SOME ERGODICITY AND TRUNCATION BOUNDS FOR A SMALL SCALE MARKOVIAN SUPERCOMPUTER MODEL

Rostislav Razumchik
Federal Research Center "Computer Science and Control"
of the Russian Academy of Sciences
44-2 Vavilova Str., Moscow 119333, Russian Federation,
Peoples Friendship University of Russia (RUDN University),
6 Miklukho-Maklaya Str., Moscow 117198, Russian Federation
Email: rrazumchik@ipiran.ru, razumchik-rv@rudn.ru

Alexander Rumyantsev
Institute of Applied Mathematical Research,
Karelian Research Centre of RAS
11 Pushkinskaya Str.
Petrozavodsk 185910, Russian Federation
Email: ar0@krc.karelia.ru

KEYWORDS

supercomputer model, multiserver job model, transient analysis, speed of convergence to steady state

ABSTRACT

In this paper we address the transient analysis of a markovian two-server supercomputer model where customers are served by a random number of servers simultaneously. The Markov process, which described the model's evolution, is of quasi–birth–death type. It is shown that, at least under low load conditions, the logarithmic norm method can be used to obtain ergodicity bounds for the model. This allows one to solve both the stability detection problem (i.e. determine when the computations of the time–dependent performance measures can be terminated) and the truncation problem (i.e. locate the level at which the infinite system of Kolmogorov forward equations must be truncated in order to guarantee certain accuracy). An illustrative numerical example is provided.

INTRODUCTION

Most of the standard computing systems are now parallel. Ranging from multicore battery-powered devices to large-scale datacenters and supercomputers, all these are capable of processing compute load in parallel way. As a response to this trend, multiserver queueing models are actively studied in recent decades.

From a software perspective, parallel computing technologies are used to empower the software engineers with relevant hardware capabilities. An essential feature present both in a supercomputer and on, say, laptop is the possibility to run a specific code simultaneously on a number of cores/servers which results in overall computation time reduction. However, from the modeling perspective, such a model, called simultaneous service multiserver system (a.k.a. multiserver job model, cluster model, hereinafter referred as *supercomputer model*) is known to be hard to analyze (Harchol-Balter, 2021).

The distinctive feature of the supercomputer model is simultaneous occupation and simultaneous release of a (random) number of servers by a customer in a rigid way (Filippopoulos and Karatza, 2007) (in contrast to classical single-server customers), causing the workload process to be non-work-conserving (Rumyantsev and Morozov, 2017). While supercomputer model is used to analyze the stability and performance characteristics of supercomputers (Morozov and Rumyantsev, 2011; Rumyantsev and Morozov, 2017), these models are also applicable to the study of social service systems (Brill and Green, 1984; Kim, 1979).

Stability (Rumyantsev and Morozov, 2017; Rumyantsev, 2020) and performance characteristics of supercomputer model are difficult to obtain even for small scale instances (Filippopoulos and Karatza, 2007; Chakravarthy and Karatza, 2013), and therefore a significant number of problems associated with such systems are open (Harchol-Balter, 2021, 2022). Surprisingly, among the most interesting problems related to the supercomputer model recently announced in (Harchol-Balter, 2022), performance analysis in transient regime is not enumerated. Yet time-dependent characteristics are rather important for energy consumption analysis, which goes in the context of energy efficiency studies focused on supercomputer model's energy-performance tradeoff, see e.g. (Rumyantsev et al., 2021). The present paper is a step in this important direction.

Here consideration is given to time–dependent performance characteristics (e.g. average number of customers at the time instant t) of the *small scale* supercomputer model in *markovian* case. While in general this problem can be considered from the aspects of computation time, accuracy, complexity, storage etc., we establish *upper bounds* on the transient performance under light load, which constitutes the main contribution of the paper. Below we give the motivation for the importance of this result.

In transient regime, basic performance characteristics can be obtained as the solution the infinite system (see (7)) of ordinary differential equations (ODE). As such, solution techniques are closure approximations (Taaffe and Ong, 1987; Clark, 1981; Massey and Pender, 2013), uniformization (Van Dijk et al., 2018) as well as various differential equation solvers (Arns et al., 2010), to name a few. However, to improve the effi-

Communications of the ECMS, Volume 36, Issue 1,
Proceedings, ©ECMS Ibrahim A. Hameed, Agus Hasan,
Saleh Abdel-Afou Alaliyat (Editors) 2022
ISBN: 978-3-937436-77-7/978-3-937436-76-0(CD) ISSN 2522-2414

ciency of any solution technique, two questions need to be addressed. The first question concerns *stability detection*: time-dependent performance computation (from ODE) can be terminated once the model has reached stability regime. The second question deals with the *truncation*: the infinite system of ODEs needs to be truncated at a finite level before numerical technique can be applied. Both questions are addressed in the present paper using the well–known *logarithmic norm* method, see e.g. (Zeifman et al., 2021, Section 2). Compared to the previous studies, here we give one more evidence that the method can be applied to the analysis of the so-called quasi–birth–death (QBD) processes (with finite number of phases). Even though the particular case considered in the present paper can be treated using other methods from the literature (see, for example, (Burak and Korytkowski, 2020)), comparison of the methods, not being the goal of the paper, was not undertaken.

For the sake of brevity, further attention is paid only to the case of homogeneous servers. Although the generalization to the non-homogeneous case is possible (by following the lines drawn in the present paper), the generalization to n-server case is not straightforward and will be considered elsewhere.

The paper is organized as follows. In the next section, the detailed problem statement is given. Section 3 reviews the necessary theory, which is used in the Sections 4-6 to obtain the solutions. Some results of the numerical experiments and the main conclusions of the research are briefly summarized in Sections 7-8.

Notation

In what follows by $\|\cdot\|$ we denote the l_1-norm, i.e. if $\vec{x}$ is an $(l+1)$-dimensional column vector, then $\|\vec{x}\| = \sum_{k=0}^{l} |x_k|$. The choice of operator norms will be the one induced by the l_1-norm on column vectors i.e. $\|\mathbb{H}\| = \sup_{0\le j\le l} \sum_{i=0}^{l} |h_{ij}|$ for a linear operator (matrix) $\mathbb{H}$. The vectors throughout the paper are regarded as column vectors (dimensions are clear from the context), $\vec{1}^{\mathrm{T}}$ — row vector of 1's with T denoting the matrix transpose, $\mathbb{I}$ — identity matrix.

THE MODEL AND THE PROBLEM STATEMENT

The schematic representation of the model considered in this paper is best given by the Fig. 1 in (Filippopoulos and Karatza, 2007). For the sake of completeness we reproduce it below.

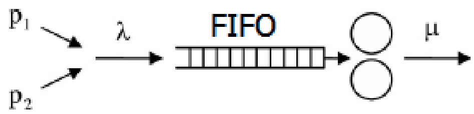

Figure 1: Two identical servers process customers from the infinite–capacity queue on the FIFO basis. For each customer i servers (with the probability p_i) are required to start service

The system consists of two identical servers and a queue of infinite capacity which follows the FIFO scheduling discipline. Customers of two classes arrive to the system according to the Poisson flow at joint rate $\lambda = \lambda p_1 + \lambda p_2$, where p_i is the probability of class-i customer arrival, $i = 1, 2$, and $p_1 + p_2 = 1$. Class-i customer requires i servers simultaneously for the same service time exponentially distributed with the rate μ (independent of customer class). In case there are insufficient resources (servers) available, the head-of-queue customer prevents subsequent ones (if any) from entering service. We assume the most interesting case $p_1 \in (0, 1)$, since at the endpoints the model degenerates either to to the classic $M/M/1$, or to classic $M/M/2$ queue already studied using the logarithmic norm method, see (Zeifman et al., 2019b). Thus we intentionally do not consider the boundary cases $p_1 = 0$ and $p_1 = 1$.

It is assumed that the customer class becomes known upon customer's arrival and remains unchanged during customer's sojourn time in the system. However, in the model only two oldest (in the order of arrival) customer classes are tracked. This is correct, since other customers (if any) may have generic class which indeed becomes known only upon arrival to the *head of the queue*, for more discussion of this issue see (Rumyantsev and Morozov, 2017). As such, we adopt from (Rumyantsev and Morozov, 2017) the following three–dimensional continuous-time Markov process describing the system,

$$\{S(t) = (N(t), X_1(t), X_2(t)), t \ge 0\}, \tag{1}$$

where $N(t) \ge 0$ is the number of customers in the system, and $X_i(t)$ is the class of ith oldest customer in the system at time $t \ge 0$, if any. The state space $\mathcal{X}$ of the model consists of subsets (levels)

$$\mathcal{X} = \{0\} \cup \mathcal{X}_1 \cup \mathcal{X}_2 \ldots,$$

where $\{0\}$ denotes an empty system, and the set $\mathcal{X}_n$, $n \ge 1$ corresponds to all possible states with n customers in the system. Thus $\mathcal{X}_1$ consists of two states $\{(1, 1), (1, 2)\}$ (we do not include the empty component); $\mathcal{X}_n$ consists of four triplets (n, x_1, x_2), where n is the total number of customers in the system, and $x_1, x_2 \in \{1, 2\}$ are the classes of oldest, second oldest customers, respectively.

As such, the process (1) is the irreducible QBD process with infinitesimal generator $\mathbb{Q}$ of the form

$$\mathbb{Q} = \begin{pmatrix} \mathbb{N}_0 & \mathbb{L}_0 & 0 & 0 & \ldots \\ \mathbb{M}_0 & \mathbb{N} & \mathbb{L} & 0 & \ldots \\ 0 & \mathbb{M} & \mathbb{N} & \mathbb{L} & \ldots \\ 0 & 0 & \mathbb{M} & \mathbb{N} & \ldots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

where the blocks are defined as

$$\mathbb{N}_0 = \begin{pmatrix} -\lambda & \lambda p_1 & \lambda p_2 \\ \mu & -\lambda-\mu & 0 \\ \mu & 0 & -\lambda-\mu \end{pmatrix}, \ \mathbb{M}_0 = \begin{pmatrix} 0 & 2\mu & 0 \\ 0 & 0 & \mu \\ 0 & \mu & 0 \\ 0 & 0 & \mu \end{pmatrix},$$

$$\mathbb{L}_0 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ \lambda p_1 & \lambda p_2 & 0 & 0 \\ 0 & 0 & \lambda p_1 & \lambda p_2 \end{pmatrix}, \ \mathbb{L} = \lambda \mathbb{I},$$

$$\mathbb{N} = \begin{pmatrix} -2\mu-\lambda & 0 & 0 & 0 \\ 0 & -\mu-\lambda & 0 & 0 \\ 0 & 0 & -\mu-\lambda & 0 \\ 0 & 0 & 0 & -\mu-\lambda \end{pmatrix},$$

$$\mathbb{M} = \begin{pmatrix} 2\mu p_1 & 2\mu p_2 & 0 & 0 \\ 0 & 0 & \mu p_1 & \mu p_2 \\ \mu p_1 & \mu p_2 & 0 & 0 \\ 0 & 0 & \mu p_1 & \mu p_2 \end{pmatrix}.$$

Denote the time–dependent probability distribution of the Markov process (1) by $p_{n;i,j}(t)$ i.e.

$$\begin{aligned} p_{n;i,j}(t) &= \mathrm{P}\{N(t) = n, X_1(t) = i, X_2(t) = j\}, \ n \geq 2, \\ p_{1,i}(t) &= \mathrm{P}\{N(t) = 1, X_1(t) = i\}, \\ p_0(t) &= \mathrm{P}\{N(t) = 0\}. \end{aligned}$$

Let

$$\vec{p}_n(t)^{\mathrm{T}} = (p_{n;1,1}(t), p_{n;1,2}(t), p_{n;2,1}(t), p_{n;2,2}(t)), \ n \geq 2,$$

$$\vec{p}_1(t)^{\mathrm{T}} = (p_{1,1}(t), p_{1,2}(t)),$$

and

$$\vec{p}(t)^{\mathrm{T}} = \left(p_0(t), \vec{p}_1(t)^{\mathrm{T}}, \vec{p}_2(t)^{\mathrm{T}}, \dots\right).$$

It is known (Brill and Green, 1984), that when the inequality $\lambda < 2\mu/(2 - p_1^2)$ holds, the model is stable and thus $\vec{p}(t) \to \vec{p}$ (element-wise) as $t \to \infty$, where the vector $\vec{p}$ can be found from the system of global balance equations $\mathbb{Q}^{\mathrm{T}}\vec{p} = \vec{0}$, $\vec{p}^{\mathrm{T}}\vec{1} = 1$. The exact procedures to compute the entries of $\vec{p}$ are already available in the literature (see, for example, (Filippopoulos and Karatza, 2007)). In what follows we are basically interested in the two questions:
(i) given $\epsilon > 0$, find t^* such that

$$\left\|\vec{p}(t) - \vec{p}\right\| < \epsilon \text{ for } t > t^*; \tag{2}$$

(ii) given $\epsilon > 0$, find a positive integer N^* such that $\left\|\vec{p}(t) - \vec{p}^*(t)\right\| < \epsilon$, where $\vec{p}^*(t)$ denotes the time-dependent probability distribution vector of the supercomputer model with the queue of finite capacity N^* (this results in changes for the corresponding generator matrix $\mathbb{Q}^*$ given explicitly in (18)).

AUXILIARY RESULTS

In order to construct the upper bounds for $\left\|\vec{p}(t) - \vec{p}\right\|$ (and, as will be seen, for $\left\|\vec{p}(t) - \vec{p}^*(t)\right\|$ as well) we will use the notion of the logarithmic norm of locally integrable operator functions and (known) estimates for the differential equations. Consider an ODE system[1]

$$\frac{d}{dt}\vec{y}(t) = \mathbb{H}(t)\vec{y}(t), \ t \geq 0, \tag{3}$$

where the entries of the matrix $\mathbb{H}(t) = (h_{ij}(t))_{i,j=1}^{\infty}$ are locally integrable on $[0, \infty)$ and $\mathbb{H}(t)$ is bounded in the sense that $||\mathbb{H}(t)||$ is finite for any fixed t. Then

$$\frac{d}{dt}||\vec{y}(t)|| \leq \gamma\left(\mathbb{H}(t)\right)||\vec{y}(t)||, \tag{4}$$

where

$$\gamma(\mathbb{H}(t)) = \sup_i \left\{ h_{ii}(t) + \sum_{j \neq i} |h_{ji}(t)| \right\}. \tag{5}$$

is called the logarithmic norm of $\mathbb{H}(t)$. Thus from (4) one gets the following upper bound[2]:

$$||\vec{y}(t)|| \leq e^{\int_0^t \gamma(\mathbb{H}(u))\,du}||\vec{y}(0)||. \tag{6}$$

STABILITY DETECTION

Given any proper initial condition $\vec{p}(0)$, the Kolmogorov forward equations for the time-dependent distribution $\vec{p}(t)$ of (1) can be written as

$$\frac{d}{dt}\vec{p}(t) = \mathbb{Q}^{\mathrm{T}}\vec{p}(t), \ t \geq 0, \tag{7}$$

with the normalization condition

$$\vec{p}(t)^{\mathrm{T}}\vec{1} = 1. \tag{8}$$

It is straightforward to check that the logarithmic norm $\gamma(\mathbb{Q}^{\mathrm{T}})$ is always positive. Thus the right part of (6) grows with t and is useless in solving both (i) and (ii). As such, a specific transformation is needed which we describe below.

Fix a positive constant, say c, and a non-decreasing sequence of positive numbers $\{d_i, i \geq 1\}$ with $d_1 = 1$. Let $D_i = \prod_{n=1}^{i} d_n$. Introduce two infinite matrices, say $\mathbb{D}$ and $\mathbb{C}$, having the form:

$$\mathbb{D} = diag(D_1, D_2, D_2, \underbrace{D_3, D_3, D_3, D_3}_{4\ entries}, \dots, \underbrace{D_n, D_n, D_n, D_n}_{4\ entries}, \dots),$$

[1]The definitions and results, which are stated without any details below, can be fully recovered from, for example, (Zeifman, 1995, Appendix).

[2]It is worth mentioning, that for the bound (6) to hold, it is not necessary for $H(t)$ to be bounded to all $t \geq 0$. In such a case the right part of (5) is the generalization of the logarithmic norm (see (Zeifman et al., 2019a)).

$$\mathbb{C} = \begin{pmatrix} c & c & c & c & \dots \\ 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Now consider two stochastic vectors, $\vec{p}^*(t)$ and $\vec{p}^{**}(t)$, solving (7). It follows from (8) that the solution of (7) is not affected by performing row operations on $\mathbb{Q}^{\mathrm{T}}$ and, in particular, adding or subtracting some constant (componentwise) from a single row in the matrix $\mathbb{Q}^{\mathrm{T}}$. As such, we get the following easily verifiable identity:

$$\frac{d}{dt}\mathbb{D}(\vec{p}^*(t) - \vec{p}^{**}(t)) = \underbrace{\mathbb{D}\left(\mathbb{Q}^{\mathrm{T}} - \mathbb{C}\right)\mathbb{D}^{-1}}_{=\mathbb{A}}\mathbb{D}(\vec{p}^*(t) - \vec{p}^{**}(t)).$$

Assume now that such $\mathbb{D}$ (i.e. the sequence $\{d_i, i \geq 1\}$) and $\mathbb{C}$ exist that the logarithmic norm of the matrix $\mathbb{A}$ is negative[3], i.e.

$$\gamma(\mathbb{A}) < 0. \tag{9}$$

Then, since $d_i \geq 1$, it follows from (6) that

$$\left\|\vec{p}^*(t) - \vec{p}^{**}(t)\right\| \leq e^{\gamma(\mathbb{A})t}\left\|\mathbb{D}(\vec{p}^*(0) - \vec{p}^{**}(0))\right\|. \tag{10}$$

By plugging into the previous inequality $\vec{p}$ instead of $\vec{p}^{**}(t)$ and solving (2) for t, one gets the solution of (i):

$$t^* = \frac{1}{\gamma(\mathbb{A})}\ln\left(\frac{\epsilon}{\left\|\mathbb{D}(\vec{p}(0) - \vec{p})\right\|}\right). \tag{11}$$

It remains to establish the conditions of existence of $\mathbb{D}$ and $\mathbb{C}$ that guarantee $\gamma(\mathbb{A}) < 0$, which is done in the next section.

Note that the condition (9) allows one to obtain useful insights into the model. For example, assuming that $\inf_{i\geq 1}(i^{-1}D_{i+1}) > 0$, for the average number $\mathsf{E}N(t)$ of customers at instant t we have:

$$\mathsf{E}N(t) = \sum_{n=1}^{\infty} n\vec{p}_n(t)^{\mathrm{T}}\vec{1} \leq \left[\inf_{i\geq 1}\frac{D_{i+1}}{i}\right]^{-1}\left\|\mathbb{D}\vec{p}(t)\right\|. \tag{12}$$

By left-multiplying (7) with $\mathbb{D}$, one obtains the upper bound for $\left\|\mathbb{D}\vec{p}(t)\right\|$:

$$\left\|\mathbb{D}\vec{p}(t)\right\| \leq e^{\gamma(\mathbb{A})t}\left\|\mathbb{D}\vec{p}(0)\right\| + \frac{c}{\gamma(\mathbb{A})}\left(e^{\gamma(\mathbb{A})t} - 1\right). \tag{13}$$

Once the initial condition $\vec{p}(0)$ is fixed, by plugging (13) into (12) one immediately gets the upper bounds for the average number of customers in the model at instant t and for its steady state value. Since the Little's law holds for the considered model, one has at once also the upper bound for the average response time.

[3]Note, that even when $\gamma(\mathbb{A}) < 0$, the matrix $\mathbb{A}$ may have negative row elements off the main diagonal.

LOGARITHMIC NORM

Let us proceed establish the conditions for (9) to hold good. Put $c = \mu$ and $d_i = d > 1$ for all $i \geq 2$. Then the right part of (5) (after plugging $\mathbb{A}$ instead of $\mathbb{H}(t)$) becomes equal to

$$\lambda(d-1) - \mu + \mu\frac{d+1}{d^2} = \frac{f(d)}{d^2}. \tag{14}$$

It can be shown, that if $f(d) < 0$, then the range of possible value of the customer's arrival rate λ is limited to the interval

$$\lambda \in \left(0, \mu \cdot \min\left(\frac{1+\sqrt{5}}{2}, \frac{2}{2-p_1^2}\right)\right). \tag{15}$$

But even when (15) holds, it may happen[4] that $f(d) > 0$. It is well-known that $f(d) = 0$ has three distinct real roots if its discriminant, say Δ, is positive. Assume[5] $\Delta > 0$. Descartes rule of signs shows that two of the three roots of $f(d) = 0$ are positive; denote them δ_1 and δ_2 ($\delta_1 < \delta_2$). Since $f(0) > 0$ and $f(\infty) > 0$, then $f(d) < 0$ always for $d \in (\delta_1, \delta_2)$. Since $f(1) > 0$ and $f'(1) < 0$, then $(\delta_1, \delta_2) \subset (1, \infty)$. Moreover, Sturm's theorem applied to the polynomial $f(d)$ shows that $(\delta_1, \delta_2) \subset \left(1, 1 + \frac{\mu}{\lambda}\right)$. Now, in order to find the proper value of $\gamma(\mathbb{A})$, it remains to choose $d \in (\delta_1, \delta_2)$, which minimizes (14) i.e.

$$\gamma(\mathbb{A}) = \min_{d\in(\delta_1,\delta_2)}\left\{\lambda(d-1) - \mu + \mu\frac{d+1}{d^2}\right\}. \tag{16}$$

For any $p_1 \in (0,1)$ it can be checked numerically, that the assumption $\Delta > 0$ holds whenever

$$\lambda \in \left(0, \alpha(p_1)\cdot\mu \cdot \min\left(\frac{1+\sqrt{5}}{2}, \frac{2}{2-p_1^2}\right)\right), \tag{17}$$

where α is the curve with the box markers in the Fig. 2 (upper curve).

The approximation for the function α allows one to estimate the fraction of the stability interval, in which it is possible[6] to make the logarithmic norm $\gamma(\mathbb{A})$ negative: it varies from $\approx 29\%$ (when $p_1 \approx 0$) to $\approx 15\%$ (when $p_1 \approx 1$).

TRUNCATION BOUNDS

Fix a positive integer N^* and assume that the total number of customers in the model can never be greater than N^* i.e. those customers which find the queue full are considered as lost. It is convenient to describe the new model dynamics by the QBD, say $\{S^*(t), t \geq 0\}$, which is identical to (1), except for the fact that the new generator, say

[4]Because the cubic equation $f(d) = 0$ (in one variable d) does not have three real roots.

[5]If $\Delta = 0$ all roots of $f(d) = 0$ are real, with one root repeated. Due to the Descartes rule of signs the latter is positive. Thus, since $f(0) > 0$ and $f(\infty) > 0$, $f(d)$ can never be negative for $d \in (0, \infty)$.

[6]When $c = \mu$ and $d_i = d > 1$ for all $i \geq 2$.

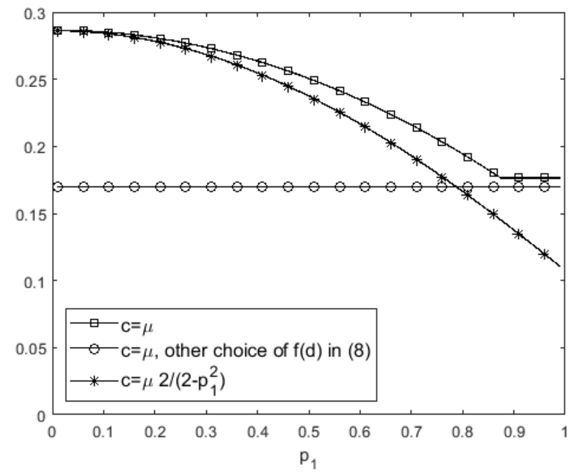

Figure 2: Dependencies of peak allowed values of α on the value of p_1.

$\mathbb{Q}^*$, has the form

$$\mathbb{Q}^* = \begin{pmatrix} \mathbb{N}_0 & \mathbb{L}_0 & 0 & 0 & 0 & \dots \\ \mathbb{M}_0 & \mathbb{N} & \mathbb{L} & 0 & 0 & \dots \\ 0 & \mathbb{M} & \mathbb{N} & \mathbb{L} & 0 & \dots \\ 0 & 0 & \mathbb{M} & -diag(\mathbb{M}\vec{1}) & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (18)$$

Denote the time–dependent probability vector of the new model states with $\vec{p}^*(t)$; by analogy with $\vec{p}(t)$, we define it as

$$\vec{p}^*(t)^{\mathrm{T}} = \left(p_0^*(t), \vec{p}_1^*(t)^{\mathrm{T}}, \vec{p}_2^*(t)^{\mathrm{T}}, \dots \vec{p}_{N^*}^*(t)^{\mathrm{T}}, 0, \dots\right).$$

Let us assume that the initial state probability distributions of both QBDs $\{S(t), t \geq 0\}$ and $\{S^*(t), t \geq 0\}$ are identical. Then, by applying the well-known truncation technique for birth–and–death processes (see, for example, (Zeifman et al., 2014; Satin et al., 2017)), one has that

$$\left\|\mathbb{D}(\vec{p}(t) - \vec{p}^*(t))\right\| \leq (\|{-\mathbb{M}} - \mathbb{N}\| + \lambda d_{N^*+1})\, D_{N^*} \times \\ \times \frac{e^{\gamma(\mathbb{A})t} - 1}{\gamma(\mathbb{A})} \sup_{u \in [0,t]} p_{N^*}^*(u). \quad (19)$$

Once $\vec{p}^*(0)$ (equal to $\vec{p}(0)$) is chosen, the right-most term on the right-hand side of (19) can be computed from the (practically) finite system of ODEs $\frac{d}{dt}\vec{p}^*(t) = (\mathbb{Q}^*)^{\mathrm{T}}\vec{p}^*(t)$, with the conventional numerical methods. Clearly, since $d_i \geq 1$, $\left\|\vec{p}(t) - \vec{p}^*(t)\right\| \leq \left\|\mathbb{D}(\vec{p}(t) - \vec{p}^*(t))\right\|$. Thus (provided that the appropriate values of d_i and c are known for the given values of λ, p_1 and μ) one can answer the question (ii) from (19) by applying simple exhaustive search.

In order to highlight the usefulness of such results like (19), note that the value of N^*, which solves (ii), does not necessarily guarantee, for example, that the average number $\mathsf{E}N^*(t) = \sum_{n=1}^{\infty} n\vec{p}^*(t)^{\mathrm{T}}\vec{1}$ of customers in the model with the finite queue is close to $\mathsf{E}N(t)$. But since

$$|\mathsf{E}N(t) - \mathsf{E}N^*(t)| \leq \frac{1}{\inf_{i \geq 1} \frac{D_{i+1}}{i}} \left\|\mathbb{D}(\vec{p}(t) - \vec{p}^*(t))\right\|,$$

one can again use (19) and exhaustive search to detect the value of N^*, which makes $\mathsf{E}N^*(t)$ as close to $\mathsf{E}N(t)$ as required.

NUMERICAL EXPERIMENT

In order to illustrate the findings of the previous sections, consider the following simple example. Fix the service rate $\mu = 1$ and the arrival rate $\lambda = 0.25$. Let $p_1 = 0.55$ i.e. almost half of the customers require two servers. Under these conditions the model is stable and the load is equal to ≈ 0.21. From the Fig. 2 it can be seen that for $p_1 = 0.55$ and $\lambda = 0.25$ it is possible to choose c and $d > 1$ such that the logarithmic norm $\gamma(\mathbb{A})$ is negative. Clearly, $c = \mu$. Computations from (16) yield $d = 2.65$ and $\gamma(\mathbb{A}) = -0.0677$. Assuming that initially the model is empty, one can apply (11) to compute the instant t^* beyond which the system can be considered stable with the error less than $\epsilon = 10^{-2}$. This value ($t^* \approx 73.5$) is depicted as the vertical dashed line in the Fig. 3 alongside with the time–dependent probabilities $p_2(t)$, $p_3(t)$ and $p_4(t)$.

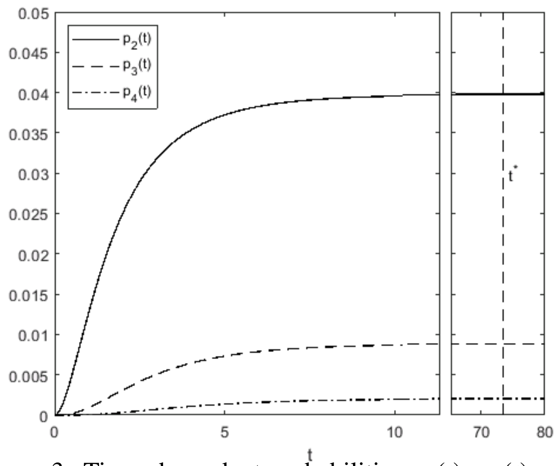

Figure 3: Time–dependent probabilities $p_2(t)$, $p_3(t)$ and $p_4(t)$ as the function of time t

The probabilities depicted in the Fig. 3 are the solutions of the systems (7), which was truncated at some high level, chosen arbitrarily. If now one uses (19) to detect the proper value of N^* (with the same ϵ), then one finds that $N^* = 16$. In the Fig. 4 one can see the behaviour of the average number of customers in the model at instant t and its upper bound according to (12). Using the steady state value of the upper bound (it is equal to ≈ 5.54) one immediately obtains from the Little's law that the customer's average sojourn time in the system is below ≈ 22.16. As it is commented below, this bound can be significantly improved by choosing other values of c and d; for example, for $d = 2.14$ and $c = 0.437$ we have $\gamma(\mathbb{A}) \approx -0.1522$ and $\lambda^{-1}\mathsf{E}N \approx 5.37$.

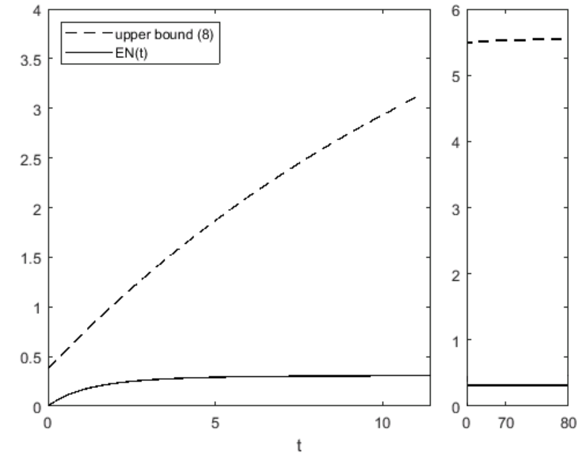

Figure 4: Average number $\mathsf{E}N(t)$ of customers in the model at instant t and its upper bound according to (12)

SUMMARY

Even though the two questions (i) and (ii) considered above concerned certain upper bounds, the logarithmic norm method can also be used to obtain inequalities similar to (6) but reversed. This naturally leads to the solution of some questions, which involve lower bounds (see (Zeifman et al., 2018b)); for example, find an estimate of the instant $t_* > 0$ such that the model cannot become stable earlier than at $t = t_*$.

All the analysis in the paper was carried out under the assumption that the model parameters are time–independent. The most promising aspect of the adopted method is that it allows generalizations (at least) towards the time–varying arrival rate. This can clearly be seen from (6), which permits the generator to be time–dependent. Yet the conditions under which the logarithmic norm $\gamma(\mathbb{A}(t))$ is negative, are to be found. This seems to be possible at least in those settings, when the arrival rate function is bounded (such are known for a long time, (Calzarossa and Serazzi, 1985)).

One of the drawbacks of the obtained solutions to the questions (i) and (ii) is that the obtained upper bounds are not sharp. Specifically, with respect to the value t^* (see (10)–(11)), it means that the model becomes stable earlier than at instant $t = t^*$. This is clearly seen from the figures in the previous section. According to the general theory (see, for example, (Zeifman et al., 2018a; Satin et al., 2020)) this effect is due to the fact that the matrix $\mathbb{A}$ has negative row elements off the main diagonal. In our experiments we were unable to detect the values of d_i, $i \geq 2$, which would fix the issue.

Probably the most serious defect in the solutions is that they are applicable not in the whole stability region of the model: with $c = \mu$ and $d_i = d$, $i \geq 2$, low system's load (below 0.3 when $p_1 \approx 0$ and below 0.15 when $p_1 \approx 1$) is the only feasible region (see (17) and Fig. 2). Yet there do exist several ways to transform it. One can manipulate $f(d)$ in (14), or vary the values of c and d_i, $i \geq 2$. For example, by noting that $\frac{f(d)}{d^2} < \lambda(d-1) - \mu + \frac{2\mu}{d}$ one can repeat the derivations of Section 5 and obtain that the feasible region is $\lambda \in (0, 0.1635\mu)$ (see the line with the circle markers in the Fig. 2). It is narrower than (17): below 0.17 when $p_1 \approx 0$ and below 0.09 when $p_1 \approx 1$. If one keeps $f(d)$ in (14) and puts $c = \mu\frac{2}{2-p_1^2}$ instead of $c = \mu$ in $\mathbb{A}$ (keeping $d_i = d > 1$ for all $i \geq 2$ unchanged), then the feasible region is transformed but again becomes narrower than the original (see the line with the asterisk markers in the Fig. 2). Our preliminary analysis shows that the good choice for c seems to be $c = \mu \cdot g(d, p_1)$, where a function g is defined in the semi-infinite strip $(d, p_1) \in (1, \infty) \times (0, 1)$ and is everywhere positive and non-increasing in d. Computations show that when the feasible region is made wider (by manipulating c), the value of $\gamma(\mathbb{A})$ becomes smaller (invoking too pessimistic upper bounds). Thus, if the described steps are followed, a trade-off must exist between the applicability and the desired accuracy. Filling the gaps here, as well as the choice of d_i, are the directions of further research.

Acknowledgements The research has been prepared with the support of Russian Science Foundation according to the research project No. 21-71-10135.

REFERENCES

Arns, M., P. Buchholz, and A. Panchenko. 2010. On the numerical analysis of inhomogeneous continuous–time Markov chains. INFORMS Journal on Computing. Vol. 22. No. 3. Pp. 416–432.

Brill, P. H., and L. Green. 1984. Queues in which customers receive simultaneous service from a random number of servers: A system point approach. Management Science. Vol. 30. No. 1. Pp. 51–68. DOI: 10.1287/mnsc.30.1.51

Burak, M. R., and P. Korytkowski. 2020. Inhomogeneous CTMC birth–and–death models solved by uniformization with steady–state detection. ACM Transactions on Modeling and Computer Simulation. Vol. 30. No. 3. Pp. 1–18.

Calzarossa, M., and G. Serazzi. 1985. A characterization of the variation in time of workload arrival patterns. IEEE Trans. Comput. Vol. C-34. No. 2. Pp. 156–162.

Chakravarthy, S.R., and H.D. Karatza. 2013. Two–server parallel system with pure space sharing and markovian arrivals. Computers & Operations Research. Vol. 40. No. 1. Pp. 510–519.

Clark, G. M. 1981. Use of Polya distributions in approximate solutions to nonstationary $M/M/s$ queues. Communications of the ACM. Vol. 24. No. 4. Pp. 206–217.

Filippopoulos, D., and H. Karatza. 2007. An $M/M/2$ parallel system model with pure space sharing among rigid jobs. Mathematical and Computer Modelling. Vol. 45. Pp. 491–530.

Harchol-Balter, M. 2021. Open problems in queueing theory inspired by datacenter computing. Queueing Systems. 2021. Vol. 97. Pp. 3–37.

M. Harchol-Balter. 2022. The multiserver job queueing model. Queueing Systems. DOI: 10.1007/s11134-022-09762-x.

Kim, S. 1979. $M/M/s$ queueing system where customers demand multiple server use. PhD Thesis. Southern Methodist University.

Massey, W. and J. Pender. 2013. Gaussian skewness approximation for dynamic rate multiserver queues with abandonment. Queueing Systems. Vol. 75. No. 2. Pp. 243–27.

Morozov, E. V., and A. S. Rumyantsev. 2011. Multi–server models to analyze high performance cluster. Transactions of Karelian Research Centre of the Russian Academy of Sciences. No. 5. Pp. 75–85. (in Russian)

Rumyantsev, A., and E. Morozov. 2017. Stability criterion of a multiserver model with simultaneous service. Ann. Oper. Res. Vol. 252. No. 1. Pp. 29–39.

Rumyantsev A. 2020. Stability of Multiclass Multiserver Models with Automata-type Phase Transitions. Proceedings of the Second International Workshop on Stochastic Modeling and Applied Research of Technology (SMARTY 2020). pp 213225. URL: http://ceur-ws.org/Vol-2792/#paper16

Rumyantsev A. et al. 2021. A Three-Level Modelling Approach for Asynchronous Speed Scaling in High-Performance Data Centres. Proceedings of the Twelfth ACM International Conference on Future Energy Systems. Association for Computing Machinery, New York, NY, USA, pp 417423. DOI: 10.1145/3447555.3466580

Satin, Y., K. Kiseleva, S. Shorgin, V. Korolev, and A. Zeifman. 2017. Two–sided truncations for the $M_t/M_t/S$ queueing model. Proceedings of the 31st European Conference on Modelling and Simulation. Pp. 635–641.

Satin, Ya., A. Zeifman, and G. Shilova. 2020. On approaches to constructing limiting regimes for some queuing models. Informatika i ee Primeneniya — Inform. Appl. Vol. 15. No. 2. Pp. 19–24.

Taaffe, M. R., and K. L. Ong. 1987. Approximating nonstationary $Ph(t)/M(t)/s/c$ queueing systems. Annals of Operations Research. Vol. 8. No. 1. Pp. 103–116.

Van Dijk, N. M., S. P. J. van Brummelen, and R. J. Boucherie. 2018. Uniformization: Basics, extensions and applications. Performance evaluation. Vol. 118. Pp. 8-32.

Zeifman, A., Y. Satin, I. Kovalev, R. Razumchik, and V. Korolev. 2021. Facilitating numerical solutions of inhomogeneous continuous time Markov chains using ergodicity bounds obtained with logarithmic norm method. Mathematics. Vol. 9. No. 1. Art. ID 42. 20 p.

Zeifman, A.I. 1995. Upper and lower bounds on the rate of convergence for nonhomogeneous birth and death processes. Stoch. Proc. Appl. Vol. 59. Pp. 157–173.

Zeifman, A., Y. Satin, K. Kiseleva, and V. Korolev. 2019. On the rate of convergence for a characteristic of multidimensional birth–death process. Mathematics. Vol. 7. Iss. 5. Art. ID 477. 10 p.

Zeifman, A., Y. Satin, V. Korolev, and S. Shorgin. 2014. On truncations for weakly ergodic inhomogeneous birth and death processes. Int. J. Appl. Math. Comput. Sci. Vol. 24. Pp. 503–518.

Zeifman, A., A. Sipin, V. Korolev, G. Shilova, K. Kiseleva, A. Korotysheva, and Y. Satin. 2018. On sharp bounds on the rate of convergence for finite continuous–time markovian queueing models. LNCS. Vol. 10672. Pp. 20–28.

Zeifman, A. I., V. Y. Korolev, Y. A. Satin, and K. M. Kiseleva. 2018. Lower bounds for the rate of convergence for continuous–time inhomogeneous Markov chains with a finite state space. Statistics & Probability Letters. Vol. 137. Pp. 84–90.

Zeifman, A., Y. Satin, K. Kiseleva, V. Korolev, and T. Panfilova. 2019. On limiting characteristics for a non-stationary two-processor heterogeneous system. Applied Mathematics and Computation. Vol. 351(C). Pp. 48–65.

AUTHOR BIOGRAPHIES

ROSTISLAV RAZUMCHIK received his Ph.D. degree in Physics and Mathematics in 2011. Since then, he has worked as the leading research fellow at the Institute of Informatics Problems of the Federal Research Center "Computer Science and Control" of the Russian Academy of Sciences (FRC CSC RAS). His current research activities are focused on queueing theory and its applications for performance evaluation of stochastic systems. His email address is rrazumchik@ipiran.ru

RUMYANTSEV ALEXANDER received his PhD from Petrozavodsk State University. He is now a researcher in the Institute of Applied Mathematical Research of the Karelian Research Centre of the Russian Academy of Sciences. His research interests include Stochastic Processes, Queueing Systems, High–Performance and Distributed Computing, Multi-Core and Many-Core Systems. His email address is ar0@krc.karelia.ru.