



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms



Probability, Networks and Algorithms

Performance analysis of differentiated resource-sharing
in a wireless ad-hoc network

F. Roijers, J.L. van den Berg, M.R.H. Mandjes

REPORT PNA-E0807 AUGUST 2008

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO). CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2008, Stichting Centrum voor Wiskunde en Informatica
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

Performance analysis of differentiated resource-sharing in a wireless ad-hoc network

ABSTRACT

In this paper we model and analyze a relay node in a wireless ad-hoc network; the capacity available at this node is used to both transmit traffic from the source nodes (towards the relay node), and to serve traffic at the relay node (so that it can be forwarded to successor nodes). Clearly, when a specific node is used more heavily than others, it is prone to becoming a performance bottleneck. In this paper we consider the situation that the relay node obtains a share of the capacity that is m times as large as the share that each source node receives. The main performance metrics considered are the workload at the relay node and the average overall flow transfer time, i.e., the average time required to transmit a flow from a source node via the relay node to the destination. Our aim is to find expressions for these performance metrics for a general resource-sharing ratio m , as well as a general flow-size distribution. The analysis consists of the following steps. First, for the special case of exponential flow sizes we analyze the source-node dynamics, as well as the workload at the relay node by a fluid-flow queueing model. Then we observe from extensive numerical experimentation over a broad set of parameter values that the distribution of the number of active source nodes is actually insensitive to the flow-size distribution. Using this remarkable (empirical) result as an approximation assumption, we obtain explicit expressions for both the mean workload at the relay node and the overall flow transfer time, both for general flow-size distributions.

2000 Mathematics Subject Classification: 60K25, 90B18

Keywords and Phrases: resource sharing; bottleneck; feedback fluid-queue; insensitivity; transfer time

Performance analysis of differentiated resource-sharing in a wireless ad-hoc network

Frank Roijers^{a,b} Hans van den Berg^{a,c} Michel Mandjes^{b,d}

^aTNO Information and Communication Technology, The Netherlands

^bUniversity of Amsterdam, Korteweg-de Vries Institute for Mathematics, The Netherlands

^cUniversity of Twente, Department of Design and Analysis of Communication Systems (DACS), The Netherlands

^dCenter for Mathematics and Computer Science (CWI), The Netherlands

Abstract

In this paper we model and analyze a relay node in a wireless ad-hoc network; the capacity available at this node is used to both transmit traffic from the source nodes (towards the relay node), and to serve traffic at the relay node (so that it can be forwarded to successor nodes). Clearly, when a specific node is used more heavily than others, it is prone to becoming a performance bottleneck. In this paper we consider the situation that the relay node obtains a share of the capacity that is m times as large as the share that each source node receives.

The main performance metrics considered are the workload at the relay node and the average overall flow transfer time, i.e., the average time required to transmit a flow from a source node via the relay node to the destination. Our aim is to find expressions for these performance metrics for a general resource-sharing ratio m , as well as a general flow-size distribution.

The analysis consists of the following steps. First, for the special case of exponential flow sizes we analyze the source-node dynamics, as well as the workload at the relay node by a fluid-flow queueing model. Then we observe from extensive numerical experimentation over a broad set of parameter values that the distribution of the number of active source nodes is actually *insensitive* to the flow-size distribution. Using this remarkable (empirical) result as an approximation assumption, we obtain explicit expressions for both the mean workload at the relay node and the overall flow transfer time, both for general flow-size distributions.

Key words: resource sharing; bottleneck; feedback fluid-queue; insensitivity; transfer time

1 Introduction

Recent developments in wireless communication technology have opened up the possibility of operating *ad-hoc networks*. These networks have the key feature that they can be deployed instantly without a fixed infrastructure or pre-advanced configuration. The underlying communication technology is usually based on shared medium access (e.g., IEEE 802.11), i.e., neighboring nodes share the same radio capacity. In more detail, considering a given relay node, this means that the capacity available at this node is used both to transmit traffic from the source nodes (towards the relay node), and to serve traffic at the relay node (so that it can be forwarded to successor nodes). In the standard protocol every source node gets the same share of the capacity as the relay node. This explains why in case a certain node is used as a relay node more heavily than other nodes, this node is likely to become a performance bottleneck.

The present paper investigates a family of resource-sharing policies that grant a larger share of the available capacity to the relay node. We do so by considering a *fluid-flow model* of a single relay node that is fed by multiple source nodes. Source nodes become active at random time instances and start transmitting flows (where individual packets are approximated by fluid) to a destination via the relay node; after a source node has transmitted its flow to the relay node it becomes inactive. The ‘resource-sharing ratio’ m indicates the share of the overall capacity c that the relay node obtains relative to the share that is allocated to each of the source nodes. More precisely, if n source nodes are active, then the relay node obtains service rate $mc/(m+n)$ while each source node receives service rate $c/(m+n)$. If the aggregate rate of traffic flowing into the relay node exceeds the service rate of the relay node, work is backlogged at the relay node in a buffer (of infinite size). It is stressed that the relay node can only claim its full share $c/(m+n)$ if either the number of active source nodes exceeds m (that is, $n \geq m$), or if the relay node is backlogged. Otherwise the relay node is allocated a share $c/2$, while each source node obtains $c/(2n)$ (and hence the buffer remains empty). An important consequence of these allocation rules is that the system is *work-conserving*.

A special case of the above-mentioned model is the standard situation in which the relay node gets the same share as each source node, i.e., $m = 1$, which was treated in detail in [3, 17]. The analysis in these papers relies heavily on the fact that for $m = 1$ the resource sharing only depends on the number of active source nodes, and *not* on the current buffer content; this is in contrast with the case of general m , where knowledge of both the number of active sources as well as the current level are needed in order to determine the rate allocation, as pointed out above. It turned out that for $m = 1$ the behavior of the source nodes can be related to a $M/G/1$

Processor Sharing-queue with state-dependent service-rates; it is well-known that this model exhibits insensitivity to the flow-size distribution (apart from its mean). In [3, 17] the authors derive explicit expressions for the workload at the relay node and the overall flow transfer time, i.e., the time required to send a flow from a source node via the relay node to a destination.

1.1 Contribution

This paper presents analytical expressions for a number of performance metrics of the fluid model for general $m \in (0, \infty)$. It is stressed once more that this general case is significantly more difficult than the special case $m = 1$ of [3, 17] due to the fact that the resource sharing between the source nodes and the relay node is influenced by the workload at the relay node, and it is not solely determined by the number of active source nodes as is the case for $m = 1$.

In our analysis we first address the case of exponentially distributed flow-sizes. As the resource sharing is influenced by the workload at the relay node, the number of active source nodes itself does not constitute a Markov chain. However, the model falls within the class of so-called ‘fluid queues with feedback’ (or: ‘feedback fluid-queues’), as introduced in e.g., [16]. Using the theory developed there, we derive the joint distribution of the number of active source nodes and the workload at the relay node. In addition, we provide an analysis of the overall flow transfer time.

For the case of generally distributed flow-sizes we cannot rely on the framework developed in [16], as the exponentiality assumptions imposed there are crucial. For generally distributed flow sizes, however, we empirically observed the remarkable property that the distribution of the number of active source nodes is *insensitive* to the flow-size distribution (apart from its mean) for *any* general resource-sharing ratio $m \in (0, \infty)$ (i.e., not just for $m = 1$). This insensitivity claim is supported by a sizeable set of simulation experiments, corresponding to a broad range of parameters settings, including combinations of heavy load, a highly variable flow-size distribution, and a large resource-sharing ratio m . We consider the observed insensitivity as a highly surprising fact, since (for $m > 1$) the process describing the number of active source nodes does not fulfill the usual criteria for insensitivity (where we remark that for $m \in [0, 1]$ the insensitivity can be proven in a relatively elementary way).

By using the (conjectured) insensitivity as an approximation assumption, and by relating the workload of the total system to the workload in an appropriate M/G/1-queue, we derive insightful closed-form expressions for the mean workload at the relay node. In particular, we show that the mean workload at the relay node for gen-

eral flow-size distributions is proportional to the mean workload for exponentially distributed flow-sizes; here the associated multiplicative factor is independent of the system parameters (including m), apart from the Coefficient of Variation of the flow-size distribution. In addition, we present an approximation for the overall flow transfer time for general flow-size distributions. The resulting expressions allow for easy numerical evaluation and are thoroughly validated by simulations of the fluid model.

1.2 Literature

Notably, just a rather limited number of studies consider the flow-level analysis of ad-hoc networks, focusing on metrics such as the transfer time of a flow (typically consisting of multiple data packets). These metrics are particularly interesting for elastic flows, as for this type of traffic they relate more naturally to the perceived quality of end-users than packet-level statistics (such as packet delays). For flow-level performance studies of a single-hop ad-hoc network, we refer to e.g., [9, 15]. More specifically, in [15] a closed-form expression for the mean flow-transfer time is obtained by considering the system as a Processor Sharing queue with state-dependent service rates.

As mentioned before, [3, 17] study the present fluid model for the special case that $m = 1$. The focus of [20] was on the validation of the fluid model for $m \geq 1$; by system simulations incorporating all details of the IEEE 802.11e wireless LAN technology (see, e.g., [11]) it was demonstrated that the fluid model accurately captures the resource sharing amongst source nodes and a common relay node. In [2] the fluid model for $m = 1$ is analyzed in the special case of regularly varying (that is, *heavy-tailed*) flows. The tail asymptotics of the overall flow transfer time are derived by sample-path arguments; it is proven that the tail behaves roughly as the residual flow size. In [19] a versatile infinite-state Markov reward model is proposed to investigate the impact of different resource-sharing strategies for exponentially distributed flow-sizes. The authors numerically compute the performance measures such as the distribution of the number of active source nodes and the workload at the relay node; they do not study flow-level performance measures, e.g., flow-transfer times.

1.3 Outline

The outline of this paper is as follows. Section 2 describes our model in full detail and introduces the performance metrics that are analyzed in this paper. In addition,

it presents a number of preliminary results on the *total* workload in the ad-hoc network, which are exploited later. The model is analyzed in Section 3 for the special case of exponentially distributed flow-sizes and in Section 4 for generally distributed flow-sizes.

For the case of exponentially distributed flow-sizes we derive the joint distribution of the number of active source nodes and the buffer workload in Section 3.1, relying on the concept of fluid queues with feedback. Furthermore, we obtain the mean workload at the relay node (Section 3.2) and the mean overall flow transfer time (Section 3.3) to transmit a flow from source node to a destination via the relay node. In Section 4 the analysis is extended to the case of generally distributed flow-sizes. A crucial step in this analysis is played by our claim that the source-node behavior is insensitive to the flow-size distribution apart from its mean for general m , see Section 4.1. This result is then used to obtain the mean workload (Section 4.2) at the relay node and the mean overall flow transfer time (Section 4.3). Section 4.4 presents a calculation scheme that is useful to evaluate all the expressions involved.

Section 5 presents an extensive numerical validation of the analyses of the earlier sections. The scenarios considered are introduced in Section 5.1. Numerical results, validating the performance formulae for exponentially distributed flows-sizes, are given in Section 5.2. For generally distributed flow-sizes we present empirical evidence of the insensitivity claim in Section 5.3 (for a broad set of parameters, including extreme situations), as well as validation of the performance formulae in Section 5.4. Finally, Section 6 concludes this paper.

2 Model and performance metrics

Section 2.1 describes the model in more detail. Section 2.2 presents some preliminary analysis, in particular a relation between the total workload in the system and the workload in an $M/G/1$ -queue. Section 2.3 introduces the performance metrics.

2.1 Fluid model

We consider a network with a large number of source nodes which may become active and transfer flows of data (files) to destinations via a common relay node, see Figure 1. Flow transfers are initiated according to a Poisson process with rate λ ('flow arrival rate'). Flow sizes (in terms of fluid or bits) are distributed according to distribution F with mean f , second moment f_2 (assumed to be finite), and Coefficient of Variation (cov) c_F , i.e., $c_F^2 := \text{Var}(F) / (\mathbb{E}F)^2 = f_2 / f^2 - 1$. Let N_t denote




Fig. 1. Network model.

the number of active source nodes at time t and W_t the workload at the relay node at time t .

The total transmission capacity of the system is denoted by c and it is shared among the active source nodes and the relay node. If the aggregate rate of traffic flowing from the source nodes to the relay node exceeds the rate out of the relay node, traffic is stored in a buffer of infinite size which is served by the relay node in a FCFS-fashion.

The resource-sharing ratio between the relay node and the source nodes is denoted by m , i.e., the relay node obtains capacity $mc/(N_t + m)$ if $N_t \geq m$ sources are present or if the buffer of the relay node is backlogged, i.e., $W_t > 0$. Otherwise, the relay node and the set of active source nodes each obtain half the capacity ($c/2$), i.e., the aggregate input rate at the relay node is equal to the output rate. The source nodes always equally share the capacity not used by the relay node. Observe that the entire capacity c is always used if there is work in the system, so that the system is work-conserving. The resource sharing between the source nodes and the relay node is summarized in Table 1. The column ‘drift’ indicates the sign of the *net* input rate into the buffer at the relay node, i.e., it indicates whether the buffer content increases (+), decreases (–) or remains constant (0). Notice from Table 1 that the resource sharing at epoch t depends on both N_t and W_t . Figure 2 presents a sample-path example of the resource sharing in case $m = 2$.

Table 1

Resource sharing between source nodes and relay node.

Number of active sources	$W_t = 0$			$W_t > 0$		
	source	relay	drift	source	relay	drift
$N_t < m$	$c/2N_t$	$c/2$	0	$c/(m + N_t)$	$mc/(m + N_t)$	–
$N_t = m$	$c/2N_t$	$c/2$	0	$c/(m + N_t)$	$mc/(m + N_t)$	0
$N_t > m$	NA	NA	NA	$c/(m + N_t)$	$mc/(m + N_t)$	+

Observe that in our model a flow may be present (and receive service) at both the source node and the relay node: at flow initiation the source node immediately starts transmitting fluid to the relay node and parts of the flow may be present at

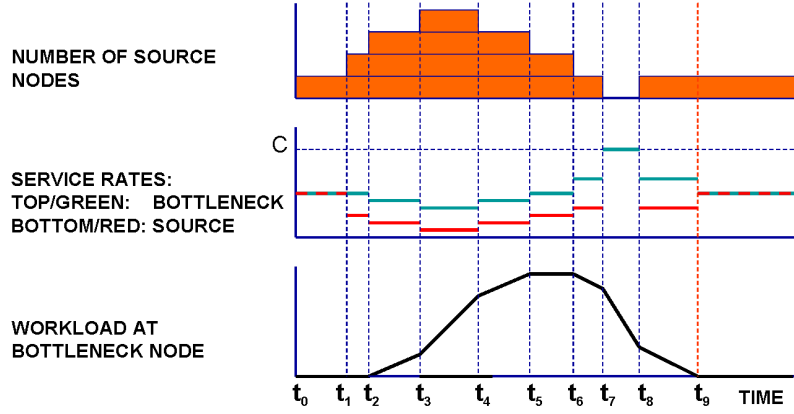


Fig. 2. Example of resource sharing for allocation $m = 2$.

both source and relay node. At some point in time the source node transmits the last (infinitesimally small) ‘particle’ of the flow to the relay node, and then the source node becomes inactive; this epoch is referred to as the ‘arrival of the last particle at the relay node’-epoch, or alternatively as the ‘source-departure’-epoch. In case the relay node is not backlogged, this last particle will be served instantly and the flow transfer is completed. If the relay node is backlogged, the last particle suffers an additional delay before the flow transfer time from source node to destination is completed.

2.2 Total workload of the system

The total workload W_{total} in our model, i.e., the sum of the workloads at the source nodes and the relay node, can be related to the workload in an $M/G/1$ -queueing system by making the following observations:

- (1) Each flow is served essentially *twice*; once by the source node to transmit the flow to the relay node, and once by the relay node to forward the flow to the destination node.
- (2) The overall system is *work-conserving*, i.e., the total service rate of the system is always c if there is work in the system.

As an aside, we remark that these observations immediately lead to the stability criterion $2\rho < 1$, with $\rho := f/c$.

Now it easily seen that the following lemma holds.

Lemma 1 $\mathbb{E}W_{\text{total}}$ corresponds to the virtual waiting-time in an $M/G/1$ -queue with service requirement distribution F with first moment $2f/c$ and second moment

$4f_2/c^2$. Then,

$$\mathbb{E}W_{\text{total}} = \left(\frac{2\rho}{1-2\rho} \right) \cdot \frac{1}{2}(c_F^2 + 1) \frac{2f}{c} = \left(\frac{2\rho}{1-2\rho} \right) \cdot (c_F^2 + 1) \frac{f}{c}. \quad (1)$$

Proof The expression for $\mathbb{E}W_{\text{total}}$ follows directly from the Pollaczek-Khintchine mean value formula. ■

2.3 Performance metrics

The main performance metrics of interest are the steady-state buffer workload W_{buffer} at the relay node and the overall flow transfer time D_{overall} , i.e., the time required to completely transfer a flow from source node to destination. The overall flow transfer time is the sum of two other performance metrics: (i) the time (D_{source}) a source requires to completely transfer a flow to the relay node, and (ii) the delay at the relay node (D_{buffer}^*) experienced by the last particle of the flow (the asterisk indicates that the performance measure relates to the last particle of a flow).

3 Analysis for exponentially distributed flow-sizes

This section presents an analysis of the model for the special case of exponentially distributed flow-sizes, which allows for a detailed analysis of the mean performance metrics. First, in Section 3.1 we derive the joint distribution of the number of active source nodes and the workload at the relay node at an arbitrary epoch. Next, this result is used in Section 3.2 to determine the workload at the relay node at the arrival epoch of the last particle of a flow. Finally, based on the results obtained in Sections 3.1 and 3.2, we present in Section 3.3 an approximation for the mean delay $\mathbb{E}D_{\text{buffer}}^*$ of the last particle of a flow that added to with the mean delay $\mathbb{E}D_{\text{source}}$ yields the mean overall flow transfer time $\mathbb{E}D_{\text{overall}}$.

3.1 Joint distribution of the number of source nodes and the buffer workload

The source-node dynamics of the model of Section 2.1 does not constitute a Markov chain; the transition rates depend on both the number of active source nodes N_t and on the buffer workload W_t , or, more precisely, whether W_t is positive or not. Put differently, there is *feedback* from the buffer content to the source behavior, in that information on the buffer content is needed to describe the source-node dynamics.

Thus the stationary joint distribution of the number of source nodes and the buffer workload can be obtained by modeling it as a so-called *feedback fluid-queue*, analogously to [16]. This distribution is obtained in two steps: first we consider the simplified fluid model *without* feedback (i.e., the source nodes evolve independently of the buffer content), and then the obtained results are translated in terms of the model *with* feedback. An extensive treatment of this derivation is presented in Appendix A; below we restrict ourselves to an excerpt.

First we consider the fluid-queue *without feedback*, i.e., the relay node always obtains share m of the capacity; let \bar{N}_t and \bar{W}_t respectively denote the number of active source nodes and buffer workload at epoch t in this model (in the sequel all quantities with bars ‘ $\bar{\cdot}$ ’ relate to the model without feedback). Observe that \bar{N}_t constitutes a Markov chain and has generator matrix \bar{Q} (as given in (A.1) in the Appendix). Further, let \bar{R} be a diagonal matrix where \bar{R}_n is the *net* input rate into the relay node if n sources are active, i.e., $\bar{R}_n := (n - m)c/(n + m)$. We define the stationary distribution of (\bar{N}_t, \bar{W}_t) as

$$F_n(x) := \lim_{t \rightarrow \infty} \mathbb{P}(\bar{W}_t \leq x; \bar{N}_t = n) = \mathbb{P}(\bar{W} \leq x; \bar{N} = n).$$

To facilitate the analysis we here assume that a maximum n_{max} is imposed on the number of source nodes that may be simultaneously active; flows that are initiated if already n_{max} other source nodes are active are blocked. Observe however that if one chooses n_{max} sufficiently large (which we will do in the sequel), this will not have any significant impact on the model.

The buffer workload has to satisfy the Kolmogorov forward equations

$$\vec{F}'(x)\bar{R} = \vec{F}(x)\bar{Q}'.$$

The spectral expansion of the solution is given by

$$\vec{F}(x) = \sum_{j=0}^{n_{max}} a_j \vec{v}_j \exp(z_j x)$$

where (z_j, \vec{v}_j) is an eigenvalue-eigenvector pair, i.e., a scalar and vector that solve $z_j \vec{v}_j \bar{R} = \vec{v}_j \bar{Q}$. For details on how to obtain the coefficients a_j see Appendix A. Clearly, $\mathbb{P}(\bar{W} \leq x) = \sum_n F_n(x)$. For stability we assume that the average net input rate is negative, i.e., $\sum_n \omega_n \bar{R}_n < 0$, where ω denotes the distribution of \bar{N} , i.e., $\omega_n := F_n(\infty)$.

The joint distribution of the workload and number of sources present for the fluid model *with* feedback follows from the corresponding distribution of the *without* feedback, by the crucial observation that both models behave identically *during*

busy periods, see [16]. Then the joint distribution $G_n(x)$ of (N, W) for the model with feedback is found after some sort of rescaling the distribution $F_n(x)$ of (\bar{N}, \bar{W}) :

$$G_n(x) = \frac{F_n(x) - \sum_k F_k(0)}{1 - \sum_k F_k(0)} \mathbb{P}(W > 0) + \mathbb{P}(W = 0; N = n). \quad (2)$$

For the derivation of the probabilities $\mathbb{P}(W > 0)$ and $\mathbb{P}(W = 0; N = n)$ in (2) we refer to Appendix A. The stationary distribution π of the number of active source nodes N follows from (2) as

$$\pi_n = G_n(\infty). \quad (3)$$

Let π^a (π^d) denote the stationary distribution of the active number of source nodes present at a flow-arrival epoch (left behind at a source-departure epoch, which coincides with the arrival epoch of the last particle of a flow at the relay node). Then we have the following result.

Lemma 2 π, π^a and π^d are identical.

Proof $\pi = \pi^a$ follows directly from the PASTA-property. $\pi^d = \pi^a$ as N_t is a birth-death process and the number of up-crossings of level n balances the number of down-crossings of level n . ■

3.2 Flow transfer time and the buffer workload

This subsection first explains how to compute $\mathbb{E}D_{\text{source}}$, and then provides two expressions for $\mathbb{E}W_{\text{buffer}}$, one of them being based on the theory developed in Section 3.1. Finally it addresses the expected workload $\mathbb{E}W_{\text{buffer}}^*$ at the relay node at the epoch that a last particle of a flow arrives at the relay node, by considering the sum of the mean workload $\mathbb{E}W_{\text{buffer}}$ present at flow initiation and the mean workload increase $\mathbb{E}\Delta W_{\text{buffer}}$ at the buffer of the relay node during a flow transfer.

3.2.1 Expected flow transfer time $\mathbb{E}D_{\text{source}}$

The expected number of active source nodes is given by

$$\mathbb{E}N = \sum_{n=0}^{\infty} n\pi_n \quad (4)$$

(which we can approximate by $\sum_{n=0}^{n_{max}} n\pi_n$, but as we had chosen n_{max} large these expressions effectively coincide). The flow transfer time follows from Little's law:

$$\mathbb{E}D_{\text{source}} = \mathbb{E}N/\lambda \quad (5)$$

(or, more precisely, $\mathbb{E}N/(\lambda(1 - \pi_{n_{max}}))$), but again this will not give a substantially different answer if n_{max} is chosen sufficiently large).

3.2.2 Expected buffer workload at an arbitrary moment

The expected buffer workload $\mathbb{E}W_{\text{buffer}}$ at the relay node upon flow arrival can be obtained in two manners which are both presented below: in the first place by relying on the workload distribution, i.e., Expression (2), and in the second place as a direct application of Lemma 1.

Buffer workload at an arbitrary moment, using the workload distribution. Observe that $\pi_n^{-1}G_n(x)$ is the workload distribution conditional on n source nodes being active. Then, the expected conditional buffer workload is obtained by

$$\begin{aligned} \mathbb{E}[W_{\text{buffer}}|N = n] &= \pi_n^{-1} \int_0^\infty x dG_n(x) \\ &= \pi_n^{-1} \frac{1}{1 - \sum_i F_i(0)} \mathbb{P}(W > 0) \sum_{j=0, j \neq n+}^{n_{max}} \frac{a_j(\vec{v}_j)_n}{z_j}, \end{aligned}$$

see Appendix A. Hence, the expected unconditional buffer workload is:

$$\mathbb{E}W_{\text{buffer}} = \frac{1}{1 - \sum_i F_i(0)} \mathbb{P}(W > 0) \sum_{n=0}^{n_{max}} \sum_{j=0, j \neq n+}^{n_{max}} \frac{a_j(\vec{v}_j)_n}{z_j}. \quad (6)$$

Buffer workload at an arbitrary moment, using the relation between the model and the workload in the corresponding M/G/1 FCFS-queue. We define W_{sources} as the aggregate workload at all active source nodes, and W_{total} and W_{buffer} as in respectively Section 2.2 and 3.2. Recall that work at a source node still needs to be served twice, i.e., by the source node and relay node. Then,

$$W_{\text{total}} = W_{\text{sources}} + W_{\text{buffer}}. \quad (7)$$

Observe that W_{total} does not depend on the resource-sharing ratio m ; in fact, it does not even depend on the service discipline as long as it is work-conserving.

As $\mathbb{E}W_{\text{total}}$ is given by (1) with $c_F^2 = 1$, we are left to derive $\mathbb{E}W_{\text{sources}}$. This follows due to the memoryless property of the exponential distribution of the flow sizes. The

expected amount of work at an active source node (i.e., the residual of the flow) equals $2f/c$. Furthermore the expected number of source nodes simultaneously active is given by Expression (4). Hence

$$\mathbb{E}W_{\text{sources}} = \mathbb{E}N \cdot 2f/c,$$

and the expected workload at the buffer of the relay node is:

$$\mathbb{E}W_{\text{buffer}} = \left(\frac{2\rho}{1-2\rho} - \mathbb{E}N \right) \cdot 2f/c. \quad (8)$$

Remark 3 Relation between Expressions (6) and (8) for $\mathbb{E}W_{\text{buffer}}$. It is interesting to note that Expression (6) depends on all the eigenvalue-eigenvectors (z_j, \vec{v}_j) for $j = \{0, \dots, n_{\max}\}$; on the other hand Expression (8) depends just through $\mathbb{E}N$ on the normalized eigenvector that corresponds to the zero eigenvalue, cf. Expressions (4) and (3). Due to the implicitness of the eigenvalue-eigenvector pairs (z_j, \vec{v}_j) and corresponding constants a_j , it is not a priori obvious from these expressions that they match. \diamond

3.2.3 Buffer workload at the arrival epoch of the last particle of a flow at the relay node

In this section we derive the expected buffer workload $\mathbb{E}W_{\text{buffer}}^*$ at the epoch that the last particle of a flow arrives at the relay node. Note that the mean buffer workload on flow initiation coincides with the mean workload at an arbitrary epoch (PASTA). Hence, W_{buffer}^* can be obtained using the following relation:

$$\mathbb{E}W_{\text{buffer}}^* = \mathbb{E}W_{\text{buffer}} + \mathbb{E}\Delta W_{\text{buffer}}, \quad (9)$$

where ΔW_{buffer} denotes the buffer increase during the transfer time D_{source} of an arbitrary flow.

We are left to derive $\mathbb{E}\Delta W_{\text{buffer}}$. Let ΔW_{total} denote the increase in workload in the total system during the flow transfer time D_{source} by a source node, and let $\Delta W_{\text{sources}}$ denote the increase (which may be negative) of the aggregate workload of all source nodes during D_{source} . By (7) we evidently have the following relation:

$$\Delta W_{\text{total}} = \Delta W_{\text{sources}} + \Delta W_{\text{buffer}}. \quad (10)$$

Lemma 4 $\mathbb{E}\Delta W_{\text{buffer}} = \mathbb{E}\Delta W_{\text{total}}$.

Proof We have to prove that $\mathbb{E}\Delta W_{\text{sources}} = 0$, i.e., the expected amount of work at the source nodes present upon arrival of a flow coincides with the amount present

at the corresponding source-departure epoch. This property follows directly from two observations. First, due to Lemma 2, the expected number of source nodes at the flow-initiation epoch equals the expected number of source nodes at the epoch of the arrival of the last particle of a flow at the relay node. Second, the expected residual flow-sizes at these instances coincide due to the memoryless property of the exponentially distributed flow-sizes. ■

Lemma 4 and Relation (10) lead to the following proposition.

Proposition 5 *The expected increase of the buffer workload $\mathbb{E}\Delta W_{\text{buffer}}$ during the transfer time D_{source} of a flow is given by*

$$\mathbb{E}\Delta W_{\text{buffer}} = (\mathbb{E}N + 1) \cdot 2f/c - \mathbb{E}N/\lambda. \quad (11)$$

Proof Due to Lemma 4 we are left to compute $\mathbb{E}\Delta W_{\text{total}}$ during the flow transfer time (with mean $\mathbb{E}D_{\text{source}}$) of a tagged flow. The input into the total system is the result of initiations of new flows (including the tagged flow) which arrive at rate λ , each bringing along an amount of work with expected value $2f/c$ (cf. Lemma 1). The expected number of arrivals (including the tagged flow) is $\lambda\mathbb{E}D_{\text{source}} + 1$, and consequently the expected input into the total system is $(\lambda\mathbb{E}D_{\text{source}} + 1) \cdot 2f/c$.

The expected output is $\mathbb{E}D_{\text{source}}c$, as is readily verified by the following two observations. First, the total system is non-empty during the flow transfer time D_{source} as at least the tagged flow is served during D_{source} . Second, the total system is work-conserving and serves at rate c . Writing the expressions in terms of $\mathbb{E}N$ using (5) proves the lemma. ■

Notice that the expected workload $\mathbb{E}W_{\text{buffer}}^*$ of Expression (9), which is the sum of Expressions (8) and (11), only depends on the resource-sharing ratio m via $\mathbb{E}N$. Recall that $\mathbb{E}N$ is given by Expression (4) that can be determined by Expressions (2) and (3).

An interesting result follows from rewriting Expressions (8) and (11) in terms of $\mathbb{E}N$ and considering their ratio. It turns out that, remarkably, the proportionality constant does not depend on m .

Corollary 6 *The expected workload increase $\mathbb{E}\Delta W_{\text{buffer}}$ at the relay node during a flow transfer is proportional to the expected workload $\mathbb{E}W_{\text{buffer}}$ at flow arrival:*

$$\mathbb{E}\Delta W_{\text{buffer}} = \frac{1 - 2\rho}{2\rho} \mathbb{E}W_{\text{buffer}},$$

and Expression (9) can be written as

$$\mathbb{E}W_{\text{buffer}}^* = \frac{1}{2\rho} \mathbb{E}W_{\text{buffer}}.$$

3.3 Mean delay $\mathbb{E}D_{\text{buffer}}^*$ of the last particle and mean overall flow transfer time $\mathbb{E}D_{\text{overall}}$

At the moment that the source node has transmitted the full flow into the relay node, the last fluid particle enters the buffer at the relay node, and then the source node becomes inactive. In this subsection we present an approximation for the expected buffer delay $\mathbb{E}D_{\text{buffer}}^*$ of this last particle.

Recall that the last particle does not experience any buffer delay if the buffer is empty. In case the buffer is non-empty, the buffer delay D_{buffer}^* of the last particle is the time required by the relay node to serve the amount of work W_{buffer}^* present upon arrival of that last particle. Recall from Section 2.1 that during D_{buffer}^* the relay node uses the entire resource-sharing ratio m . Hence, during D_{buffer}^* the behavior of the system behaves as the model without feedback presented in Appendix A.1.

Conditional buffer delay of the last particle at the relay node. Let $Y_n(\tau)$ denote the conditional buffer delay, i.e., the time required by the relay node with resource-sharing ratio m to serve an amount τ of fluid if initially n source nodes are active. Here we again assume that there is a maximum n_{\max} imposed on the number of source nodes that may be simultaneously active, as in Section 3.1. Let \bar{Q} denote the generator matrix without feedback as in (A.1) and $M(s) := -\bar{Q} + sR$ where

$$R := \text{diag} \left\{ 1, \frac{m}{m+1}, \frac{m}{m+2}, \dots, \frac{m}{m+n_{\max}} \right\}.$$

Proposition 7 *The expected conditional time required to serve an amount τ of fluid, if initially n source nodes are active, is given by*

$$\mathbb{E}Y_n(\tau) = A_{(n,0)}\tau + \sum_{j=1}^{n_{\max}} \frac{A_{(n,j)}}{s_j} e^{s_j\tau} - \sum_{j=1}^{n_{\max}} \frac{A_{(n,j)}}{s_j} \quad (12)$$

where s_j denote the eigenvalues of $R^{-1}\bar{Q}$. The constants $A_{(n,j)}$ follow from the partial-fraction expansion of

$$s\phi_n(s) = \frac{\det M_{-n}(s)}{\det M(s)}.$$

where $M_{-n}(s)$ is defined as $M(s)$ with the n -th column replaced by $\vec{1}$.

Proof See Appendix B. ■

In Expression (12) the eigenvalues s_j for $j \geq 1$ are negative, and hence $\mathbb{E}Y_n(\tau)$ is approximately linear in τ .

Approximation of the buffer delay. We now use Proposition 7 to approximate the buffer delay experienced by the last particle of the flow. By definition the expected buffer delay $\mathbb{E}Y_n(\tau)$ of the last particle can be expressed as

$$\mathbb{E}D_{\text{buffer}}^* = \sum_{n=0}^{n_{\max}} \pi_n^d \int_0^\infty \mathbb{E}Y_n(\tau) w_n^*(\tau) d\tau,$$

where $w_n^*(\tau)$ is the probability density function of the amount of work at the buffer at a source-departure epoch leaving behind n source nodes, and π_n^d coincides with π (due to Lemma 2). Unfortunately, we do not have the density function $w_n^*(\tau)$.

If one assumes, however, that $\mathbb{E}Y_n(\tau)$ is linear in τ , the conditional buffer delay roughly looks like

$$\int_0^\infty \mathbb{E}Y_n(\tau) w_n^*(\tau) d\tau \approx \mathbb{E}Y_n(\mathbb{E}W_{\text{buffer}}^*). \quad (13)$$

Then we obtain the following approximation for the expected delay D_{buffer}^* of the last particle.

Approximation 8 *The buffer delay of the last particle can be approximated by*

$$\mathbb{E}D_{\text{buffer}}^* \approx \sum_{n=0}^{n_{\max}} \pi_n \mathbb{E}Y_n(\mathbb{E}W_{\text{buffer}}^*). \quad (14)$$

Remark 9 Special case $m = 1$. As mentioned earlier, the special case m equals 1 was studied in [3]. For exponential flow-sizes and $m = 1$ and $n_{\max} \rightarrow \infty$ a closed-form expression for the equivalent of (12) is available, namely Expression (33) of [6]. It is seen that Expressions (33) of [6] and (12) are very similar in nature, in the sense that both expressions consist of a linear term and in addition a term that is exponentially decaying in τ . ◇

Now we have derived expressions for all the parts of the main performance metric from the user perspective: the mean expected overall flow transfer time $\mathbb{E}D_{\text{overall}}$ is

$$\mathbb{E}D_{\text{overall}} = \mathbb{E}D_{\text{source}} + \mathbb{E}D_{\text{buffer}}^* \quad (15)$$

where $\mathbb{E}D_{\text{source}}$ is given by (5) and $\mathbb{E}D_{\text{buffer}}^*$ is approximated by (14).

4 Analysis for general flow-size distributions

This section treats the analysis the performance metrics for *generally* distributed flow-sizes. The analysis presented in this section borrows elements from the approach followed in the previous section for the case of exponentially distributed flow-sizes.

Our analysis relies heavily on knowledge of the stationary distribution of the number of active source nodes, together with the expected residual flow-sizes at the source nodes. In Section 4.1 we present an approximation assumption that states that the distribution of the number of active source nodes is insensitive to the flow-size distribution, and we also state two (related) properties concerning the expected residual flow-sizes at the source nodes. (These claims will be thoroughly assessed in Section 5.3.) The approximation assumption relates to a general resource-sharing ratio $m \geq 0$, and we use it to derive the mean workload at the relay node in Section 4.2. In Section 4.3 we then consider the buffer delay of the last particle of a flow, which, together with earlier results, enables us to compute the mean overall flow transfer time. Finally, Section 4.4 presents an overview of all the expressions required to evaluate the performance metrics, in the form of a calculation scheme.

4.1 Steady-state behavior of the active source nodes

By extensive simulations of the fluid model, for $m \geq 0$, we observed the striking property (i) that the source-node behavior (in terms of the distribution of the number of active source nodes) is *insensitive* to the flow-size distribution, i.e., only the mean flow-size plays a role. In addition, our simulations revealed that the system exhibits two other characteristics that are closely related to insensitive systems, i.e., (ii) the residual flow-sizes at the source nodes are very well approximated by the ‘usual’ excess life distribution as known from renewal theory, and (iii) the residual flow-sizes are nearly independent of the number of active source nodes. The numerical evidence for these claims is presented in detail in Section 5.3 (see Figures 5–8). We assessed the properties for wide ranges of the parameter settings including (extremely) heavy loads, various flow-size distributions, and high resource-sharing ratios. The simulations indicate that the claim (i) is exact, whereas claims (ii) and (iii) seem to hold as a very accurate approximation; in fact it took a huge number of replications to show that those statements were not exact. This motivates the use of the three properties (i)–(iii) as *approximation assumptions*.

Let us now formally state the approximation assumptions. As mentioned above, convincing support is provided by the extensive simulation results, to be presented

in Section 5.3, but it is noted that in Remark 11 we formally *prove* that for the special case $m \in [0, 1]$ the assumption holds true.

Assumption 10

(i) *The stationary distribution of the number of active source nodes is insensitive to the flow-size distribution apart from its mean and is given by Expression (3).*

(ii) *The expected residual flow-size $\mathbb{E}[F^r]$ of a flow at a source node coincides with the expected residual flow-size of a renewal process, i.e.,*

$$\mathbb{E}[F^r] = (1/2)(c_F^2 + 1) \cdot f = \frac{f_2}{2f}. \quad (16)$$

(iii) *The number of active source nodes N_t and their expected residual flow-sizes $\mathbb{E}[F^r]$ are independent.*

Properties as those mentioned in Assumption 10 are well-known to hold for stationary symmetric queues, cf., e.g., Kelly [12], Cohen [7], Bonald and Proutière [4], and more recently the work of Zachary [23]. However, the service discipline of our model is *not* symmetric; the requirement that the service rate only depends on N_t is not fulfilled as the service rate also depends on the workload W_t at the buffer of the relay node. We would therefore like to stress that the (empirically observed) insensitivity of Assumption 10.(i) is a highly remarkable property: to the authors’ best knowledge there are no results on other insensitive queueing-systems that do not have a symmetric service-discipline. It is a subject for further research to formally prove this insensitivity.

Remark 11 Insensitivity of the source-node behavior for $m \in [0, 1]$. For $m \in [0, 1]$ the relay node obtains a service rate less than or equal to the share that each active source node obtains. Therefore, the relay node always obtains its entire share m if there is work in the system; consequently, the resource sharing only depends on the number of active sources (and no information on the buffer content is needed).

Hence, the behavior of the source nodes is described by a Processor Sharing queue with state-dependent service rates $nc/(n + m)$ if n source nodes are active. This model is a special case of the so-called Generalized Processor Sharing queue described by Cohen in [7] for which he presented a joint stationary probability/density function of the number of active sources nodes N and their residual service requirements. Via this result we obtain the distribution of the number of source nodes

$$\pi_n = (1 - \rho)^{m+1} \cdot \rho^n \prod_{k=1}^n \frac{m + k}{k}.$$

In [7] it was shown that the stationary distribution is *insensitive* to the flow-size distribution and that it is independent of the residual flow-sizes. Further, [7] establishes

that the residual flow-sizes are distributed according to the excess life distribution from renewal theory, see e.g. Expression (16), and that the residual flow sizes and the number of active source nodes are independent. \diamond

4.2 Mean buffer workload

For the expected buffer workload $\mathbb{E}W_{\text{buffer}}$ we use the relation $\mathbb{E}W_{\text{total}} = \mathbb{E}W_{\text{sources}} + \mathbb{E}W_{\text{buffer}}$. The mean total workload $\mathbb{E}W_{\text{total}}$ is given by (1), and $\mathbb{E}W_{\text{sources}}$ follows from Assumption 10 and equals

$$\mathbb{E}W_{\text{sources}} = \mathbb{E}N \cdot \mathbb{E}[F']. \quad (17)$$

We obtain the following expression for the expected workload in the buffer.

Lemma 12 *Under Assumption 10,*

$$\mathbb{E}W_{\text{buffer}} = \left(\frac{2\rho}{1-2\rho} - \mathbb{E}N \right) \cdot (c_F^2 + 1)f/c. \quad (18)$$

Note that Expression (18) coincides with Expression (8) as $c_F^2 = 1$ for exponentially distributed flow-sizes. Corollary 13 results from considering the ratio of Expressions (18) and (8).

Corollary 13 *Under Assumption 10, the buffer workload in case of general flow-size distribution relates to the workload in case of exponential flow-size distribution, with the same mean f , in the following manner. In self-evident notation,*

$$\mathbb{E}W_{\text{buffer}} = \frac{(c_F^2 + 1)}{2} \mathbb{E}W_{\text{buffer}}^{\text{exp}}. \quad (19)$$

The important implication of relation (19) is that it entails that the expected buffer workload for general flow-size distributions is *proportional to* the expected buffer workload for exponential flow-sizes. It is stressed that the proportionality constant just includes the cov, and, importantly, that this factor is independent of the resource-sharing ratio m (but recall that, evidently, $\mathbb{E}W_{\text{buffer}}^{\text{exp}}$ *does* depend on m).

Using Expression (9), the mean buffer workload $\mathbb{E}W_{\text{buffer}}^*$ at the arrival of the last particle is the sum of $\mathbb{E}W_{\text{buffer}}$ and the mean workload increase $\mathbb{E}\Delta W_{\text{buffer}}$. The latter is derived in Section 3.2, and observe that the derivation is independent of the flow-size distribution, i.e., Expression (11) holds for general flow-size distributions. As a consequence, imposing Assumption 10, $\mathbb{E}N$ is still given by (4).

4.3 Mean buffer delay of the last particle in case of general flow-size distributions

For the mean buffer delay $\mathbb{E}D_{\text{buffer}}^*$ of the last particle we use Approximation (14) which is derived in Section 3.3, although the approximation is derived assuming exponentially distributed flow-sizes. This procedure is motivated by considering the two ways in which the flow-size distribution has impact on the buffer delay.

- First, it affects the buffer workload W_{buffer}^* seen by the last particle, but recall that this effect could (under Assumption 10) be captured, and resulted in Expression (11), see the remarks at the end of Section 4.2.
- The second effect is on the transient behavior during buffer delay D_{buffer}^* where the resource sharing depends on the number of active source nodes and their residual flow sizes. Recall that during the entire D_{buffer}^* the relay node continuously obtains ratio m , which is, importantly, a symmetric service-discipline (therefore corresponding to an insensitive invariant distribution). Small flows have a small delay anyway, whereas long jobs will see a number of source nodes that is (by approximation) distributed according to this invariant. This argumentation suggests that the impact of the distribution of the flow-sizes remains limited.

4.4 Calculation scheme of the performance metrics for general flow-sizes

In order to facilitate easy evaluation of all performance metrics involved, we present in Table 2 an overview of the expressions required to calculate the performance metrics. For each performance metric we state the equation number of the corresponding expression in this paper, which other expressions are required to calculate this expression, and how the expression was derived (where FQWF is an abbreviation of ‘fluid-queue with feedback’).

5 Numerical results

This section serves three goals: (i) to numerically illustrate the behavior of the system as described in Section 2.1 (or, more particularly, to assess the impact of the ratio m under various loads, and for various flow-size distributions), (ii) to provide empirical evidence for Assumption 10, and (iii) to validate the approximations proposed in the previous sections. Section 5.1 introduces the scenarios and the parameter settings used in the numerical examinations. Section 5.2 presents results for exponentially distributed flow-sizes and general resource-sharing ratios (cf. Section 3). Numerical support of Assumption 10 is provided in Section 5.3. Finally,

Table 2

Performance metrics calculation scheme for general flow-size distributions.

Performance metric	corresponding expression	required expressions	derived by
$\mathbb{E}N$	(4)	(2)	FQWF, Conj. 10
$\mathbb{E}D_{\text{source}}$	(5)	(4)	FQWF, Conj. 10
$\mathbb{E}W_{\text{buffer}}$	(19)	(4)	FQWF, Conj. 10 or Expr. (7)
$\mathbb{E}\Delta W_{\text{buffer}}$	(11)	(4)	Expr. (7)
$\mathbb{E}W_{\text{buffer}}^*$	(9)	(19), (11)	Expr. (7)
$\mathbb{E}D_{\text{buffer}}^*$	(14)	(9), (12)	Laplace transforms
$\mathbb{E}D_{\text{overall}}$	(15)	(5), (14)	Expr. (7)

Section 5.4 focuses on the performance metrics for general flow-size distributions (cf. Section 4).

5.1 Validation scenarios

The analytical expressions were evaluated by Matlab. Simulations of the fluid-model were performed by a simulator which was built in Delphi.

We assume a system capacity $c = 5$ mbit/s and an average flow size f of 0.12 mbits. The average file size of 0.12 mbits in the model corresponds to files of 10 packets of each 1500 bytes in an actual communication system (cf. [3, 20]). For the flow-size distributions we used Deterministic, Erlang-4, Exponential and Hyper-Exponential (with balanced means – see e.g. [21] – and a coefficient of variation (cov) of 2, 4, and 16). The resource-sharing ratio m is varied over the values $\{0, 1, 2, 3, 5, 10\}$. The load ρ is varied from very low loads (0.024) up to almost saturation (0.48), by varying the flow-initiation rate λ ; recall that, for stability, $\rho < \frac{1}{2}$ is required. Finally, the maximum number of active source nodes n_{max} is set to 20, which is chosen so large that it effectively means that there is no blocking.

Sufficiently many replications have been simulated in order to obtain small confidence intervals. In all experiments the confidence interval's half-width divided by the estimate should be below 5%. We remark that this ratio is only close to 5% for high loads *and* high cov (i.e., cov= 16), in the other cases the ratio is usually well below 1%.

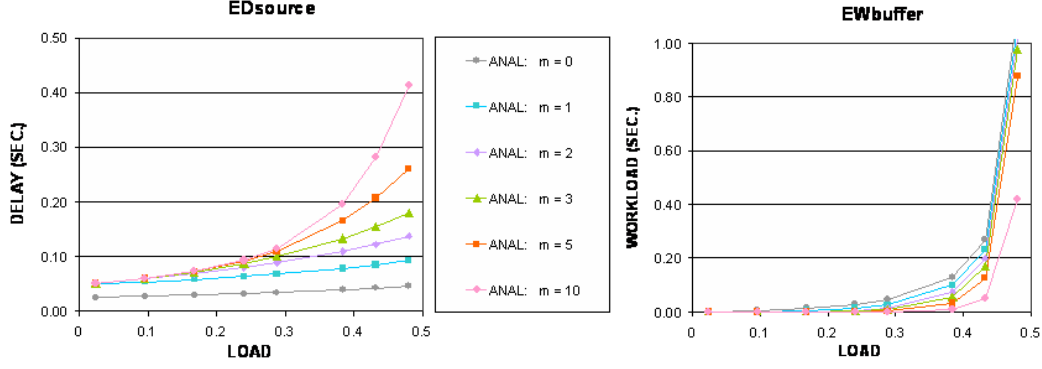


Fig. 3. Exponential flow-sizes. Left: $\mathbb{E}D_{\text{source}}$. Right: $\mathbb{E}W_{\text{buffer}}$.

5.2 Results for exponentially distributed flow-sizes

This section presents numerical results for exponentially distributed flow-sizes. Figure 3 presents the mean flow transfer time $\mathbb{E}D_{\text{source}}$ (left graph) and the buffer workload $\mathbb{E}W_{\text{buffer}}$ (right graph) for different values of resource-sharing ratio m as a function of the load. Recall that the results are exact for exponentially distributed flows; therefore we do not compare these results with simulations. The graphs illustrate the influence of resource-sharing ratio m : a small ratio m implies that the source nodes obtain a large share of the capacity resulting in short flow transfer times $\mathbb{E}D_{\text{source}}$ for the source nodes. On the other hand for small ratios m the relay node obtains a small share of the capacity which results in a larger mean buffer workload $\mathbb{E}W_{\text{buffer}}$.

Figure 4 presents the approximation of the mean buffer delay $\mathbb{E}D_{\text{buffer}}^*$ (left graph) and the mean overall flow transfer time $\mathbb{E}D_{\text{overall}}$ (right graph). Approximation (14) of the buffer delay performs very well: it is close to the simulation results. The right graph shows the overall performance $\mathbb{E}D_{\text{overall}}$. The small error between the analysis and simulation results is solely due to the approximation of the buffer delay




Fig. 4. Exponential flow-sizes. Left: $\mathbb{E}D_{\text{buffer}}^*$. Right: $\mathbb{E}D_{\text{overall}}$.

$\mathbb{E}D_{\text{buffer}}^*$, as $\mathbb{E}D_{\text{source}}$ is exact. Also, the graph illustrates that the $\mathbb{E}D_{\text{overall}}$ improves for increasing m . Observe that the curves of $\mathbb{E}D_{\text{overall}}$ for different values of m are close to each other for $\rho < 0.4$; this indicates that the trade-off between $\mathbb{E}D_{\text{source}}$ and $\mathbb{E}D_{\text{buffer}}^*$ is more or less balanced for these regimes.

5.3 Numerical evidence for Assumption 10

Recall that Assumption 10 consists of three parts, which we will validate separately. For $m \in [0, 1]$ these results can be proven, see Remark 11; in this section we therefore focus on $m > 1$.

To validate part (i) of the assumption (i.e., the distribution of the source nodes is insensitive to the flow-size distribution), we compare the stationary distribution of the number of active source nodes obtained by simulations of the fluid model (for various flow-size distributions) with the exact results for the exponential flow-size distribution. Figures 5, 6 and 7 show the results for high loads $\rho \in \{0.38, 0.48\}$ and resource-sharing ratio $m \in \{2, 5, 10\}$. We observe that the distributions from simulation and the analysis coincide for all ranges, i.e., the analytically obtained probability of n active source nodes falls within each of the confidence intervals of the simulated probability, for *all* flow-size distributions. These results offer strong support for the first part of Assumption 10 for all loads and resource-sharing ratios. Observe also the remarkable shapes of the stationary distribution; in particular consider the shape for high load and high resource-sharing ratio in Figure 7.

Figure 8 presents numerical results for the second part of Assumption 10 (i.e., the expected residual flow-size at a source node coincides with the expected residual excess flow-size from renewal theory). The numerical results from the simulation and the analysis are very close together, although the proposed mean residual flow-

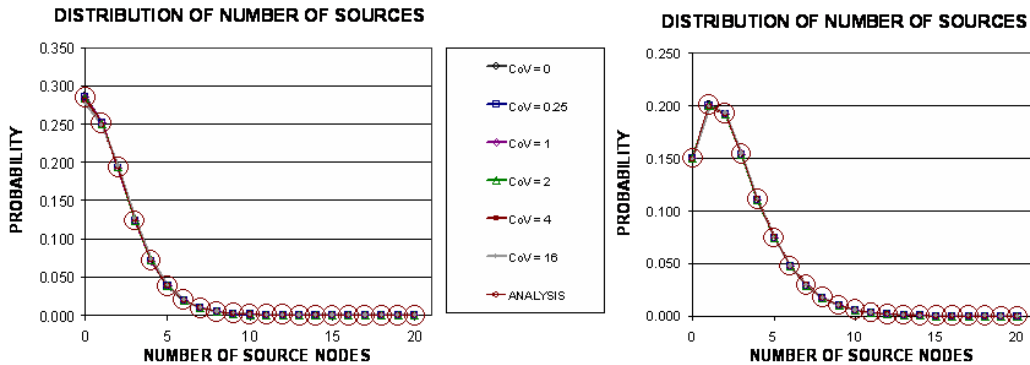


Fig. 5. Distributions of the number of source nodes for ratio $m = 2$ for different flow-size distributions. Left: $\lambda = 16, \rho = 0.38$. Right: $\lambda = 20, \rho = 0.48$.

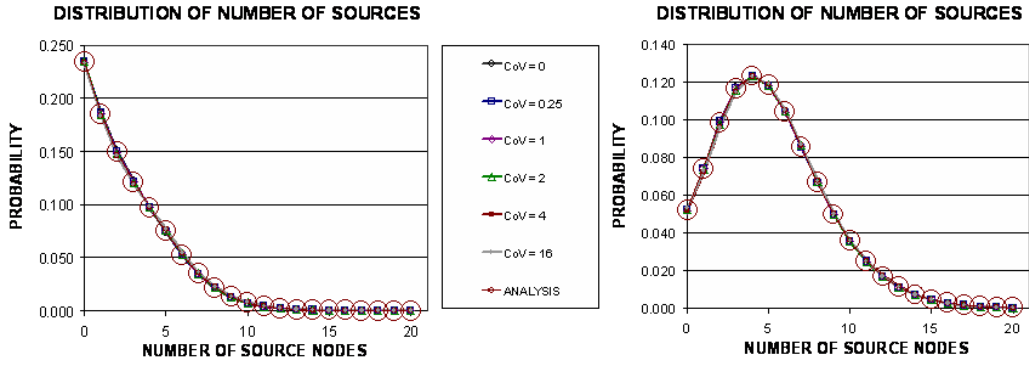


Fig. 6. Distributions of the number of source nodes for ratio $m = 5$ for different flow-size distributions. Left: $\lambda = 16, \rho = 0.38$. Right: $\lambda = 20, \rho = 0.48$.




Fig. 7. Distributions of the number of source nodes for ratio $m = 10$ for different flow-size distributions. Left: $\lambda = 16, \rho = 0.38$. Right: $\lambda = 20, \rho = 0.48$.

size distribution given by Expression (16) is not always within the confidence interval of the mean residual flow-sizes at the source nodes as obtained from our simulations. For scenarios with a high load and a flow-size distribution with cov

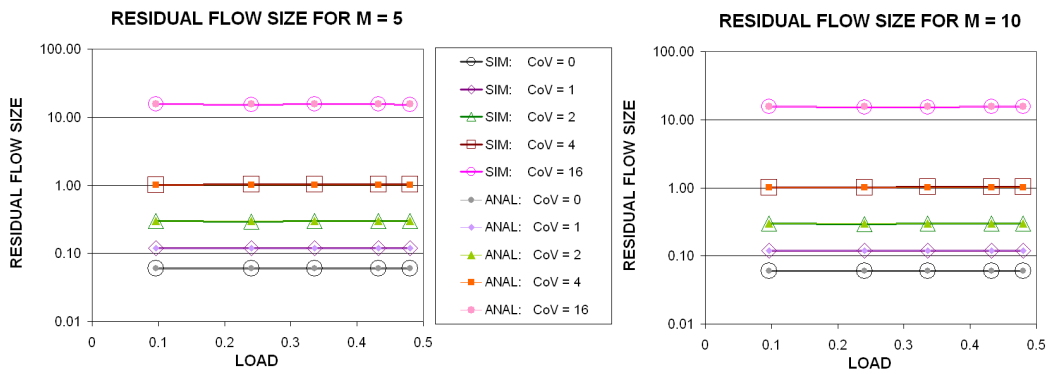


Fig. 8. Residual flow-sizes at the source nodes for different flow-size distributions. Left: $m = 5$. Right: $m = 10$.

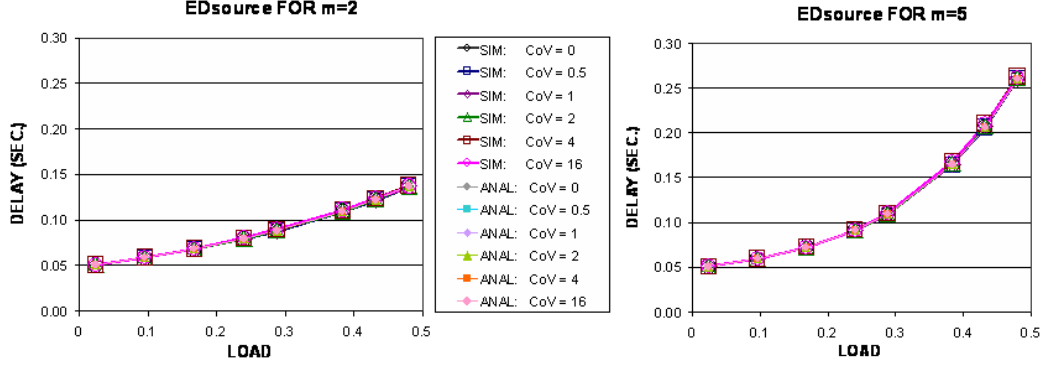


Fig. 9. Flow transfer time $\mathbb{E}D_{\text{source}}$ for general flow-size distributions. Left: $m = 2$. Right: $m = 5$.

> 1 we observe that the simulated mean residual flow size is slightly larger than the analytical value. In summary, the curves are very close together, and it seems justified to use the claimed as an approximation assumption.

The third part of Assumption 10 deals with the independence of the number of active source nodes and their residual flow-sizes. This assumption is mainly required to support Expression (17), i.e., $\mathbb{E}W_{\text{sources}} = \mathbb{E}N \cdot \mathbb{E}[F']$. Therefore we directly sampled $\mathbb{E}W_{\text{sources}}$, $\mathbb{E}N$, and $\mathbb{E}[F']$ from simulations; there we observed that both sides of Expression (17) match (up to a high level of precision), and the same applies to the mean residual flow size $\mathbb{E}[F']$ (which follows from Expression (16)) and the directly sampled value.

5.4 Results for general flow-size distributions

This section focuses on the validation of the analysis for general flow-size distributions. We do so by comparing the output of our calculation scheme with estimates obtained from fluid simulations. We present the main performance metrics for general flow-size distributions and general resource-sharing ratio m . Each graph in Figures 9–12 shows the performance metrics as a function of the load for various flow-size distributions. In each graph the resource-sharing ratio m is fixed: left $m = 2$ and right $m = 5$. The effects of the load and resource-sharing ratio on the performance metrics are similar to the results for the exponential case of Section 5.2 and will not be discussed again.

Figure 9 presents the results for the mean flow transfer time $\mathbb{E}D_{\text{source}}$. Note that, due to the (empirically observed) insensitivity of the distribution of the number of active source nodes, the curves coincide for the different flow-size distributions; the deviations between analysis and simulations are overall less than a percent, for




Fig. 10. Buffer workload $\mathbb{E}W_{\text{buffer}}$ for general flow-size distributions. Left: $m = 2$. Right: $m = 5$.

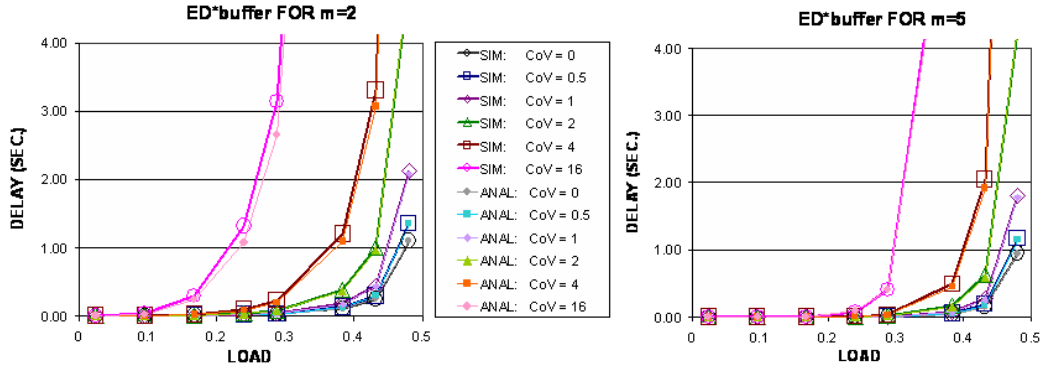


Fig. 11. Buffer delay $\mathbb{E}D_{\text{buffer}}^*$ for general flow-sizes distributions. Left: $m = 2$. Right: $m = 5$.

cov = 16 less than two percent. The mean buffer workload is shown in Figure 10. Recall that, under Assumption 10, the analysis should be exact for both $\mathbb{E}D_{\text{source}}$ and $\mathbb{E}W_{\text{buffer}}$. At this point we see that, for high cov's, the analysis overestimates $\mathbb{E}W_{\text{buffer}}$, as is illustrated by the numerical results.

Figure 11 shows the approximation of the delay $\mathbb{E}D_{\text{buffer}}^*$ of the last particle. It is observed that the resulting curves are close to the results of the fluid-model simulations. This strengthens the explanation in Section 4.3 that, although the analytical derivation of the expected conditional delay $\mathbb{E}Y_n(\tau)$ relies on exponentially distributed flow-sizes, the flow-size distribution hardly affects the outcome. As a result, the approximation $\mathbb{E}D_{\text{buffer}}^*$ gives a good approximation for general flow-sizes. Finally, Figure 12 presents the results for the mean overall flow transfer time $\mathbb{E}D_{\text{overall}}$. As it is the sum of the exact $\mathbb{E}D_{\text{source}}$ and the (accurately) approximated $\mathbb{E}D_{\text{buffer}}^*$, it implies that the overall flow transfer time has a remarkably good fit.

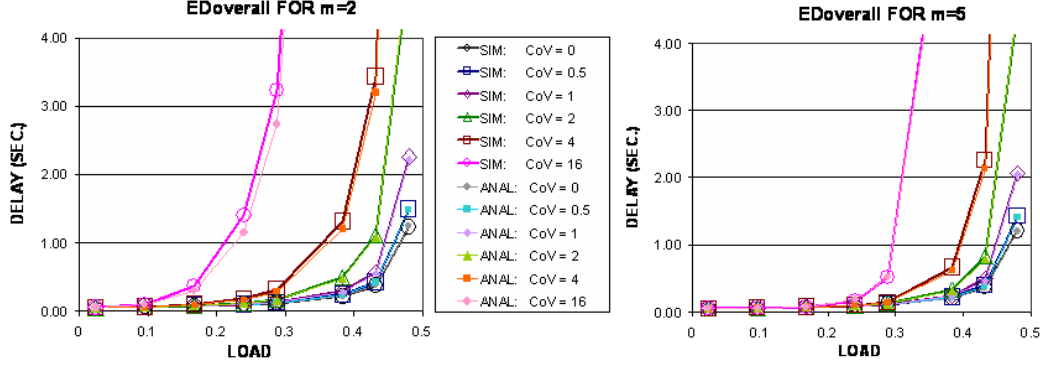


Fig. 12. Overall flow transfer time ED_{overall} for general flow-size distributions. Left: $m = 2$. Right: $m = 5$.

6 Concluding remarks and further research

In this work we presented a method to analyze the impact of the resource-sharing policy in a wireless ad-hoc network. We considered a setting where source nodes transmit flows to destinations via a common relay node. We obtained explicit expressions for the means of a number of performance metrics, such as the transfer time of a flow and the workload at the relay node.

The source-node behavior does not constitute a Markov chain (for $m > 1$), but, when assuming exponential flow-sizes, the joint distribution of the number of active source nodes and the workload can be analyzed using feedback fluid queues. We claim the remarkable fact that the obtained stationary distribution of the number of active source nodes is even valid for *generally* distributed flow-sizes, as we argue that the source-node behavior is actually insensitive to the flow-size distribution (apart from its mean); the latter claim is supported through extensive simulation experiments. Under this insensitivity claim we derived a number of expressions (some of them being exact, others approximations) for the performance metrics under consideration. Again by simulation it was shown that these expressions are highly accurate over a broad set of parameter values.

Further research includes the generalization to include state-dependent capacity c_n and resource-sharing ratio m_n which is seen in real communication systems, see, e.g., [20]. Another possible extension concerns *service-based qos-differentiation* where source nodes can obtain different shares of the capacity based on the priorities of their services; these priorities can even be dynamically adjusted based on the advertised buffer content per node. Another interesting topic for future research relates to models with multiple hops. This introduces so-called ‘hidden nodes’, and as a result there is not a single resource c shared by all nodes, but multiple resources shared by non-disjoint subsets of nodes.

References

- [1] D. Anick, D. Mitra, and M.M. Sondhi. Stochastic theory of a data handling system with multiple sources. *Bell System Technical Journal*, 61:1871–1894, 1982.
- [2] R. Bekker and M.R.H. Mandjes. A fluid model for a relay node in an ad-hoc network: the case of heavy-tailed input. *Mathematical Methods in Operations Research*, to appear, 2008. Available as cwi-report at <http://ftp.cwi.nl/CWIreports/PNA/PNA-E0703.pdf>.
- [3] J.L. van den Berg, M.R.H. Mandjes, and F. Roijers. Performance modeling of a bottleneck node in an IEEE 802.11 ad-hoc network. In T. Kunz and S.S. Ravi, editors, *Proceedings of AdHoc-Now2006*, volume 4104 of LNCS, pages 321–336, 2006.
- [4] T. Bonald and A. Proutière. Insensitivity in processor-sharing networks. *Performance Evaluation*, 49:193–209, 2002.
- [5] S. Borst, O. Boxma, and N. Hegde. Sojourn times in finite-capacity processor-sharing queues. In *Proceedings of 1st EURO-NGI Conference*, 2005.
- [6] E.G. Coffman jr., R.R. Muntz, and H. Trotter. Waiting time distributions for processor-sharing systems. *Journal of the ACM*, 17:123–130, 1970.
- [7] J. Cohen. The multiple phase service network with generalized processor sharing. *Acta Informatica*, 12:245–284, 1979.
- [8] A. Elwalid and D. Mitra. Analysis and design of rate-based congestion control of high speed networks, I: stochastic fluid models, access regulation. *Queueing Systems*, 9:29–64, 1991.
- [9] C.H. Foh and M. Zukerman. Performance analysis of the IEEE 802.11 MAC protocol. In *Proceedings of European Wireless '02*, 2002.
- [10] R. Howard. *Dynamic Probabilistic Systems*, volume 1. Wiley, New York, 1971.
- [11] IEEE p802.11e 2005. Amendment 8: Medium access control (MAC) quality of service enhancements, 2005.
- [12] F.P. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.
- [13] L. Kosten. Stochastic theory of data-handling systems with groups of multiple sources. *Performance of Computer Communication Systems*, pages 321–331, 1984.
- [14] R.B. Lenin and P.R. Parthasarathy. A computational approach for a fluid queue driven by a truncated birth-death process. *Journal Methodology and Computing in Applied Probability*, 2:373–392, 2004.
- [15] R. Litjens, F. Roijers, J.L. van den Berg, R.J. Boucherie, and M.J. Fleuren. Performance analysis of wireless LANS: An integrated packet/flow level approach. In *Proceedings of ITC-18*, pages 931–940, 2003.
- [16] M.R.H. Mandjes, D. Mitra, and W.R.W. Scheinhardt. Models of network

- access using feedback fluid queues. *Queueing Systems*, 44:365–398, 2003.
- [17] M.R.H. Mandjes and F. Roijers. A fluid system with coupled input and output and its application to bottlenecks in ad hoc networks. *Queueing Systems*, 56:79–92, 2007.
- [18] D. Mitra. Stochastic theory of a fluid model of producers and consumers coupled by a buffer. *Advances in Applied Probability*, 20:646–676, 1984.
- [19] A. Remke, B.R. Haverkort, and L. Cloth. A versatile infinite-state markov reward model to study bottlenecks in 2-hop ad hoc networks. In *Proceedings of QEST*, pages 63–72, 2006.
- [20] F. Roijers, J.L. van den Berg, and M.R.H. Mandjes. Fluid-flow modeling of a relay node in an IEEE 802.11 wireless ad-hoc network. In L. Mason, T. Drwiega, and J. Yan, editors, *Proceedings of ITC-20*, volume 4516 of LNCS, pages 321–334, 2007.
- [21] H.C. Tijms. *Stochastic models: an algorithmic approach*. Wiley & Sons, 1994.
- [22] E. van Doorn, A. Jagers, and J. de Wit. A fluid reservoir regulated by a birth death-process. *Stochastic Models*, 4:457–472, 1987.
- [23] S. Zachary. A note on insensitivity in stochastic networks. *Journal of Applied Probability*, 44:238–248, 2007.

A Analysis of source-node behavior for exponentially distributed flow-sizes

This section presents a more comprehensive analysis of the results presented in Section 3.1, i.e., the stationary distribution of the number of active source nodes and the buffer workload.

We assume that flow-sizes are exponentially distributed with mean f . The source-node behavior of the model of Section 2.1 is not an autonomous process; the transition rates depend on both the number of active source nodes N_t and on whether the buffer workload W_t is positive or not. Hence, N_t does not constitute a Markov chain as it requires *feedback* of the workload W_t , e.g., see [16].

We analyze the source-node behavior analogously to [16]. First we analyze the fluid-queue *without feedback*, i.e., the system in which the relay node is always allotted a share $mc/(n+m)$ of the capacity (when there are n source nodes transmitting); random variables (and other quantities) corresponding to the model without feedback are denoted with a bar ‘-’ on top. Now, \bar{N}_t constitutes a Markov chain and the joint distribution of (\bar{N}_t, \bar{W}_t) is derived in terms of a system of linear differential equations as in the seminal studies on fluid queues [1, 13]. The result without feedback is extended to the case with feedback by the important observation that

the behavior *during busy periods* of both models coincide. Finally, the joint distribution (N_t, W_t) of the model with feedback follows from rescaling the distribution (\bar{N}_t, \bar{W}_t) of the model without feedback. In the following we make this procedure precise.

A.1 Fluid-queue without feedback

First, we consider the model of Section 2.1 without feedback, i.e., the relay node *always* obtains its entire resource-sharing ratio m . We introduce \bar{W}_t as the buffer workload at time t for a system without feedback. Consequently, \bar{N}_t does not depend on \bar{W}_t , and therefore \bar{N}_t constitutes a Markov chain with generator matrix \bar{Q} , given through

$$\bar{Q}(i, j) := \begin{cases} \lambda & \text{if } j = i + 1, \\ ic/((m + i)f) & \text{if } j = i - 1, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.1})$$

where the diagonal elements are such that the rowsums are 0. For technical reasons we first assume that the resource-sharing ratio m is non-integer; the case of integer values of m is explained at the end of this section.

Let \bar{R} be a diagonal matrix where \bar{R}_n is the *net* input rate into the relay node, i.e., if n sources are active then

$$\bar{R}_n = \frac{n - m}{n + m}c.$$

Denote by D (U) all states with negative (positive) drift, i.e., $\bar{R}_n < 0$ for all $n \in D$ and $\bar{R}_n > 0$, for all $n \in U$. Let \bar{Q}_{DD} , \bar{Q}_{UU} be submatrices obtained by partitioning of \bar{Q} according to the ‘up states’ and ‘down states’.

We define the stationary distribution of (\bar{N}_t, \bar{W}_t) as

$$F_n(x) := \lim_{t \rightarrow \infty} P(\bar{W}_t \leq x; \bar{N}_t = n) = P(\bar{W} \leq x; \bar{N} = n).$$

For the analysis we assume that a maximum n_{max} is imposed on the number of source nodes that may be simultaneously active; flows that are initiated if already n_{max} other source nodes are active are blocked. The buffer workload satisfies the Kolmogorov forward equations $\vec{F}'(x)\bar{R} = \vec{F}(x)\bar{Q}'$. The spectral expansion of the

solution is given by

$$\vec{F}(x) = \sum_{j=0}^{n_{max}} a_j \vec{v}_j \exp(z_j x)$$

where (z_j, \vec{v}_j) is an eigenvalue-eigenvector pair, i.e., a scalar and vector that solve $z_j \vec{v}_j \bar{R} = \vec{v}_j \bar{Q}$. Clearly, $\mathbb{P}(\bar{W} \leq x) = \sum_n F_n(x)$. Further, let ω denote the stationary distribution of \bar{N}_t without feedback, hence $\omega_n = F_n(\infty)$. For stability we assume that the average net input rate is negative, that is, $\sum_n \omega_n \bar{R}_n < 0$. As our generator matrix \bar{Q} corresponds to a birth-death chain, all eigenvalues z_j are real [22].

Following Mitra [18], the number of negative eigenvalues n_+ in a stable system (i.e., $\sum_n \omega_n \bar{R}_n < 0$) is equal to the number of states with positive drift, i.e., $n_+ = n_{max} - \lceil m \rceil$; exactly one eigenvalue has value zero and the remaining eigenvalues are positive. In the remainder we label the eigenvalues z_j such that $z_j < 0$ for $j \in \{0, \dots, n_+ - 1\}$, $z_{n_+} = 0$, and $z_j > 0$ for $j \in \{n_+ + 1, \dots, n_{max}\}$.

The coefficients a_j are calculated as follows. When $z_j > 0$, then $a_j = 0$ as the distribution function should be in $[0, 1]$. The other coefficients a_j are computed from $F_i(0) = 0$ for all up states i . Further observe that

$$\omega = \frac{\vec{v}_{n_+}}{\langle \vec{v}_{n_+}, \vec{1} \rangle},$$

where $\langle \cdot, \cdot \rangle$ denotes the (standard) inner product. For computationally efficient numerical schemes, see e.g. [14].

Elwalid and Mitra [8] presented explicit expressions for a number of quantities related to the busy and idle periods of the workload at the relay node. A busy (idle) period is the period during which the workload at the relay node is positive (zero). A busy period starts when the system is empty and N_t becomes larger than m by a flow initiation. A busy period ends when the buffer becomes empty, and then N_t is in a state in D .

Denote by \vec{P} the distribution of \bar{N} at the end of the busy period. Then, due to Expression (5.9) of [8],

$$\vec{P} = \frac{1}{\langle \vec{F}_D(0) \bar{Q}_{DD}, \vec{1} \rangle} \vec{F}_D(0) \bar{Q}_{DD}.$$

Note that the (i, j) entry of $-(\bar{Q}_{DD})^{-1}$ is the mean time spent in state j by N_t , if the process started in state i , before leaving the set D , see e.g. [10]. Then, the mean idle

period $\mathbb{E}\bar{I}$ is given by

$$\mathbb{E}\bar{I} = \left\langle -\vec{P}(\bar{Q}_{DD})^{-1}, \vec{1} \right\rangle. \quad (\text{A.2})$$

The mean busy period $\mathbb{E}\bar{B}$ is obtained from $\sum_{i \in D} F_i(0) = \mathbb{E}\bar{I}/(\mathbb{E}\bar{B} + \mathbb{E}\bar{I})$:

$$\mathbb{E}\bar{B} = \mathbb{E}\bar{I} \cdot \frac{1 - \sum_{n \in D} F_n(0)}{\sum_{n \in D} F_n(0)}.$$

Remark 14 Integer-valued resource-sharing ratio m . If the resource-sharing ratio m has an integer value, state m has zero drift, or, more precisely, $R_m = 0$. Therefore R is singular. In this situation the Kolmogorov forward equations consist of n_{max} differential equations and 1 supplementary algebraic equation. This algebraic equation results from the state with drift zero and hinders obtaining the eigenvalue-eigenvector pairs. Observe that the state with zero drift does not influence the workload distribution and is basically *redundant*. In Appendix A.1 of [18], Mitra proposes how to reduce the dimension of the system of differential equations by 1 to (obtain a proper system), by eliminating the redundant algebraic equation. Further, it is proven that the eigenvalues of the reduced form coincide with the original form and it is shown how the eigenvectors of the original system are obtained from the reduced system. \diamond

A.2 Fluid queue with feedback

Here we consider the model of Section 2.1 which includes feedback of the workload W_t at the relay node. As remarked before, the number of source nodes no longer constitutes a Markov chain. We are interested in the stationary buffer workload denoted by $G_n(x) := \mathbb{P}(W \leq x; N = n)$. Let random variable B (I , respectively) denote a busy (idle) period in the system with feedback, and \vec{P} the distribution of N at the end of a busy period.

Note that the distributions \vec{P} and $\vec{\bar{P}}$ are identical, and also the busy periods B and \bar{B} have the same distribution. Hence,

$$\mathbb{P}(W \leq x; N = n | W > 0) = \mathbb{P}(\bar{W} \leq x; \bar{N} = n | \bar{W} > 0).$$

As a consequence, the stationary distribution $G_n(x)$ of the buffer workload and number of source nodes is

$$\begin{aligned} G_n(x) &= \mathbb{P}(\bar{W} \leq x; N = n | \bar{W} > 0) \mathbb{P}(W > 0) + \mathbb{P}(W = 0; N = n) \\ &= \frac{F_n(x) - \sum_k F_k(0)}{1 - \sum_k F_k(0)} \mathbb{P}(W > 0) + \mathbb{P}(W = 0; N = n). \end{aligned} \quad (\text{A.3})$$

To complete (2) we require expressions for $\mathbb{P}(W > 0)$ and $\mathbb{P}(W = 0; N = n)$. Here $\mathbb{P}(W > 0)$ follows from

$$\mathbb{P}(W > 0) = \frac{\mathbb{E}B}{\mathbb{E}I + \mathbb{E}B}.$$

Also

$$\mathbb{E}I = \left\langle -\vec{P}(Q_{DD})^{-1}, \vec{1} \right\rangle,$$

cf. Equation (A.2), where Q_{DD} is the square generator matrix of dimension n_+ for the states with downwards drift in case $W_t = 0$, i.e.,

$$Q_{DD}(i, j) := \begin{cases} \lambda & \text{if } j = i + 1, \\ c/2 & \text{if } j = i - 1, \\ 0 & \text{otherwise.} \end{cases}$$

Further, $\mathbb{P}(W = 0; N = n)$ corresponds to the n -th element of $-\vec{P}(Q_{DD})^{-1}$. Finally, the stationary distribution π of the number of active source nodes N_t follows from Equation (A.3), as $\pi_n = G_n(\infty)$.

B Proof of Proposition 7

This proof is essentially along the lines of [5]. Recall that $Y_n(\tau)$ denotes the conditional buffer delay, i.e., the time required by the relay node with resource-sharing ratio m to serve an amount of work τ if n other jobs are present upon the start. By $\phi_n(s)$ we denote the Laplace transform of $\mathbb{E}Y_n(\tau)$, i.e.,

$$\phi_n(s) = \int_0^\infty e^{-s\tau} \mathbb{E}Y_n(\tau) d\tau.$$

We obtain an expression for $\phi_n(s)$ by conditioning on the next possible event, namely: flow arrival, source departure of one of the n source nodes, or the relay node completes the service of amount τ . First we define $\nu_n := \lambda_n + \mu_n$, with

$$\lambda_n := \lambda \mathbf{1}\{n < n_{max}\}; \quad \mu_n := \frac{c}{f} \frac{n}{m+n}.$$

Then we obtain the following expression:

$$\begin{aligned} \phi_n(s) = & \int_0^\infty e^{-s\tau} \int_{\tau \cdot \frac{m+n}{m}}^\infty v_n e^{-v_n t} \tau \cdot \frac{m+n}{m} dt d\tau + \\ & \int_0^\infty e^{-s\tau} \int_0^{\tau \cdot \frac{m+n}{m}} v_n e^{-v_n t} \left\{ t + \frac{\lambda_n}{v_n} \mathbb{E}Y_{n+1} \left(\tau - \frac{m}{m+n} t \right) \right. \\ & \left. + \frac{\mu_n}{v_n} \mathbb{E}Y_{n-1} \left(\tau - \frac{m}{m+n} t \right) \right\} dt d\tau. \end{aligned}$$

After elementary algebra, this is rewritten as the following system of linear equations:

$$\frac{1}{s} = -\lambda_n \phi_{n+1}(s) + \left(\frac{m}{m+n} s + v_n \right) \phi_n(s) - \mu_n \phi_{n-1}(s), \quad (\text{B.1})$$

or, in matrix notation, $\vec{1} = s \cdot M(s) \overrightarrow{\phi}(s)$ where $M(s) := -\bar{Q} + sR$ with \bar{Q} as in (A.1) and

$$R := \text{diag} \left\{ 1, \frac{m}{m+1}, \frac{m}{m+2}, \dots, \frac{m}{m+n_{\max}} \right\}.$$

It is readily verified that the equation $\det M(s) = 0$ coincides with $\det(R^{-1}Q - sI) = 0$. In other words: the roots of $\det M(s) = 0$ are the eigenvalues of $R^{-1}Q$. As Q is singular, one of the eigenvalues of $R^{-1}Q$ is 0, say s_0 . Further, a straightforward application of ‘Geršgorin’ yields that all eigenvalues $s_0, \dots, s_{n_{\max}}$ are real, non-positive and unique.

The Laplace transform ϕ_n can be solved from the linear system by applying Cramer’s rule to $s\overrightarrow{\phi}(s) = (M(s))^{-1}\vec{1}$, i.e.,

$$s\phi_n(s) = \frac{\det M_{-n}(s)}{\det M(s)}. \quad (\text{B.2})$$

where $M_{-n}(s)$ is defined as $M(s)$ with the n -th column replaced by $\vec{1}$. The denominator of the right-hand side of (B.2) is a polynomial of degree $n_{\max} + 1$ in s . The above considerations entail

$$s\phi_n(s) = \frac{A_{(n,0)}}{s} + \sum_{j=1}^{n_{\max}} \frac{A_{(n,j)}}{s - s_j}$$

where the constants $A_{(n,j)}$ follow from the partial-fraction expansion at $s_0, \dots, s_{n_{\max}}$. Then,

$$\phi_n(s) = \frac{A_{(n,0)}}{s^2} + \sum_{j=1}^{n_{\max}} \frac{A_{(n,j)}}{s_j} \frac{1}{s - s_j} - \sum_{j=1}^{n_{\max}} \frac{A_{(n,j)}}{s_j} \frac{1}{s}.$$

Finally, inverting the individual parts gives the desired result:

$$\mathbb{E}Y_n(\tau) = A_{(n,0)}\tau + \sum_{j=1}^{n_{max}} \frac{A_{(n,j)}}{s_j} e^{s_j\tau} - \sum_{j=1}^{n_{max}} \frac{A_{(n,j)}}{s_j}.$$