M. F. Neuts. *Structured Stochastic Matrices of M/G/1-type and their Applications*. Marcel Dekker, New York, NY, 1989.

[17] B. Patuwo, R. Disney, and D. McNickle. The effect of correlated arrivals on queues. *IIE Transactions*, 25:105–110, 1993.

[18] A. Riska and E. Smirni. Exact aggregate solutions for M/G/1-type Markov processes. In *Proc. Int. Conf. on Measurement and Modeling of Computer Systems (ACM SIGMETRICS 2002)*, pages 86–96. ACM Press, 2002.

[19] A. Riska and E. Smirni. MAMSolver: a matrix-analytic methods tools. In T. Field, P. Harrison, J. Bradley, and U. Harder, editors, *Proc. 12th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, volume 2324 of *LNCS*, pages 205–211, 2002.

[20] A. Riska and E. Smirni. M/G/1-type Markov processes: A tutorial. In *Tutorials of the IFIP WG7.3 Int. Symposium on Computer Performance Modeling, Measurement and Evaluation*, volume 2459 of *LNCS*, pages 36–63, 2002.

[21] R. Sadre and B. Haverkort. Characterizing traffic streams in networks of MAP/MAP/1 queues. In *Proc. 11th GI/ITG Conf. on Measuring, Modelling and Evaluation of Computer and Communication Systems*, pages 195–208, Aachen, Germany, 2001.