VU Research Portal

Optimal job splitting in parallel processor sharing queues

Hoekstra, G.; van der Mei, R.D.; Bhulai, S.

published in

Stochastic Models
2012

DOI (link to publisher)

[10.1080/15326349.2012.646555](https://doi.org/10.1080/15326349.2012.646555)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Hoekstra, G., van der Mei, R. D., & Bhulai, S. (2012). Optimal job splitting in parallel processor sharing queues. *Stochastic Models*, 28(1), 144-166. <https://doi.org/10.1080/15326349.2012.646555>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Stochastic Models

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/Istm20>

Optimal Job Splitting in Parallel Processor Sharing Queues

G. J. Hoekstra^{a b}, R. D. van der Mei^{b c} & S. Bhulai^{b c}

^a Thales, Innovation Research & Technology, Huizen, The Netherlands

^b CWI, Probability and Stochastic Networks, Amsterdam, The Netherlands

^c Department of Mathematics, VU University, Amsterdam, The Netherlands

Published online: 03 Feb 2012.

To cite this article: G. J. Hoekstra, R. D. van der Mei & S. Bhulai (2012) Optimal Job Splitting in Parallel Processor Sharing Queues, *Stochastic Models*, 28:1, 144-166, DOI: [10.1080/15326349.2012.646555](https://doi.org/10.1080/15326349.2012.646555)

To link to this article: <http://dx.doi.org/10.1080/15326349.2012.646555>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

OPTIMAL JOB SPLITTING IN PARALLEL PROCESSOR SHARING QUEUES

G. J. Hoekstra^{1,2}, R. D. van der Mei^{2,3}, and S. Bhulai^{2,3}

¹Thales, Innovation Research & Technology, Huizen, The Netherlands

²CWI, Probability and Stochastic Networks, Amsterdam, The Netherlands

³Department of Mathematics, VU University, Amsterdam, The Netherlands

□ *The main barrier to the sustained growth of wireless communications is the Shannon limit that applies to the channel capacity. A promising means to realize high-capacity enhancements is the use of multi-path communication solutions to improve reliability and network performance in areas that are covered by a multitude of overlapping wireless access networks. Despite the enormous potential for capacity enhancements offered by multi-path communication techniques, little is known about how to effectively exploit this. Motivated by this, we study a model where jobs are split and downloaded over N multiple parallel networks, each of which is modeled as a processor sharing (PS) queue. Each job is fragmented, according to a fixed splitting rule $\underline{\alpha} = (\alpha_1, \dots, \alpha_N)$ and re-assembled at the receiving end. The complex correlation structure between the sojourn times at the PS nodes makes an exact detailed mathematical analysis of the model impossible. Therefore, in this article we propose a simple and fast approximation for the splitting rule $\underline{\alpha}^*$ that minimizes the expected job-download time. Our approximation is validated extensively by simulations. The results show that the outcomes are extremely accurate over a wide range of parameter combinations.*

Keywords Concurrent access; File splitting; Flow-level performance; Processor sharing; Traffic splitting.

Mathematics Subject Classification Primary 68M20, 60K25; Secondary 90B22.

1. INTRODUCTION

The Shannon limit on channel capacity is already closely approached by some of today's wireless networks, leaving complex signal processing techniques room for only modest improvements in the data transmission

Received June 2010; Accepted April 2011

Address correspondence to R. D. van der Mei, CWI, Probability and Stochastic Networks, VU University, Amsterdam, The Netherlands and Department of Mathematics, Amsterdam, The Netherlands; E-mail: mei@cwi.nl

rate^[6]. Concurrently using multiple, possibly different, networks then becomes an alternative to increase the overall data rate, because (a) the spectrum is regulated among various frequency bands and corresponding communication network standards, and (b) the overall spectrum usage remains relatively low over a wide range of frequencies^[8]. In areas that are covered by multiple wireless access networks, the concurrent use of multiple networks simultaneously opens up tremendous possibilities for increasing capacity, improving reliability, and enhancing Quality of Service (QoS). Despite the enormous potential for quality improvement, only little is known about how to fully exploit this potential. A main requirement for the widespread use of traffic-splitting algorithms for concurrent access is that the algorithms are simple, yet effective. Motivated by this, we concentrate on splitting rules that (a) require only aggregated information on the network status, and (b) are easy to enforce by deciding upon the splitting ratio of the jobs only once. To this end, we study static splitting rules that determine how jobs should be partitioned into fractions $(\alpha_1, \dots, \alpha_N)$, transferred over the different overlapping networks and reassembled, to minimize the expected download time of the entire job. We study this file-splitting problem in a queueing theoretical context. Modeling network performance using processor sharing (PS) based models^[4,19,21] is applicable to a variety of communication networks, including CDMA, 1xEV-DO, WLAN, and UMTS-HSDPA. It was shown that PS models can actually model file transfers over WLANs accurately^[14], taking into account the complex dynamics of the protocol-stack, including their interactions.

In the literature on telecommunication systems, the concurrent use of multiple network resources in parallel was already described for a Public Switched Digital Network (PSDN)^[7]. Here inverse multiplexing was proposed as a technique to perform the aggregation of multiple independent information channels across a network to create a single higher-rate information channel. Various approaches have appeared to exploit multiple transmission paths in parallel. Examples are using multi-element antennas, as adopted by the IEEE802.11n standard^[1], splitting at the physical layer or switching datagrams at the link layer^[5,16], and also using multiple TCP sessions in parallel to a file-server^[20]. In the latter case each available network transports part of the requested data in a separate TCP session. Previous work has indicated that downloading from multiple networks concurrently may not always be beneficial^[9], but in general significant performance improvements can be realized^[11,13,15]. Under these circumstances of using a combination of different network types, the transport layer approaches in particular have shown their applicability^[15], as they allow appropriate link layer adaptations for each TCP session. In addition, many papers have studied traffic distribution algorithms in a

more theoretical framework. In Ref.^[10] the authors investigate the same model as the one under consideration, but without the presence of background traffic and with the Join the Shortest Queue (JSQ) policy instead of the static splitting rules considered in the present article. In Refs.^[17,18], the author analyzes a similar model but with FCFS queues and with probabilistic splitting. We further refer to Altman et al.^[2], who consider routing policies in a distributed versus centralized environment. In general our queueing model falls within the framework of fork-join queueing networks, see Ref.^[3] for an extensive overview. In a recent article^[12], the theoretical foundation for a tail-optimal splitting rule is provided for light foreground load that is shown to work well with respect to both the tail asymptotics and the mean sojourn times.

In this article we study static job-splitting rules $\underline{\alpha} = (\alpha_1, \dots, \alpha_N)$, where a job of size τ is split into N tasks of size $\alpha_i \tau$ ($i = 1, \dots, N$), where the i th task is processed by PS node i , and reassembled upon completion of all the tasks. In addition, we assume the presence of background traffic at each of the nodes. The goal is to find a splitting rule $\underline{\alpha}^*$ that minimizes $\mathbb{E}[S_0^{\underline{\alpha}}]$, where $S_0^{\underline{\alpha}}$ is the total processing time of an entire foreground job, which generally depends on the file-size distributions and on the characteristics of the background traffic streams. Unfortunately, this model does not allow for an exact analysis. The complexity lies in the fact that the sojourn times of the fragments in the different PS nodes are generally correlated. Therefore, we develop a new approximation for $\mathbb{E}[S_0^{\underline{\alpha}}]$, combining light- and heavy-traffic asymptotics, which then leads to an approximation for $\underline{\alpha}^*$. The approximation is validated by extensive simulations over a wide range of parameter combinations, including light- and heavy-tailed job-size distributions, and mixtures of light- and heavy-load scenarios on foreground and background traffic. These simulations demonstrate that the differences between the approximated optimal splitting rule and the estimated optimum with respect to the expected foreground sojourn time are extremely small for a wide range of the parameter settings.

The organization of this article is as follows. In Section 2 the model is described and the notation is introduced. In Section 3 we analyze the performance of the model and use these insights to develop a new approximation method for determining the optimal split $\underline{\alpha}^*$. In Section 4 the accuracy of the approximation method is discussed in detail. Finally, in Section 5 we address a number of topics for further research.

2. MODEL DESCRIPTION

We consider a job-split model consisting of N parallel PS nodes, $PS_1, \dots, PS_N$, operating at the same speed (see Figure 1). Each of these PS nodes in our model corresponds to a communication network. Files are

FIGURE 1 Illustration of the job-split PS-model.

modeled as jobs that are fragmented into tasks. There are $N + 1$ traffic streams: stream 0 is called the foreground stream and streams $1, \dots, N$ are called the background streams. Jobs of background stream i are not fragmented and are served exclusively at PS_i ($i = 1, \dots, N$). Jobs of the foreground stream are fragmented into tasks upon arrival according to a fixed splitting rule $\underline{\alpha} = (\alpha_1, \dots, \alpha_N)$ where $\sum_{i=1}^N \alpha_i = 1$ and $\alpha_i \geq 0$, $i = 1, \dots, N$; thus a foreground job of size $B = \tau$ is split into N tasks of size $\alpha_i \tau$, and the i th task is processed by PS_i ($i = 1, \dots, N$). Once all tasks have been completed, they are reassembled, which completes the processing of the job. Our probabilistic and load assumptions are as follows: arrivals of jobs in all streams are according to independent Poisson processes with rates λ_i , $i = 0, 1, \dots, N$. The total arrival rate is denoted by $\Lambda = \lambda_0 + \lambda_1 + \dots + \lambda_N$. For all streams, each job size B is an independent sample from a general distribution with k th moment $\beta^{(k)} = \mathbb{E}[B^k]$, for $k = 1, 2, \dots$. Denote the background load of stream i by $\rho_i = \lambda_i \beta^{(1)}$ ($i = 0, 1, \dots, N$), and denote the total load offered to the system by $\rho = \rho_0 + \rho_1 + \dots + \rho_N$. The utilization of node i is denoted by

$$\xi_i := \rho_i + \alpha_i \rho_0, \quad \text{and let } \xi := \max_{i=1, \dots, N} \xi_i. \quad (1)$$

For stability, it is assumed that $\xi < 1$. Note that if $\rho > 1$ the stability condition may pose restrictions on the choice of the splitting rule $\underline{\alpha}$. Denote by $A := \{\underline{\alpha} : \xi < 1\}$, i.e., the set of combinations for which the stability conditions are met. A splitting rule $\underline{\alpha}$ is called *feasible* if $\underline{\alpha} \in A$.

For an arbitrary foreground job, denote by $S_i^{\underline{\alpha}}$ the sojourn time of its i th task operating under the splitting rule $\underline{\alpha}$. This is the time it takes the

i th task to complete service at PS_i . Denote $\underline{S}^\alpha = (S_1^\alpha, \dots, S_N^\alpha)$. The sojourn time of a foreground job through the job-split model is denoted by

$$S_0^\alpha := \max \{S_1^\alpha, \dots, S_N^\alpha\}. \quad (2)$$

Our purpose is to find a splitting rule $\underline{\alpha}^* = (\alpha_1^*, \dots, \alpha_N^*) \in A$ such that

$$\mathbb{E}[S_0^{\underline{\alpha}^*}] = \min_{\underline{\alpha} \in A} \mathbb{E}[S_0^\alpha]. \quad (3)$$

For a non-negative random variable X with finite positive first moment, the squared coefficient of variation is denoted c_X^2 . Finally, heavy-traffic limits for $\xi \uparrow 1$ are taken such that the total arrival rate Λ is increased while the service-time distribution, the splitting rule $\underline{\alpha}$ and the *proportions* between the arrival rates $\lambda_0, \lambda_1, \dots, \lambda_N$ remain fixed. Note that in this limiting regime, not all nodes tend to become unstable as $\xi \uparrow 1$. More precisely, node i becomes unstable for $\xi \uparrow 1$ only if $\xi_i = \xi$, and otherwise node i remains stable as $\xi \uparrow 1$. Denote the set of potentially unstable queues by $U := \{i : \xi_i = \xi\}$, where ξ_i and ξ are defined in (1).

3. ANALYSIS

In general, the “cost function” $\mathbb{E}[S_0^\alpha]$ does not allow for an exact expression, and the optimization problem defined in (2) and (3) cannot be solved explicitly. The mathematical complexity is caused by the correlations between the sojourn times $S_1^\alpha, \dots, S_N^\alpha$ of the jobs at the different nodes. This dependence is caused by the fact that the (fragmented) foreground tasks arrive at the nodes simultaneously, and by the fact that their sizes $\alpha_1 B, \dots, \alpha_N B$ are correlated. For this reason, in this section we will develop heuristic methods to approximate $\mathbb{E}[S_0^\alpha]$ and the optimal splitting rule $\underline{\alpha}^*$. The approximation of $\underline{\alpha}^*$ is based on an interpolation between two components. The first component is based on the concept of reduced-load equivalence (RLE), and works well in light-traffic scenarios. The second component is based on heavy-load asymptotics, and complements the RLE-based approximation for heavy-load scenarios.

In Section 3.1 we formulate some known results on multiclass PS models and present a number of simulation results that lead to observations that are useful for later reference. In Section 3.2 we outline the RLE-based approximation and in Section 3.3 we present the heavy-load approximation. Subsequently, in Section 3.4 both approximations for $\underline{\alpha}^*$ are combined into our composed-split approximation, which interpolates $\underline{\alpha}$ between these two components.

3.1. Preliminaries

Considering node i in isolation, it is easy to see that this node can be modelled as a two-class M/G/1-PS model, where class-1 represents the tasks originating from the foreground traffic and class-2 the background traffic. Class-1 tasks arrive according to a Poisson process with rate λ_0 and size $\alpha_i B$, where B is the size of an arbitrary job. Similarly, class-2 jobs arrive according to a Poisson process with rate λ_i and job size B . For this model, it is known that for given $B = \tau$, the conditional expected sojourn time of a foreground task of size $\alpha_i \tau$ at node i is given by:

$$\mathbb{E}[S_i^{\underline{\alpha}} | B = \tau] = \frac{\alpha_i \tau}{1 - \xi_i}, \quad \text{and hence,} \quad \mathbb{E}[S_i^{\underline{\alpha}}] = \frac{\alpha_i \beta^{(1)}}{1 - \xi_i} \quad (i = 1, \dots, N), \quad (4)$$

where $\xi_i = \beta^{(1)}(\lambda_i + \alpha_i \lambda_0)$ is the utilization of node i . However, despite the fact that the conditional mean sojourn times $\mathbb{E}[S_i^{\underline{\alpha}} | B = \tau]$ of the *individual* tasks at each of the N nodes are known, an exact expression for the mean sojourn time of *entire* foreground job, $\mathbb{E}[S_0^{\underline{\alpha}}]$ (defined in (2) and (3)), which is defined as the maximum of the (correlated) per task sojourn times is not known.

Prior to developing the approximations for the optimal splitting rule $\underline{\alpha}^*$, we perform numerical experiments based on simulations to gain insight in the optimization problem. This will lead to a number of important observations that will turn out to be useful for later reference. As an illustrative example, for the case $N = 2$ and $\beta^{(1)} = 1$, Figure 2 shows the behavior of $\mathbb{E}[S_0^{\underline{\alpha}}]$ for the traffic splitting rule $\underline{\alpha} = (\alpha, 1 - \alpha) \in A$ as a function of α ($0 < \alpha < 1$), for different background and foreground load scenarios. To highlight the impact of the service-time distributions, results are shown for the extreme cases of deterministic service times (with $c_B^2 = 0$) and Pareto-2 distributed service times with $\Pr\{B > x\} = \frac{1}{4x^2}$ for $x > 1/2$, so that $c_B^2 = \infty$.

Figure 2 illustrates the behavioral differences between the various systems. In the absence of background traffic, the curves exhibit wedge-shaped behavior around their optimum, both for light ($\rho_0 = 0.1$) and moderate ($\rho_0 = 0.9$) foreground load. For non-zero but mild background load and light foreground load, the curve is nearly constant around its optimum. This is not the case if the background load becomes highly asymmetric ($\rho_1 = 0, \rho_2 = 0.9$) in the presence of light foreground load; the system is not stable for all values of α and exhibits a sharp increase in the mean sojourn time for slightly underestimated values of α (below 0.95) and is rather forgiving to overestimations where it coincides with the system without background load. Considering the outcomes depicted in Figure 2, it is clear that heavy foreground traffic yields a bended curve, showing that the cost function is highly sensitive to the choice α .

FIGURE 2 $\mathbb{E}[S_0^2]$ as a function of the traffic splitting ratio α , based on simulations (color figure available online).

The most interesting observation from Figure 2 is that both the cost function $\mathbb{E}[S_0^2]$ and optimal splitting ratio α^* are nearly insensitive to the job-size distribution. This observation is quite remarkable: Although it is well-known that the per-node mean sojourn times are (fully) insensitive to the distribution of the job sizes (see Eq. (4)), no general results are known for the expected value of the maximum of the per-task sojourn times, which in general are mutually dependent.

The question arises as to what the impact of the job-size distribution is on the correlations between the per-node sojourn times of the foreground tasks of sizes $\alpha_1 B, \dots, \alpha_N B$. Recall that the correlations between the sojourn times are caused by the foreground traffic, because (a) the foreground traffic stream generates simultaneous arrivals of tasks at each of the nodes, and (b) the job sizes of the per-node tasks $\alpha_1 B, \dots, \alpha_N B$ are stochastically dependent. Intuitively, one may expect that the higher the foreground load, the stronger the correlations.

To validate this, we have performed simulation experiments for a two-node model, with $\beta^{(1)} = 1$ and split rule $\underline{\alpha} = (1/2, 1/2)$, where ρ_0 , ρ_1 , and ρ_2 are parameterized as follows $\rho_1 = 1 - 2\rho_0/3$, $\rho_2 = 1/2 - \rho_0/3$, and where ρ_0 is varied between 0 and $3/2$. In this way, ρ_0 is varied over the interval $[1/10, 3/2]$ such that the total load $\rho = \rho_0 + \rho_1 + \rho_2$ is kept fixed at value $\rho = 3/2$, while the ratios between ρ_1 and ρ_2 are fixed to $\rho_1 = 2\rho_2$. For each foreground job, we have calculated the statistical correlation between the two per-node tasks (both of size $\tau/2$), and the mean sojourn times of the foreground jobs. Simulations have been run for 10^{11} jobs, which led to extremely narrow confidence intervals (not shown here).

FIGURE 3 The correlations between the sojourn times of the foreground traffic and $\mathbb{E}[S_0^\alpha]$ as a function of ρ_0 with fixed total load $\rho = 1.5$, based on simulations. (a) Correlation as a function of ρ_0 and (b) Sojourn time as a function of ρ_0 (color figure available online).

Figure 3(a) shows the empirical correlation between the sojourn times of the foreground jobs considered as a function of ρ_0 , where the service-time distributions are varied as deterministic ($c_B^2 = 0$), exponential ($c_B^2 = 1$), Erlang-2 ($c_B^2 = 1/2$), two-phase hyper-exponential with $c_B^2 = 4$ and $c_B^2 = 16$ (and balanced means), Pareto-3 (with $\Pr\{B > x\} = (1 + x/2)^{-3}$ for $x > 0$ and hence $c_B^2 = 3$), and Pareto-2 (with $\Pr\{B > x\} = \frac{1}{4x^2}$ for $x > 1/2$ and hence $c_B^2 = \infty$). Note that in Figure 3(a) the results for Pareto-3 and the two-phase hyper-exponential with $c_B^2 = 4$ are so close that they can hardly be distinguished. Moreover, the results in Figure 3 are so strikingly similar that they almost entirely overlap.

The results depicted in Figures 3(a) and (b) lead to the following remarkable observation.

Observation 1. The correlations between the per-task sojourn times of the foreground traffic depend on the job-size distribution, whereas $\mathbb{E}[S_0^\alpha]$ is nearly insensitive to the job-size distribution.

This observation is rather intriguing. First, we observe an obvious dependence of the correlation between the per-task sojourn times with respect to the job-size distribution, where higher variability seems to imply a stronger correlation. This observation can be intuitively explained by the fact that the per-task sojourn times are positively correlated, while this correlation becomes most predominant for large job sizes, thus for outliers in the job size B . Hence, the higher the variability in the distribution of B , the more outliers in B and hence, the stronger the correlation. Second, the results show that the differences in correlations over the different service-time distributions do not manifest themselves in significant differences in $\mathbb{E}[S_0^\alpha]$. The impact of the correlations between the per-task sojourn times

“vanishes” when looking at $\mathbb{E}[S_0^{\underline{\alpha}}]$. This observation will turn out to be useful for developing an approximation for $\mathbb{E}[S_0^{\underline{\alpha}}]$, see Sections 3.3 and 3.4.

3.2. Reduced Load Approximation (RLA)

The Reduced Load Approximation (RLA) splits jobs into tasks according to the split rule (cf. Ref.^[12]):

$$\alpha_{RLA,i}^* := \frac{1 - \rho_i}{\sum_{j=1}^N (1 - \rho_j)} \quad (i = 1, \dots, N). \quad (5)$$

This simple splitting rule $\underline{\alpha}_{RLA}^* = (\alpha_{RLA,1}^*, \dots, \alpha_{RLA,N}^*)$ cuts the foreground jobs into tasks proportional to $1 - \rho_i$, i.e., the average amount of capacity not used by the background traffic at node i . Note that the RLA is insensitive to the job-size distribution (except of course for its mean $\beta^{(1)}$), which is in line with Observation 1 in Section 3.1. In addition to its attractive simplicity, the RLA (5) is asymptotically tail optimal for the case of regularly varying service-time distributions (see Ref.^[12] for details).

Extensive numerical experimentation in Ref.^[12] (and Table 1 in Section 4 below) reveals that the RLA leads to highly accurate approximations of $\underline{\alpha}^*$ if either $\rho_0 \approx 0$ or if the nodes are fairly equally loaded, but that it may become inaccurate when one or more nodes are heavily loaded while others are not (see for example the right upper corner in Table 1 below). This raises the need for a refinement of the RLA to improve its accuracy for asymmetric background-load scenarios, which is the main goal of the present article. To this end, in the next section we use and combine known heavy-traffic asymptotics for multi-class PS models and Observation 1 to derive approximations for $E[S_0^{\underline{\alpha}}]$, and hence for the optimal splitting of jobs $\underline{\alpha}^*$, under heavy-traffic assumptions.

3.3. Heavy Traffic Approximation (HTA)

In this section we will use heavy-traffic (HT) asymptotics for multi-class PS models to develop an approximation for $\mathbb{E}[S_0^{\underline{\alpha}}]$, and hence of $\underline{\alpha}^*$, that meets these asymptotic HT-properties. To formulate these HT properties, recall that node i considered in isolation can be modeled as a two-class M/G/1 PS model, where class-1 jobs (representing background jobs at node i) arrive according to a Poisson process with rate λ_0 and service times $\alpha_i B$, and where class-2 jobs (representing foreground jobs) arrive according to a Poisson process with rate λ_i and service times B . Let $S_i(\tau)$ denote the sojourn time of an arbitrary job of size τ at node i (regardless of its class). Zwart and Boxma^[22] show that for $\tau > 0$, $i \in U$, $\underline{\alpha} \in A$,

$$(1 - \xi)S_i(\alpha_i \tau) \rightarrow_d \Theta(\alpha_i \tau) \quad (\xi \uparrow 1), \quad (6)$$

where $\Theta(\zeta)$ is an exponentially distributed random variable with mean ζ . Moreover, in Ref.^[22] it is shown also that moment-wise convergence holds: For $\tau > 0$, $i \in U$, $\underline{\alpha} \in A$ and $k = 1, 2, \dots$,

$$\lim_{\zeta \uparrow 1} (1 - \zeta)^k \mathbb{E}[S_i(\alpha_i \tau)^k] = k! \alpha_i^k \tau^k. \quad (7)$$

To this end, it is important to observe that the conditional results in (6) and (7) were proven for the classical single-class M/G/1 PS node, but are also directly applicable to multi-class M/G/1 PS nodes. Then, removing the conditioning with respect to the distribution of the size of the foreground tasks at node i in Eq. (6) leads to the following result for S_i , the unconditional sojourn time for a job at node i : For $\tau > 0$, $i \in U$, $\underline{\alpha} \in A$,

$$(1 - \zeta) S_i \rightarrow_d \Theta(\alpha_i) B \quad (\zeta \uparrow 1), \quad (8)$$

where $\Theta(\alpha_i)$ is exponentially distributed with mean α_i , B is the job size, and where the random variables $\Theta(\alpha_i)$ and B are independent. Moreover, the k th moment of the unconditioned sojourn time of an arbitrary foreground job at node i has the following limiting HT-behavior: For $i \in U$, $\underline{\alpha} \in A$, $k = 1, 2, \dots$,

$$\lim_{\zeta \uparrow 1} (1 - \zeta)^k \mathbb{E}[S_i^k] = k! \alpha_i^k \beta^{(k)}. \quad (9)$$

We are now ready to use the HT-asymptotics (6) to (9) to develop a simple approximation for $\mathbb{E}[S_0^{\underline{\alpha}}]$ that works well under HT-circumstances, i.e., where $\zeta \uparrow 1$. In the absence of an exact analysis, we will construct a simple approximation for the joint probability distribution of

$$\underline{S}_{HTA}^{\underline{\alpha}}(\tau) = (S_{HTA,1}^{\underline{\alpha}}(\alpha_1 \tau), \dots, S_{HTA,N}^{\underline{\alpha}}(\alpha_N \tau)), \quad (10)$$

where $S_{HTA,i}^{\underline{\alpha}}(\alpha_i \tau)$ is the sojourn time of the i th fragment of an arbitrary foreground job of size τ (note that this fragment itself is of size $\alpha_i \tau$). To this end, note that (6) implies the following limiting behavior for the *marginal* distributions of the conditional sojourn times $S_{HTA,i}^{\underline{\alpha}}(\alpha_i \tau)$: For $i \in U$, $\underline{\alpha} \in A$, $t > 0$,

$$\lim_{\zeta \uparrow 1} \Pr\{(1 - \zeta) S_{HTA,i}^{\underline{\alpha}}(\alpha_i \tau) > t\} = \exp \left\{ -\frac{t}{\alpha_i \tau} \right\}. \quad (11)$$

Note that (11) is only valid for $i \in U$, because if $i \notin U$ then node i will not become unstable so that $(1 - \zeta) S_{HTA,i}^{\underline{\alpha}}(\alpha_i \tau) \rightarrow 0$ (a.s.), when $\zeta \uparrow 1$ (see also Remark 3.2).

Next, we develop an approximation for $\mathbb{E}[S_0^{\underline{\alpha}}]$ which satisfies the known heavy-traffic properties of the marginal per-task sojourn-time distributions formulated in (10) and (11). Regarding the correlations, recall from Observation 1 (formulated in Section 3) that $\mathbb{E}[S_0^{\underline{\alpha}}]$ at best weakly depends on the correlations between the per-task sojourn times. Therefore, although the per-task conditional sojourn times $S_{HTA,1}^{\underline{\alpha}}(\alpha_1\tau), \dots, S_{HTA,N}^{\underline{\alpha}}(\alpha_N\tau)$ are clearly not independent, we *assume* that they are. Based on this assumption, we *approximate* the distribution of $S_{HTA}^{\underline{\alpha}}(\tau)$ (defined in Eq. (10)) by assuming that $S_{HTA,1}^{\underline{\alpha}}(\alpha_1\tau), \dots, S_{HTA,N}^{\underline{\alpha}}(\alpha_N\tau)$ (a) are exponentially distributed with mean $\alpha_i\tau/(1-\xi_i)$, so that the marginal distributions satisfy the known HT behavior in (11) for $i \in U$ (see also Remark 3.2), and (b) are mutually independent: For $t_1, \dots, t_N > 0$, $\underline{\alpha} \in A$, $\tau > 0$,

$$\Pr\{S_{HTA,1}^{\underline{\alpha}}(\alpha_1\tau) > t_1, \dots, S_{HTA,N}^{\underline{\alpha}}(\alpha_N\tau) > t_N\} \approx \prod_{i=1}^N \exp\left\{-\frac{1-\xi_i}{\alpha_i\tau} t_i\right\}. \quad (12)$$

We will now use the approximation in (12), which covers the known HT-limiting behavior from (6)–(7), to derive a simple approximation for $E[S_0^{\underline{\alpha}}]$ that works well when $\xi \uparrow 1$. To this end, for $\tau > 0$, $\underline{\alpha} \in A$, define

$$S_{HTA,0}^{\underline{\alpha}}(\tau) := \max\{S_{HTA,1}^{\underline{\alpha}}(\alpha_1\tau), \dots, S_{HTA,N}^{\underline{\alpha}}(\alpha_N\tau)\}. \quad (13)$$

Then using Eq. (12), the distribution of $S_{HTA,0}^{\underline{\alpha}}(\tau)$ can be approximated as the distribution of the maximum of N independent exponentially distributed random variables with known parameters (given in Eq. (12)). See Property 1 in Appendix A for an expression for the mean value of such random variable. Using this result, the conditional cost function $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}(\tau)]$ can be readily obtained from Property 1 simply by making following substitutions in Eq. (21) from Appendix A: For $i = 1, \dots, N$,

$$X_i := S_{HTA,i}^{\underline{\alpha}}(\alpha_i\tau), \quad \text{and} \quad 1/\mu_i := \mathbb{E}[S_{HTA,i}^{\underline{\alpha}}(\alpha_i\tau)] = \frac{\alpha_i\tau}{1-\xi_i}. \quad (14)$$

Note that it is readily verified that the so-obtained approximation of $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}(\tau)]$ is *linear* in τ , so that the unconditioned cost function $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$ can be directly obtained by replacing τ by $\beta^{(1)}$ in (14). In this way, the cost function $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$ can be approximated by using Eqs. (12) to (14) and Eq. (21), which leads to the following approximation for $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$:

$$\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}] \approx \beta^{(1)} \sum_{k=1}^N (-1)^{k+1} \sum_{(i_1, \dots, i_k) \in S_k} \frac{1}{\frac{1-\xi_{i_1}}{\alpha_{i_1}} + \dots + \frac{1-\xi_{i_N}}{\alpha_{i_N}}}, \quad (15)$$

where

$$S_k := \{(i_1, \dots, i_k) : i_1, \dots, i_k \in \{1, \dots, N\}, i_1 < i_2 < \dots < i_k\}. \quad (16)$$

Note that the approximate expression for $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$ in Eq. (15) is explicit, and hence, the computation time is negligible. Notice also that the right-hand side of Eq. (15) is (fully) insensitive to the job-size distribution, which is in line with Observation 1 discussed in Section 3.1.

Next, denote by $\underline{\alpha}_{HTA}^* = (\alpha_{HTA,1}^*, \dots, \alpha_{HTA,N}^*)$ the splitting rule that minimizes $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$ among all $\underline{\alpha} \in A$, i.e.,

$$\mathbb{E}[S_{HTA,0}^{\underline{\alpha}^*}] = \min_{\underline{\alpha} \in A} \mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]. \quad (17)$$

The optimal split $\underline{\alpha}_{HTA}^*$ can then be approximated by evaluating Eq. (15) over all $\underline{\alpha} \in A$, or by some non-linear optimization method. In practice, this causes no problem as N is not too large. Note that in the context of concurrent access for wireless networks, which was the main motivation for this study, N is indeed small, say 2 or 3. We refer to Remark 4.2 for a brief discussion on the complexity of the optimization.

3.4. Composed Split Approximation (CSA)

Now we combine both approaches in the following to obtain the composed-split approximation (CSA), denoted $\underline{\alpha}_{CSA}^* = (\alpha_{CSA,1}^*, \dots, \alpha_{CSA,N}^*)$, where for $i = 1, \dots, N$,

$$\alpha_{CSA,i}^* = (1 - \kappa_i) \alpha_{RLA,i}^* + \kappa_i \alpha_{HTA,i}^*, \quad \text{with } \kappa_i := \max\{\rho_1, \dots, \rho_N\}, \quad (18)$$

and where $\alpha_{RLA,i}^*$ is given in Eq. (5), and where $\alpha_{HTA,i}^*$ can be obtained from Eqs. (15) to (17). The interpolation factor κ_i is taken such that κ_i is independent of $\underline{\alpha}$ (and of i , see also Remark 3.3), while for light-traffic scenarios the RLA is dominant and for heavy-load scenarios the HTA is dominant. We refer Remark 3.3 below for a discussion on the choice of the interpolation factor.

Remark 3.1. A complicating factor of the model is the complex correlation structure between the sojourn times of the individual tasks after a job has been split. In this context, we observe that the reduced load approximation (RLA) can be thought of as implying that there is “perfect” correlation in the sojourn times of the different tasks, and thus gives a lower bound on the mean sojourn times. To this end, note that under “perfect correlation,” for all $i, j = 1, \dots, N$, $\underline{\alpha} \in A$ and $\tau > 0$,

$$\mathbb{E}[S_i^{\underline{\alpha}} | B = \tau] = \mathbb{E}[S_j^{\underline{\alpha}} | B = \tau]. \quad (19)$$

Using Eq. (4), this set of equations (19) is readily seen to lead to the RLA defined in Eq. (5). In this way, the RLA can be seen as optimizing a lower bound for $\mathbb{E}[S_0^{\underline{\alpha}}]$.

On the contrary, the HTA can be viewed as optimizing an upper bound for $\mathbb{E}[S_0^{\underline{\alpha}}]$. To this end, note that the simulation results shown in Figure 3 suggest that for fixed $\underline{\alpha} \in A$ the per-task sojourn times are positively correlated. Loosely speaking, the maximum of positively correlated random variables is stochastically dominated by the maximum of independent random variables with identical marginal distributions. Hence, such a dominance also holds for the expected values, so that the HTA defined in Eq. (15) gives an upper bound for $\mathbb{E}[S_0^{\underline{\alpha}}]$, defined in Eq. (3). In that sense, the HTA is optimizing an upper bound.

Remark 3.2. Recall that the HT behavior in (11) holds if and only if $i \in U$, i.e., for those nodes i for which $\xi_i = \xi$, defined in (1). For $i \in U$, Eq. (11) shows that the conditional sojourn time for node i converges (both in distribution and moment-wise) to an exponential distribution with known mean. In contrast, the nodes $i \notin U$ do not become unstable for $\xi \uparrow 1$. However, note that in the HTA in (15) the marginal conditional sojourn-time distributions are approximated by (independent) exponential distributions for all $i = 1, \dots, N$. Nonetheless, note that under HT circumstances (i.e., $\xi \approx 1$) the impact of the per-task sojourn times of nodes i for $i \notin U$ on $\mathbb{E}[S_0^{\underline{\alpha}}]$ tends to vanish under HT-scalings.

Remark 3.3. The choice of the interpolation factor κ_i in Eq. (18) only depends on the background-load values. The benefit is its simplicity and the fact that κ_i does not depend on $\underline{\alpha}$, the parameter which is to be optimized. The drawback of this choice is that it does not accurately cover the HT behavior when $\xi \uparrow 1$ while $\rho_1, \dots, \rho_N$ are close to 0; this may happen when the foreground load ρ_0 is large. One way to overcome this problem is to take as the interpolating factor $\kappa_i := \xi_i$, defined in Eq. (1). The problem is that in this way, the interpolation factor itself depends on α_i which leads to a fixed-point equation to solve for α_i . We have checked the accuracy of the approximations based on $\kappa_i = \xi_i$; note that convergence of such fixed-point iteration can easily be shown to hold. Our results show that no significant improvement of the accuracy of the approximations is obtained.

4. NUMERICAL RESULTS

To assess the accuracy of the approximations for the optimal splitting rule $\underline{\alpha}$, we have performed extensive numerical experimentation, comparing the approximation results with simulations. To cover a wide

range of parameter combinations in a structured manner, we have varied the file-size distributions (deterministic, exponential, hyper-exponential, Pareto), and the load values of the foreground traffic (high, medium, low) and the background traffic (high, medium, low).

In our experiments various parameters scenarios were considered. For these parameter combinations, we calculated the following:

1. the optimal split rule $\underline{\alpha}^* = (\alpha_1^*, \dots, \alpha_{N-1}^*, 1 - \alpha_1^* - \dots - \alpha_{N-1}^*)$,
2. approximations for $\underline{\alpha}^*$, denoted

$$\underline{\alpha}_{app}^* = (\alpha_{app,1}^*, \dots, \alpha_{app,N-1}^*, 1 - \alpha_{app,1}^* - \dots - \alpha_{app,N-1}^*), \quad \text{and}$$

3. the relative difference in the mean foreground traffic processing times, defined by

$$\Delta\% := \text{abs} \left(\frac{\mathbb{E}[S_0^{\underline{\alpha}_{app}^*}] - \mathbb{E}[S_0^{\underline{\alpha}^*}]}{\mathbb{E}[S_0^{\underline{\alpha}^*}]} \right) \times 100\%. \quad (20)$$

First, we assume $N = 2$ and $\beta^{(1)} = 1$. The file-size distributions were varied as deterministic (to cover the case $c_B^2 = 0$), exponential ($c_B^2 = 1$), H_2 with $c_B^2 = 16$ (with balanced means) and Pareto-2 (with $c_B^2 = \infty$). The load of the foreground traffic ρ_0 was varied as 0.1, 0.5, 0.9, and 1.8, and the background loads ρ_1 and ρ_2 were varied as 0.1, 0.3, $\dots$, 0.9. To search for the optimal splitting rule $\underline{\alpha}^* = (\alpha, 1 - \alpha)$, we evaluated all feasible values of α with a step size 0.01, and more finely if needed. Below we will present the results of the evaluations. Tables 1 to 5 show for each feasible combination of ρ_1 and ρ_2 the corresponding values of the optimal split determined by simulation $\underline{\alpha}^* = (\alpha^*, 1 - \alpha^*)$, the approximated optimal split $\underline{\alpha}_{app}^* = (\alpha_{app}^*, 1 - \alpha_{app}^*)$ for $app \in \{RLA, HTA, CSA\}$, and the relative error in the cost function $\Delta\%$, defined in (20). To obtain highly accurate simulation results, experiments were run with extremely many jobs, up to 10^{10} if needed, leading to very narrow confidence intervals (CIs), such that all digits in the Tables 1 to 5 are significant. For compactness of the presentation the CIs are not shown. Also, note that because of the symmetry in the model for $N = 2$ only the results for $\rho_1 \leq \rho_2$ are shown.

To start, we compare the performance of the RLA (discussed in Section 3.2), the HTA (discussed in Section 3.3) and the CSA (discussed in Section 3.4). As an illustrative example, Table 1 shows the results for the case with $\rho_0 = 0.1$ and exponential job-size distributions, for a variety of background-load combinations of (ρ_1, ρ_2) .

We observe that the CSA indeed performs much better than both the RLA and the HTA. In fact, the RLA performs very well if $\rho_1 \approx \rho_2$, but tends to degrade significantly if ρ_1 and ρ_2 are far apart. This degradation

TABLE 1 Comparison of the RLA, the HTA and the CSA for the case of exponential job-size distributions and $\rho_0 = 0.1$, and $N = 2$

Reduced load approximation (RLA)															
$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{RLA}^*	$\Delta\%$	α^*	α_{RLA}^*	$\Delta\%$	α^*	α_{RLA}^*	$\Delta\%$	α^*	α_{RLA}^*	$\Delta\%$	α^*	α_{RLA}^*	$\Delta\%$
0.1	0.50	0.50	0.00	0.55	0.56	0.03	0.70	0.64	1.92	0.85	0.75	8.79	0.97	0.90	24.88
0.3				0.50	0.50	0.00	0.63	0.58	0.69	0.80	0.70	6.19	0.96	0.88	23.39
0.5							0.50	0.50	0.00	0.70	0.63	2.58	0.94	0.83	19.65
0.7										0.50	0.50	0.00	0.86	0.75	10.76
0.9													0.50	0.50	0.00
Heavy-traffic approximation (HTA)															
$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{HTA}^*	$\Delta\%$	α^*	α_{HTA}^*	$\Delta\%$	α^*	α_{HTA}^*	$\Delta\%$	α^*	α_{HTA}^*	$\Delta\%$	α^*	α_{HTA}^*	$\Delta\%$
0.1	0.50	0.50	0.00	0.55	0.63	2.29	0.70	0.78	2.79	0.85	0.91	3.01	0.97	0.99	1.71
0.3				0.50	0.50	0.00	0.63	0.67	0.66	0.80	0.85	1.50	0.96	0.98	1.54
0.5							0.50	0.50	0.00	0.70	0.73	0.32	0.94	0.96	0.93
0.7										0.50	0.50	0.00	0.86	0.87	0.12
0.9													0.50	0.50	0.00
Composed-split approximation (CSA)															
$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$
0.1	0.50	0.50	0.00	0.55	0.59	0.42	0.70	0.72	0.07	0.85	0.87	0.15	0.97	0.98	0.72
0.3				0.50	0.50	0.00	0.63	0.63	0.00	0.80	0.81	0.03	0.96	0.97	0.50
0.5							0.50	0.50	0.00	0.70	0.70	0.01	0.94	0.95	0.22
0.7										0.50	0.50	0.00	0.86	0.86	0.01
0.9													0.50	0.50	0.00

in the accuracy becomes most apparent if $\rho_1 \approx 0$ and $\rho_2 \approx 1$, showing double-digit error percentages. This was to be expected, since the RLA (5) simply splits traffic proportional to the *relative* amounts of capacity not used by the background traffic, while one may suspect that the *absolute* values of the background traffic have a large impact on the sensitivity of the choice of the splitting rule with respect to the background-load values. As expected, the HTA is doing much better in those asymmetric heavy-load scenarios, but tends to degrade somewhat when PS_1 is lightly loaded ($\rho_1 = 0.1$) and PS_2 is moderately loaded ($\rho_2 = 0.7$), leading to errors up to 3%. We observe that the CSA overall performs significantly better than the RLA and the HTA, in most cases leading to an error percentage less than 0.5%.

In short, the results in Table 1 indeed show that the usefulness of refining the simple and explicit RLA (introduced in Ref.^[12]) into the CSA, leading to extremely accurate results for most of the parameter combinations.

In the remainder of this section, we will further evaluate the accuracy of the CSA for a broad range of service-time distributions and combinations of foreground and background loads. Table 2 shows the results for light foreground load, with $\rho_0 = 0.1$. The results show that the approximations are highly accurate over a wide range of background-load combinations and service-time distributions, with errors significantly less than 1%. The least favorable results from our approximation were consistently found when the background load is highly asymmetric ($\rho_1 = 0.1$ and $\rho_2 = 0.9$), but even in those cases the results are highly accurate, with errors below 0.8%.

Tables 3 and 4 show the results for moderate foreground load values of $\rho_0 = 0.5$ and $\rho_0 = 0.9$, respectively. The results again show that the CSA

TABLE 2 Simulation results for light foreground load ($\rho_0 = 0.1$), $N = 2$

$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$
Deterministic															
0.1	0.50	0.50	0.00	0.55	0.59	0.43	0.71	0.72	0.03	0.86	0.87	0.10	0.97	0.98	0.70
0.3				0.50	0.50	0.00	0.63	0.63	0.01	0.80	0.81	0.03	0.96	0.97	0.49
0.5							0.50	0.50	0.00	0.71	0.70	0.00	0.94	0.95	0.05
0.7										0.50	0.50	0.00	0.86	0.86	0.03
0.9													0.50	0.50	0.00
Exponential															
0.1	0.50	0.50	0.00	0.55	0.59	0.42	0.70	0.72	0.07	0.85	0.87	0.15	0.97	0.98	0.72
0.3				0.50	0.50	0.00	0.63	0.63	0.00	0.80	0.81	0.03	0.96	0.97	0.50
0.5							0.50	0.50	0.00	0.70	0.70	0.01	0.94	0.95	0.22
0.7										0.50	0.50	0.00	0.86	0.86	0.01
0.9													0.50	0.50	0.00
Hyper-exponential															
0.1	0.50	0.50	0.00	0.55	0.59	0.76	0.70	0.72	0.22	0.85	0.87	0.12	0.97	0.98	0.76
0.3				0.50	0.50	0.00	0.63	0.63	0.07	0.80	0.81	0.04	0.96	0.97	0.57
0.5							0.50	0.50	0.00	0.70	0.70	0.02	0.94	0.95	0.23
0.7										0.50	0.50	0.00	0.86	0.86	0.13
0.9													0.50	0.50	0.00
Pareto															
0.1	0.50	0.50	0.00	0.55	0.59	0.52	0.70	0.72	0.15	0.85	0.87	0.23	0.97	0.98	0.73
0.3				0.50	0.50	0.00	0.62	0.63	0.03	0.79	0.81	0.01	0.96	0.97	0.57
0.5							0.50	0.50	0.00	0.70	0.70	0.01	0.94	0.95	0.25
0.7										0.50	0.50	0.00	0.86	0.86	0.03
0.9													0.50	0.50	0.00

TABLE 3 Simulation results for moderate foreground load ($\rho_0 = 0.5$), $N = 2$

$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$
Deterministic															
0.1	0.50	0.50	0.00	0.57	0.58	0.11	0.68	0.68	0.03	0.81	0.81	0.02	0.95	0.96	0.25
0.3				0.50	0.50	0.00	0.61	0.61	0.03	0.75	0.75	0.04	0.92	0.92	0.05
0.5							0.50	0.50	0.00	0.65	0.65	0.02	0.85	0.85	0.20
0.7										0.50	0.50	0.00			
Exponential															
0.1	0.50	0.50	0.00	0.57	0.58	0.13	0.68	0.68	0.04	0.81	0.81	0.01	0.95	0.96	0.28
0.3				0.50	0.50	0.00	0.61	0.61	0.03	0.75	0.75	0.06	0.92	0.92	0.06
0.5							0.50	0.50	0.00	0.65	0.65	0.04	0.85	0.85	0.22
0.7										0.50	0.50	0.00			
Hyper-exponential															
0.1	0.50	0.50	0.00	0.57	0.58	0.35	0.68	0.68	0.16	0.81	0.81	0.05	0.95	0.96	0.48
0.3				0.50	0.50	0.00	0.61	0.61	0.00	0.75	0.75	0.07	0.92	0.92	0.15
0.5							0.50	0.50	0.00	0.65	0.65	0.09	0.85	0.85	0.36
0.7										0.50	0.50	0.00			
Pareto															
0.1	0.50	0.50	0.00	0.57	0.58	0.19	0.67	0.68	0.05	0.81	0.81	0.04	0.95	0.96	0.22
0.3				0.50	0.50	0.00	0.60	0.61	0.01	0.74	0.75	0.00	0.92	0.92	0.18
0.5							0.50	0.50	0.00	0.65	0.65	0.07	0.85	0.85	0.06
0.7										0.50	0.50	0.00			

TABLE 4 Simulation results for moderate foreground load ($\rho_0 = 0.9$), $N = 2$

$\rho_1 \backslash \rho_2$	0.1			0.3			0.5			0.7			0.9		
	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$	α^*	α_{CSA}^*	$\Delta\%$
Deterministic															
0.1	0.50	0.50	0.00	0.57	0.57	0.02	0.66	0.66	0.02	0.77	0.77	0.06	0.91	0.91	0.63
0.3				0.50	0.50	0.00	0.59	0.59	0.02	0.71	0.71	0.38			
0.5							0.50	0.50	0.00						
Exponential															
0.1	0.50	0.50	0.00	0.57	0.57	0.07	0.66	0.66	0.04	0.77	0.77	0.17	0.91	0.91	0.11
0.3				0.50	0.50	0.00	0.59	0.59	0.05	0.71	0.71	0.06			
0.5							0.50	0.50	0.00						
Hyper-exponential															
0.1	0.50	0.50	0.00	0.57	0.57	0.11	0.66	0.66	0.09	0.77	0.77	0.12	0.91	0.91	0.07
0.3				0.50	0.50	0.00	0.59	0.59	0.01	0.71	0.71	0.36			
0.5							0.50	0.50	0.00						
Pareto															
0.1	0.50	0.50	0.00	0.57	0.57	0.08	0.66	0.66	0.05	0.77	0.77	0.00	0.91	0.91	0.00
0.3				0.50	0.50	0.00	0.59	0.59	0.00	0.71	0.71	0.53			
0.5							0.50	0.50	0.00						

performs extremely well in all cases considered, with errors significantly less than 1%, even for Pareto-2 distributed job sizes (thus with infinite variance) and highly asymmetric background-load scenarios. Note that in Tables 3 and 4 several combinations of (ρ_1, ρ_2) are omitted because they violate the stability conditions.

Finally, to assess the accuracy of the approximation for heavily loaded foreground traffic, additional simulation runs were conducted for $\rho_0 = 1.8$. Note that in this case, the set of α -values for which the system is still stable is limited to $\alpha \in (4/9; 1/2)$. To obtain an accurate estimate for α^* , we simulated $\mathbb{E}[S_0^{(\alpha, 1-\alpha)}]$ for different α -values with step size 0.001.

The results shown in Table 5 demonstrate that the CSA is also extremely accurate for heavy foreground load scenarios. Note that the mean sojourn times and the optimal split are indeed remarkably insensitive with respect to the job-size distributions, which supports the validity of Observation 1 in Section 3.

To assess the accuracy of the approximation for $N > 2$, we also consider a model with $N = 3$ with Pareto-2 distributed service times with $\beta^{(1)} = 1$. Table 6 shows the results for this system with $\rho_0 = 1.5$ where ρ_1 , ρ_2 , and ρ_3 are varied as 0.3, 0.5, and 0.7. For each feasible parameter setting, the table shows $\underline{\alpha}^* = (\alpha_1^*, \alpha_2^*, \alpha_3^*)^\top$ obtained via simulation, $\underline{\alpha}_{CSA}^* = (\alpha_{1CSA}^*, \alpha_{2CSA}^*, \alpha_{3CSA}^*)^\top$, and the relative error defined in Eq. (20). The results in Table 6 show that the accuracy of the CSA is again excellent for $N = 3$. To summarize, the results in Tables 1 to 6 show that the CSA, which combines the benefits of the RLA and the HTA, leads to extremely accurate approximations for $\underline{\alpha}^*$ over a wide range parameter settings.

Remark 4.1. In Section 3.1 we observed on the basis of preliminary simulation experiments that the mean sojourn times, and the optimal $\underline{\alpha}$ -values, are remarkably insensitive with respect to the service-time distributions (Observation 1), even for extremely variable job-size distributions. In this context, notice that the results shown in Tables 1 to 5 unanimously confirm this observation.

TABLE 5 Simulation results for heavy foreground load ($\rho_0 = 1.8$), with $N = 2$, $\rho_1 = 0.1$ and $\rho_2 = 0.0$

	α^*	α_{CSA}^*	$\Delta\%$
Deterministic	0.473	0.47312	0.03
Exponential	0.473	0.47312	0.06
Hyper-exponential	0.474	0.47312	0.05
Pareto	0.474	0.47312	0.06

TABLE 6 Simulation results for moderate foreground load ($\rho_0 = 1.5$), $N = 3$

$\rho_1 \backslash \rho_2$	0.1				0.3				0.5				0.7				0.9			
	$\underline{\alpha}^*$	$\underline{\alpha}_{CSA}^*$	$\Delta\%$		$\underline{\alpha}^*$	$\underline{\alpha}_{CSA}^*$	$\Delta\%$		$\underline{\alpha}^*$	$\underline{\alpha}_{CSA}^*$	$\Delta\%$		$\underline{\alpha}^*$	$\underline{\alpha}_{CSA}^*$	$\Delta\%$		$\underline{\alpha}^*$	$\underline{\alpha}_{CSA}^*$	$\Delta\%$	
0.1	0.366	0.362			0.399	0.394			0.437	0.435			0.484	0.482			0.536	0.535		
	0.366	0.362	0.39%		0.301	0.303	0.13%		0.230	0.233	0.42%		0.144	0.150	0.89%		0.051	0.053	0.47%	
	0.268	0.276			0.300	0.303			0.333	0.332			0.372	0.368			0.413	0.412		
0.3					0.333	0.334			0.370	0.370			0.415	0.414						
					0.333	0.333	0.00%		0.260	0.261	0.27%		0.170	0.172	0.34%				unstable	
					0.334	0.333			0.370	0.369			0.415	0.414						
0.5									0.292	0.293	0.24%			unstable					unstable	
									0.292	0.293										
									0.416	0.414										
0.1									$\rho_3 = 0.5$											
	0.400	0.396			0.437	0.435			0.485	0.480			0.538	0.535						
	0.400	0.395	1.03%		0.333	0.332	0.42%		0.257	0.260	0.40%		0.170	0.172	0.50%				unstable	
0.3					0.230	0.233			0.258	0.260			0.292	0.293						
					0.370	0.370			0.416	0.414										
					0.370	0.369	0.27%		0.292	0.293	0.24%			unstable					unstable	
0.1									$\rho_3 = 0.7$											
	0.440	0.435			0.484	0.482			0.538	0.535										
	0.440	0.434	1.60%		0.372	0.368	0.89%		0.292	0.293	0.50%			unstable					unstable	
0.3					0.144	0.150			0.170	0.172										
					0.415	0.414														
					0.415	0.414	0.34%							unstable					unstable	

Remark 4.2. Despite the fact the RLA for $\underline{\alpha}^*$ in (5) is explicit, the HTA is not, and can only be calculated numerically. This generally requires numerical optimization of $\underline{\alpha}$ over the set A . When N is not too large this causes no problem, because the evaluation of $\mathbb{E}[S_{HTA,0}^{\underline{\alpha}}]$ for given $\underline{\alpha}$ is explicit, and the set of $\underline{\alpha}$ -values is within the bounded set A for which standard non-linear optimization techniques are available. Note that discretization of the set A and then enumeration over all $\underline{\alpha}$ can also be done for small N . However, when N is large, the computation times may become significant. In those cases, a further simplification of the HTA seems to be needed. In this context, recall that in general not all queues become unstable as $\xi \uparrow 1$ (namely only those for which $i \in U$, see also Remark 3.2). Therefore, only the proper choice of α_i for $i \in U$ may be crucial, whereas the cost function $\mathbb{E}[S_0^{\underline{\alpha}}]$ may be expected to be fairly insensitive to the choice of α_i for $i \notin U$ as $\xi \uparrow 1$, i.e., becomes negligible in heavy traffic. This observation may lead to a dramatic reduction of the dimension of the optimization problem. Furthermore, for $N \rightarrow \infty$ one may use asymptotic results from the powerful extreme-value theory to develop approximations for the HT behavior of $\mathbb{E}[S_0^{\underline{\alpha}}]$. These observations open up possibilities for further reducing the computational complexity of the calculation of $\underline{\alpha}_{HTA}^*$. Finally, note that from an application point of view, in the context of wireless networks with concurrent access (a) N is rather small, say $N \leq 3$, and (b) the job-split ratio does not have to be (re)calculated in real time, so that the computational requirements are not very strict.

5. TOPICS FOR FURTHER RESEARCH

The results presented in this article raise a number of challenging topics for further research. First, the simulation results in Figures 2 and 3 show that $\mathbb{E}[S_0^{\underline{\alpha}}]$ is at least near-insensitive with respect to the distribution of the job size B . The question arises whether $\mathbb{E}[S_0^{\underline{\alpha}}]$ is fully insensitive to the distribution of B , similar to the marginal distributions of the per-task sojourn times (see Eq. (4)). Even extremely long simulation runs do not give a definite answer whether $\mathbb{E}[S_0^{\underline{\alpha}}]$ is fully insensitive to the distribution of B . Obtaining rigorous proofs of insensitivity properties, possibly under additional requirements on the job-size distribution, is a challenging topic for further research. Second, an intriguing observation made in Figures 3(a) and (b) is that the differences in correlations over different job-size distributions are evident but seem to have no impact on $\mathbb{E}[S_0^{\underline{\alpha}}]$. The impact of the correlations between the per-task sojourn times seems to cancel out when evaluating $\mathbb{E}[S_0^{\underline{\alpha}}]$. Currently, a full understanding of this phenomenon is lacking, and is left as a topic for follow-up research. Third, the Poisson assumption may be relaxed. In fact, we suspect that both

the RLA and the HTA can be quite easily extended for example to renewal arrival processes, but it is unclear to what extent they can be further generalized to more realistic arrival processes that include correlations between job arrivals. Derivation of such approximations is a challenging subject for further research. Finally, the job-split model assumed a fixed splitting rule $\underline{\alpha}$. However, it is likely that much better performance can be obtained if the splitting ratios are allowed to depend on the actual state of the system. Derivation and evaluation of such dynamic splitting policies address another appealing topic for follow-up research.

APPENDIX A

For self-containedness of the article, in this appendix we formulate a known result giving the expected value of maximum of a arbitrary number of exponentially distributed random variables.

Property 1. *If $X_1, \dots, X_N$ are i.i.d. exponentially distributed random variables with means $1/\mu_1, \dots, 1/\mu_N$ respectively, then*

$$E[\max\{X_1, \dots, X_N\}] = \sum_{k=1}^N (-1)^{k+1} \sum_{(i_1, \dots, i_k) \in S_k} \frac{1}{\mu_{i_1} + \dots + \mu_{i_k}}, \quad (21)$$

where S_k ($k = 1, \dots, N$) is defined in (16).

Proof. The derivation of Property 1 is trivial and requires only standard algebraic manipulations, and is left as an exercise to the reader. $\square$

To illustrate Property 1, let us work out (21) for the cases $N = 2, 3$ and 4. For $N = 2$, we have $S = \{1, 2\}$, $S_1 = \{1, 2\}$, and $S_2 = \{(1, 2)\}$, so that

$$\mathbb{E}[\max\{X_1, X_2\}] = \frac{1}{\mu_1} + \frac{1}{\mu_2} - \frac{1}{\mu_1 + \mu_2}. \quad (22)$$

For $N = 3$, we have $S = \{1, 2, 3\}$, $S_1 = \{1, 2, 3\}$, $S_2 = \{(1, 2), (1, 3), (2, 3)\}$ and $S_3 = \{(1, 2, 3)\}$, so that

$$\begin{aligned} \mathbb{E}[\max\{X_1, X_2, X_3\}] &= \left(\frac{1}{\mu_1} + \frac{1}{\mu_2} + \frac{1}{\mu_3} \right) - \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\mu_1 + \mu_3} + \frac{1}{\mu_2 + \mu_3} \right) \\ &\quad + \frac{1}{\mu_1 + \mu_2 + \mu_3}. \end{aligned} \quad (23)$$

Finally, it is readily verified that for $N = 4$ we have $S = \{1, 2, 3, 4\}$, $S_1 = \{1, 2, 3, 4\}$, $S_2 = \{(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)\}$, $S_3 = \{(1, 2, 3), (1, 2, 4), (1, 3, 4), (2, 3, 4)\}$, $S_4 = \{(1, 2, 3, 4)\}$, so that

$$\begin{aligned} & \mathbb{E}[\max\{X_1, X_2, X_3, X_4\}] \\ &= \left(\frac{1}{\mu_1} + \frac{1}{\mu_2} + \frac{1}{\mu_3} + \frac{1}{\mu_4} \right) \\ & \quad - \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\mu_1 + \mu_3} + \frac{1}{\mu_1 + \mu_4} + \frac{1}{\mu_2 + \mu_3} + \frac{1}{\mu_2 + \mu_4} + \frac{1}{\mu_3 + \mu_4} \right) \\ & \quad + \left(\frac{1}{\mu_1 + \mu_2 + \mu_3} + \frac{1}{\mu_1 + \mu_2 + \mu_4} + \frac{1}{\mu_1 + \mu_3 + \mu_4} + \frac{1}{\mu_2 + \mu_3 + \mu_4} \right) \\ & \quad - \frac{1}{\mu_1 + \mu_2 + \mu_3 + \mu_4}. \end{aligned} \quad (24)$$

ACKNOWLEDGMENTS

The authors thank Ronny Gunawan for his initial contributions. The work reported in this article was supported by the Netherlands Organisation for Scientific Research (NWO) under the Casimir project: Analysis of Distribution Strategies for Concurrent Access in Wireless Communication Networks.

REFERENCES

1. IEEE Standard 802.11n. Part 11: Wireless LAN medium access control (MAC) and physical layer specifications enhancements for higher throughput. October 2009.
2. Altman, E.; Ayesta, U.; Prabhu, B. Load balancing in processor sharing systems. In *Proceedings of the Second International Workshop on Game Theory in Communication Networks 2008, GameComm 2008, October 20, 2008, Athens Greece*. HAL—CCSD, October 2008.
3. Baccelli, F.; Massey, W.A.; Towsley, D. Acyclic fork-join queueing networks. *Journal of the ACM* **1989**, *36* (3), 615–642.
4. Borst, S.C.; Boxma, O.J.; Hegde, N. Sojourn times in finite-capacity processor-sharing queues. In *Proceedings NGI 2005 Conference*, 2005.
5. Chandra, R.; Bahl, P.; Bahl, P. Multinet: Connecting to multiple IEEE 802.11 networks using a single wireless card. In *Proceedings of IEEE INFOCOM*, 2004.
6. Cox, D. Fundamental limitations on the data rate in wireless systems. *IEEE Communications Magazine* **2008**, *46* (12), 16–17.
7. Duncanson, J. Inverse multiplexing. *IEEE Communications Magazine* **1994**, *32* (4), 34–41.
8. FCC. Report of the spectrum efficiency working group. Technical report, Federal Communications Commission Spectrum Policy Task Force, November 2002.
9. Gkantsidis, C.; Ammar, M.; Zegura, E. On the effect of large-scale deployment of parallel downloading. In *WIAPP '03: Proceedings of the Third IEEE Workshop on Internet Applications*, IEEE Computer Society: Washington, DC, 2003, p. 79.
10. Gupta, V.; Balter, M.H.; Sigman, K.; Whitt, W. Analysis of join-the-shortest-queue routing for web server farms. *Performance Evaluation* **2007**, *64* (9–12), 1062–1081.
11. Hasegawa, Y.; Yamaguchi, I.; Hama, T.; Shimonishi, H.; Murase, T. Deployable multipath communication scheme with sufficient performance data distribution method. *Computer Communications* **2007**, *30* (17), 3285–3292.

12. Hoekstra, G.; van der Mei, R.; Nazarathy, Y.; Zwart, B. Optimal file splitting for wireless networks with concurrent access. In *Network Control and Optimization, Third Euro-NF Conference, NET-COOP 2009, Eindhoven, The Netherlands, November 23–25, 2009*, volume 5894 of Lecture Notes in Computer Science. Springer Verlag: Berlin, 2009, pp. 189–203.
13. Hoekstra, G.J.; Panken, F.J.M. Increasing throughput of data applications on heterogeneous wireless access networks. In *Proceedings 12th IEEE Symposium on Communication and Vehicular Technology in the Benelux*, 2005.
14. Hoekstra, G.J.; van der Mei, R.D. Effective load for flow-level performance modelling of file transfers in wireless LANs. *Computer Communications* **2010**, *33* (16), 1972–1981.
15. Hsieh, H.Y.; Sivakumar, R. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. *Wireless Networks* January **2005**, *11* (1), 99–114.
16. Koudouris, G.P.; Agüero, R.; Alexandri, E.; Choque, J.; Dimou, K.; Karimi, H.R.; Lederer, H.; Sachs, J.; Sigle, R. Generic link layer functionality for multi-radio access networks. In *Proceedings 14th IST Mobile and Wireless Communications Summit*, 2005.
17. Lelarge, M. Tail asymptotics for discrete event systems. In *valuetools '06: Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*; New York: ACM, 2006, p. 36.
18. Lelarge, M. Packet reordering in networks with heavy-tailed delays. *Mathematical Methods of Operations Research* **2008**, *67* (2), 341–371.
19. Litjens, R.; Roijers, F.; Van den Berg, J.L.; Boucherie, R.J.; Fleuren, M.J. Performance analysis of wireless LANs: An integrated packet/flow level approach. In *Proceedings of the 18th International Teletraffic Congress—ITC18*; Berlin, Germany, 2003, pp. 931–940.
20. Rodriguez, P.; Kirpal, A.; Biersack, E. Parallel-access for mirror sites in the internet. In *INFOCOM*; 2000, pp. 864–873.
21. Wu, Y.; Williamson, C.; Luo, J. On processor sharing and its applications to cellular data network provisioning. *Performance Evaluation* **2007**, *64* (9–12), 892–908.
22. Zwart, A.P.; Boxma, O.J. Sojourn time asymptotics in the M/G/1 processor sharing queue. *Queueing Systems: Theory and Applications* **2000**, *35* (1/4), 141–166.