



HAL
open science

Lingering issues in distributed scheduling

Florian Simatos, Niek Bouman, Sem Borst

► **To cite this version:**

Florian Simatos, Niek Bouman, Sem Borst. Lingering issues in distributed scheduling. *Queueing Systems*, 2014, 77 (2), pp.243-273. 10.1007/s11134-014-9404-z . hal-01888052

HAL Id: hal-01888052

<https://hal.science/hal-01888052>

Submitted on 4 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive Toulouse Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of some Toulouse researchers and makes it freely available over the web where possible.

This is an author's version published in: <https://oatao.univ-toulouse.fr/20905>

Official URL : <https://doi.org/10.1007/s11134-014-9404-z>

To cite this version :

Simatos, Florian and Bouman, Niek and Borst, Sem Linger issues in distributed scheduling. (2014) Queueing Systems, 77 (2). 243-273. ISSN 0257-0130

Any correspondence concerning this service should be sent to the repository administrator:

tech-oatao@listes-diff.inp-toulouse.fr

Lingering issues in distributed scheduling

Florian Simatos · Niek Bouman · Sem Borst

Abstract Recent advances have resulted in queue-based algorithms for medium access control which operate in a distributed fashion, and yet achieve the optimal throughput performance of centralized scheduling algorithms. However, fundamental performance bounds reveal that the “cautious” activation rules involved in establishing throughput optimality tend to produce extremely large delays, typically growing exponentially in $1/(1-\rho)$, with ρ the load of the system, in contrast to the usual linear growth. Motivated by that issue, we explore to what extent more “aggressive” schemes can improve the delay performance. Our main finding is that aggressive activation rules induce a lingering effect, where individual nodes retain possession of a shared resource for excessive lengths of time even while a majority of other nodes idle. Using central limit theorem type arguments, we prove that the idleness induced by the lingering effect may cause the delays to grow with $1/(1-\rho)$ at a quadratic rate. To the best of our knowledge, these are the first mathematical results illuminating the lingering effect and quantifying the performance impact. In addition extensive simulation experiments are conducted to illustrate and validate the various analytical results.

This work first appeared in [25] in the proceedings of the ACM/SIGMETRICS 2013 conference. The two papers mainly differ in their organization. The present paper also contains additional technical details, especially in the proofs of Lemmas 4.2 and 4.5.

F. Simatos · N. Bouman (✉) · S. Borst
Eindhoven University of Technology, Eindhoven, The Netherlands
e-mail: n.bouman@tue.nl

F. Simatos
e-mail: f.simatos@tue.nl

S. Borst
e-mail: sem@win.tue.nl

S. Borst
Alcatel-Lucent Bell Labs, Murray Hill, NJ, USA

Keywords Delay performance · Distributed scheduling algorithms · Heavy traffic scaling · Wireless networks

Mathematics Subject Classification 68M20 · 90B18 · 90B22

1 Context and illustrative example

1.1 Context

As networks continue to grow in size and complexity, they increasingly rely on scheduling algorithms for efficient allocation of shared resources and arbitration between users. As a result, the design and analysis of scheduling algorithms for complex network scenarios has attracted significant attention over the last several years. One of the centerpieces in the scheduling literature is the celebrated MaxWeight algorithm as proposed in the seminal work [29, 30]. The MaxWeight algorithm provides throughput optimality and maximum queue stability in a variety of scenarios, and has emerged as a powerful paradigm in cross-layer control and resource allocation problems [8].

While not strictly optimal in terms of delay performance, MaxWeight algorithms do achieve so-called equivalent workload minimization and offer favorable scaling characteristics in heavy-traffic conditions [22, 26]. As a further key appealing feature, MaxWeight algorithms only need information on the queue lengths and instantaneous service rates, and do not rely on any explicit knowledge of the underlying system parameters. On the downside, solving the maximum-weight problem tends to be challenging and potentially NP-hard. This is exacerbated in a network setting, where a centralized control entity may be lacking or require global state information, creating a substantial communication overhead in addition to the computational burden. This concern is especially pertinent as the maximum-weight problem needs to be solved at a high pace, commensurate with the fast time scale on which scheduling algorithms typically need to operate.

This issue has provided a strong impetus for devising algorithms that entail *lower computational complexity and communication overhead*, but retain the *maximum stability and throughput guarantees* of the MaxWeight algorithm. Various approaches in that direction were proposed in [3, 15, 23, 24, 28, 31]. An exciting breakthrough in this quest was recently achieved in the design of random back-off schemes for wireless medium access control that seem to offer the best of both worlds. These schemes operate in a *distributed fashion, requiring no centralized control entity or global state information*, and yet, remarkably, provide the *capability of achieving throughput optimality and maximum stability*. More specifically, clever algorithms have been developed for finding the back-off rates that yield any given target throughput vector in the capacity region [11, 14]. In the same spirit, powerful algorithms have been devised for adapting the back-off rates based on queue length information, and been shown to guarantee maximum stability [13, 17, 21].

While the maximum-stability guarantees for the above-mentioned algorithms have strong appeal, the picture becomes different when we consider performance metrics such as expected queue lengths or delays. In [20], it is shown that low-complexity

schemes cannot be expected to achieve low delay in arbitrary topologies (unless P equals NP), roughly meaning that the required number of calculations prior to any transmission is super-polynomial in the number of nodes. However, this notion of delay is a transient one, and it is not exactly clear what the implications are for expected queue lengths or delays in specific networks, if any.

Performance upper bounds in [10, 12] show that for sufficiently low load the delay only grows polynomially with the number of nodes in bounded-degree interference graphs. Similar upper bounds are presented in [27] for bounded-degree conflict graphs with low load.

Performance lower bounds in [2] indicate that the “cautious” back-off functions involved in establishing maximum stability tend to produce extremely large delays, typically growing *exponentially* in $1/(1-\rho)$, with ρ the load of the system, in contrast to the usual *linear* growth. More specifically, the bounds show that the expected queue lengths grow as $\psi^{-1}(1-\rho)$ as $\rho \uparrow 1$. Here ψ^{-1} represents the inverse of the (decreasing) function ψ , specifying the probability of a node entering a back-off as a function of its current queue length. The bounds may be explained by noting that the queue lengths govern the fraction of back-off time through the function ψ . Since the fraction of back-off time cannot exceed the surplus capacity in order for the system to be stable, however, it is ultimately the amount of surplus $1-\rho$ that dictates the queue lengths through the function ψ^{-1} . We note that maximum stability has been established under the condition that the function $\psi(a)$ decays (no faster than) *inverse-logarithmically* as $a \rightarrow \infty$, i.e., $\psi(a) \sim 1/\log a$. This entails that $\psi^{-1}(s)$ grows (no slower than) *exponentially* in $1/s$ as $s \downarrow 0$, yielding the stated exponential growth of $\psi^{-1}(1-\rho)$ in $1/(1-\rho)$ as $\rho \uparrow 1$.

The above lower bounds suggest that the delay performance may be improved when the function ψ decays faster, e.g., *inverse-polynomially*: $\psi(a) \sim a^{-\beta}$, with $\beta > 0$, so that $\psi^{-1}(s) \sim s^{-1/\beta}$ as $s \downarrow 0$. The larger the value of the exponent β , the slower the growth of $\psi^{-1}(1-\rho) \sim (1-\rho)^{-1/\beta}$ as $\rho \uparrow 1$. In particular, it might seem plausible that for $\beta \geq 1$, the expected queue lengths will only exhibit the usual linear growth in $1/(1-\rho)$ as $\rho \uparrow 1$. Note that a larger value of β means that a node is more “aggressive”, in the sense that it is less likely to enter a back-off and more inclined to hold on to the medium, and hence the coefficient β will be referred to as the aggressiveness parameter. It is worth mentioning that maximum stability for the above back-off functions is not guaranteed by existing results, which do not apply for any $\beta > 0$. In fact, for $\beta > 1$, maximum stability has been shown not to hold in certain topologies [9].

In the present paper we aim to gain fundamental insight into the extent that a larger aggressiveness parameter can improve the delay performance. Our main finding is that, for large values of β , a *lingering effect* can cause the mean stationary delay to increase in heavy traffic as $1/(1-\rho)^2$. In the *infinitely persistent* case, $\beta = +\infty$, we present a complete analysis and prove that the heavy-traffic stationary delay scales as $1/(1-\rho)^2$. In the *strongly persistent* case, $\beta \in (1, \infty)$, we show that the system’s behavior is roughly similar to the behavior in the infinitely persistent case, supporting our claim that the delay scales as $1/(1-\rho)^2$ in this case as well. In the *weakly persistent* case, $\beta < 1$, fundamentally different behavior may emerge, requiring significant additional developments which go beyond the scope of the present paper.

For transparency, we focus on the simplest topology where the lingering effect occurs. This topology, described in later sections, may at first sight seem restrictive in view of recent results [13, 17, 21] which apply to general topologies. We believe, however, that our results give insight into more general situations as will be discussed in Sect. 6.

1.2 Illustrative example

Consider a network consisting of four queues which are split into two groups, in such a way that if any queue of one group is transmitting a packet, no queue of the other group may transmit, and vice versa. A group is said to be *active* if one of its queues is transmitting a packet, and a queue is said to be active if it belongs to the active group. The other group and queues are said to be *inactive*.

In our model, time is slotted and active queues adhere to the following algorithm: after each transmission, each active queue flips a coin and *advertizes a release* with probability $(1 + a)^{-\beta}$, with a the number of packets that this queue has to transmit and $\beta > 0$. If the two active queues advertize a release simultaneously, then active queues become inactive and vice versa: such a time is called a *switching time*. This simple distributed algorithm gives rise to dynamics as illustrated in Fig. 1, where the system is considered over three consecutive switching times t_1, t_2 , and t_3 . Between switching times, the packet buffers at the active queues are drained, while packets accumulate at the inactive queues. The dynamics shown in Fig. 1 are representative of the case $\beta > 1$ where a switch does not occur until both active queues are close to being empty, see Theorem 4.1 and Corollary 5.5.

Let us now give a flavor of the lingering effect. Imagine that the two active queues start with initial queue lengths of the same order, say Q . As just mentioned, queues retain the shared resource until the time T^* at which both queues are close to being empty, thus preventing other queues from activating until this time. The law of large numbers suggests that T^* is of order Q (i.e., active queues are drained linearly as in

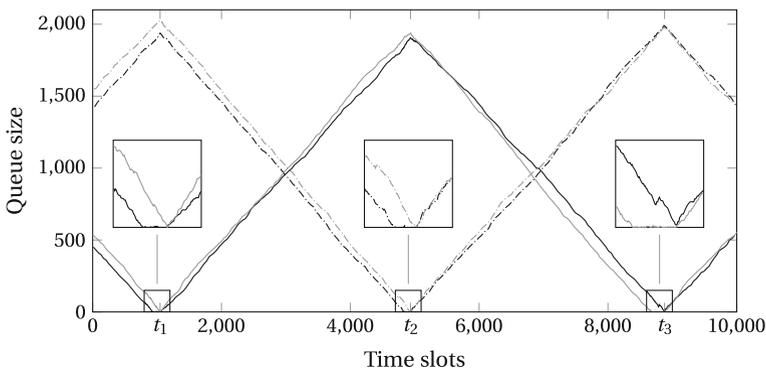


Fig. 1 A sample path on the normal time scale with $\beta = 2$, representative of the case $\beta > 1$. The three boxes zoom in to show the lingering effect. One queue hovers around zero while the other queue is yet to empty, resulting in an inefficient use of the resource

Fig. 1), but the central limit theorem suggests that up to time T^* , the first active queue to have been completely drained will be empty of the order of \sqrt{Q} units of time, while waiting for the other queue to empty. This lingering effect is illustrated in Fig. 1, which will be explained in greater detail in Sect. 3.

This leads to a fraction of the slots of the order of $(1/\sqrt{Q})$ where the shared resource is used inefficiently. This may at first sight seem negligible when Q is large, and indeed queues seem to empty at the same time on the coarse time scale of Fig. 1. However, we will actually establish in Sects. 4 and 5 that it has a significant impact in heavy traffic, causing queue lengths to grow at a rate $1/(1-\rho)^2$ as $\rho \uparrow 1$, instead of the optimal $1/(1-\rho)$.

1.3 Organization

The remainder of the paper is organized as follows. In Sect. 2 we present a detailed model description. In Sect. 3 we provide an informal discussion of the lingering effect and explain how it leads to a growth rate $1/(1-\rho)^2$ as $\rho \uparrow 1$ of the mean stationary delay. Sect. 4 is devoted to the proof of Theorem 4.1, which is the main theoretical result of the paper, and proves this quadratic growth rate in the infinitely persistent case $\beta = +\infty$. In Sect. 5 we present arguments supporting the conjecture that this quadratic growth occurs for every $\beta \in (1, \infty)$, and not just for $\beta = +\infty$. Finally, we conclude in Sect. 6 with some remarks and suggestions for further research.

2 Model description

2.1 Informal description

Let us now give a more precise definition of our model. As mentioned in the introduction, in order to analyze the lingering effect in the simplest possible setting, we focus on a symmetric system consisting of two groups of $R \geq 2$ queues. At any given point in time, one of the groups is *active* while the other is *inactive*.

Time is slotted, and inactive queues have simple dynamics, driven by independent and identically distributed numbers of packet arrivals in each slot, so that they each simply grow according to random walks with step size distribution denoted by ξ . During each time slot, active queues increase by independent amounts distributed as ξ as well, but, if at least one packet is present at the start of the slot, then an active queue also flushes exactly one packet.

Moreover, at the end of each time slot, each active queue tosses a coin and advertizes a *momentary release* with probability $\psi(a)$, with a the number of packets in the queue at the end of the time slot. This (momentary) release gives inactive queues an opportunity to become active: if all the active queues *simultaneously* advertize a release, then inactive queues become active, and active queues become inactive. Such a time is called a *switching time*. We will assume in the sequel that $\psi(a) = (1+a)^{-\beta}$ for some parameter $\beta > 0$, called the aggressiveness parameter. In particular, $\psi(a) \rightarrow 0$ as $a \rightarrow +\infty$, and so active queues are less likely to advertize a release when they are highly loaded; this mechanism thus gives priority to highly loaded queues in a

distributed fashion. Finally, there is a cost associated with advertizing a release: each time an active queue advertizes a release and is not empty, it incurs an additional increase distributed according to some random variable ζ .

This model qualitatively resembles the canonical models for queue-based medium access control mechanisms [2, 16, 17, 19, 21]. The main difference with these models is that in our model, back-off periods are infinitesimally short (hence, the releases are qualified as momentary), and the jump size ζ represents the number of packets that would have arrived during a non-zero back-off period. In contrast to the discrete-time model in [16, 19], our back-off model shares the common feature of continuous-time models that there are no collisions. This back-off model simplifies the analysis and is possible because of the simple topology considered in our paper. The key advantage of using a discrete-time model with synchronized queues is that it avoids the complication of requiring back-off periods to overlap in order to define a switching time. Moreover, simulation experiments show that our main results extend to models with non-infinitesimal back-off periods and also behave qualitatively similarly as a continuous-time model.

2.2 Parameters and heavy-traffic regime

The model described in the previous subsection is defined through four parameters: the number $R \geq 2$ of queues in each group, two integer-valued random variables $\xi, \zeta \in \mathbb{N} = \{0, 1, \dots\}$ and a real number $\beta \in (0, \infty]$, which defines the $[0, 1]$ -valued sequence $(\psi(a), a \geq 0)$ via $\psi(a) = (1 + a)^{-\beta}$ for $a \geq 0$, to be understood for $\beta = +\infty$ as $\psi(0) = 1$ and $\psi(a) = 0$ for $a > 0$. Note that only the asymptotic behavior of ψ matters, and our results could easily be extended to any ψ with $a^\beta \psi(a) \rightarrow \ell \in (0, \infty)$ as $a \rightarrow +\infty$ (when β is finite).

We assume that ξ and ζ have finite means, respectively $\mathbb{E}\xi = \rho/2$ and $\mathbb{E}\zeta = z$, and that ξ has finite variance denoted by $v = \mathbb{E}[(\xi - \rho/2)^2]$. It will be argued that the system is stable if and only if $\rho < 1$, and so we will refer to ρ as the load of the system. Note that by symmetry, each queue is active half the time, which explains the factor 2 in the definition $\rho = 2\mathbb{E}\xi$ of ρ .

In the sequel we will be interested in a heavy-traffic regime where the load ρ of the system increases to the critical value 1. Although we will not make this dependency explicit, we think of the random variable ξ as depending on ρ , i.e., we have a family of random variables $\{\xi_\rho, \rho > 0\}$ with $\mathbb{E}\xi_\rho = \rho/2$. With this in mind, we make a final assumption on the ξ 's: we assume that their second moment is uniformly bounded in ρ , i.e., $\sup_\rho \mathbb{E}\xi_\rho^2 < +\infty$.

2.3 Formal description

Because of the symmetry of the system, we do not need to label the queues individually, but only need to keep track of the state of active and inactive queues. We will consider the system embedded at switching times, and define $Q_r^a(k)$ and $Q_r^i(k)$ as the numbers of packets in the r th active and inactive queue, respectively, just after the k th switch occurred. We will be interested in the Markov chain $(\mathbf{Q}(k), k \geq 0)$ which we also write

as $\mathbf{Q} = (\mathbf{Q}^a, \mathbf{Q}^i)$ with $\mathbf{Q}^a = (Q_r^a, r \in \mathcal{R})$, $\mathbf{Q}^i = (Q_r^i, r \in \mathcal{R})$, $\mathcal{R} = \{1, \dots, R\}$, and we reserve in the sequel bold notation for vectors (of functions or numbers).

As informally described in Sect. 2.1, the dynamics of \mathbf{Q} in between two switching times are governed by two R -dimensional processes $\mathbf{S} = (S_r, r \in \mathcal{R})$ and $\mathbf{A} = (A_r, r \in \mathcal{R})$: \mathbf{S} gives the increments of the inactive queues, while \mathbf{A} gives the state of the active queues. The dynamics are as follows:

- The $2R$ processes A_r, S_r are independent;
- For each $r \in \mathcal{R}$, $(S_r(k), k \geq 0)$ is a random walk with step size distribution ξ started at 0;
- For each $r \in \mathcal{R}$, $(A_r(k), k \geq 0)$ is a space-inhomogeneous random walk with the following dynamics: for any $a \in \mathbb{N}$ and any function $f : \mathbb{N} \rightarrow [0, \infty)$, we have

$$\mathbb{E}[f(A_r(1)) \mid A_r(0) = a] = \mathbb{E}\left[f\left(Y(a) + \xi \mathbb{1}_{\{Y(a) > 0, U < \psi(Y(a))\}}\right)\right], \quad (1)$$

where $Y(a) = a + \xi - \mathbb{1}_{\{a > 0\}}$ is the number of packets at the end of the time slot, U is uniformly distributed in $[0, 1]$, and U, ξ , and ζ are independent.

Equation (1) describes the dynamics of an active queue and can be interpreted as follows. At each time slot, an active queue increases by ξ , and if not empty at the beginning of the time slot, flushes a packet, which brings the queue from state a to state $Y(a)$. If $U < \psi(Y(a))$, then we say that the active queue advertizes a release, which thus happens with the (conditional) probability $\psi(Y(a))$ as described in Sect. 2.1. If at the end of the time slot the queue is not empty and it advertizes a release, i.e., $Y(a) > 0$ and $U < \psi(Y(a))$, then the active queue also incurs an additional increase by ζ .

To define the $2R$ -dimensional process $\mathbf{Q} = (\mathbf{Q}^a, \mathbf{Q}^i)$ from \mathbf{S} and \mathbf{A} , it remains to adopt notation for the switching time, which we denote by T^* . Thus T^* is the first time at which all active queues advertize a release at the same time. Note that T^* and \mathbf{S} are independent. With these definitions, the dynamics of \mathbf{Q} as informally described in Sect. 2.1 obey the following equation: for any $\mathbf{q} = (\mathbf{q}^a, \mathbf{q}^i) \in \mathbb{N}^R \times \mathbb{N}^R$ and any function $f : \mathbb{N}^{2R} \rightarrow [0, \infty)$,

$$\mathbb{E}[f(\mathbf{Q}(1)) \mid \mathbf{Q}(0) = \mathbf{q}] = \mathbb{E}\left[f(\mathbf{q}^i + \mathbf{S}(T^*), \mathbf{A}(T^*)) \mid \mathbf{A}(0) = \mathbf{q}^a\right]. \quad (2)$$

The special case $\beta = +\infty$ will be of particular importance. Indeed, it can be analyzed exhaustively and we will show in Sect. 5 that it is representative of the system's behavior in the range $\beta > 1$. The fluid limits of the system for a fixed value of ρ and $\beta = +\infty$ have been studied in [5], where the algorithm is referred to as the random capture algorithm. Also, when $\beta = +\infty$, active queues only advertize a release when they are empty (in which case there is no additional term ζ), and so in this case we have $\mathbf{A}(T^*) = \mathbf{0}$ and $T^* = \inf\{k \geq 0 : \mathbf{A}(k) = \mathbf{0}\}$.

2.4 Notation for probabilities and expectations

Similarly as we have just done in (2), we will use in the remainder of the paper the common symbol \mathbb{E} to denote expectation with respect to the laws of \mathbf{Q} and (\mathbf{A}, \mathbf{S}) .

Initial conditions will be denoted by a subscript, and it should always be clear from the context whether we consider initial conditions of \mathbf{Q} , \mathbf{A} , or some A_r (remember that $\mathbf{S}(0)$ is always equal to $\mathbf{0}$). For instance, (1) and (2) can be rewritten as follows:

$$\mathbb{E}_a [f(A_r(1))] = \mathbb{E} [f(Y(a) + \zeta \mathbf{1}_{\{Y(a) > 0, U < \psi(Y(a))\}})]$$

and

$$\mathbb{E}_{\mathbf{q}} [f(\mathbf{Q}(1))] = \mathbb{E}_{\mathbf{q}^a} [f(\mathbf{q}^i + \mathbf{S}(T^*), \mathbf{A}(T^*))].$$

The probability distributions corresponding to these various expectations are written as \mathbb{P}_a , \mathbb{P} , $\mathbb{P}_{\mathbf{q}}$, and $\mathbb{P}_{\mathbf{q}^a}$. We also define \mathbb{P}_{∞} with corresponding expectation \mathbb{E}_{∞} as the laws of \mathbf{Q} and (\mathbf{A}, \mathbf{S}) started in the stationary distribution of \mathbf{Q} , provided \mathbf{Q} is positive recurrent.

When $\beta = +\infty$, we see based on (2) and the fact that $\mathbf{A}(T^*) = \mathbf{0}$ that $\mathbf{Q}^i(k) = \mathbf{0}$ for $k \geq 1$. In particular, when \mathbf{Q} is positive recurrent, $\mathbf{Q}^i(0)$ is \mathbb{P}_{∞} -almost surely equal to $\mathbf{0}$ and (2) therefore becomes

$$\mathbb{E}_{\infty} [f(\mathbf{Q}^a(0))] = \mathbb{E}_{\infty} [f(\mathbf{S}(T^*))] \quad (\beta = +\infty). \quad (3)$$

2.5 Additional notation τ_r , τ_{\max} , $\tau_{(r)}$, $|\cdot|$, and $\|\cdot\|$

In the remainder of the paper we define $\tau_r = \inf\{k \geq 0 : A_r(k) = 0\}$ as the time at which the r th active queue hits 0, $\tau_{\max} = \max_{r \in \mathcal{R}} \tau_r$ as the largest time at which an active queue hits 0 for the first time, and we let the $\tau_{(r)}$'s be the order statistics of the τ_r 's, i.e., $\tau_{(1)} \leq \dots \leq \tau_{(R)}$ and $\{\tau_{(r)}\} = \{\tau_r\}$ (in particular $\tau_{\max} = \tau_{(R)}$).

Let $|\cdot|$ be the L_{∞} norm and $\|\cdot\|$ be the L_1 norm, i.e., if $J \geq 1$ and $\mathbf{x} \in \mathbb{R}^J$ then $|\mathbf{x}| = \max_j |x_j|$ (which is just the absolute value for $J = 1$) and $\|\mathbf{x}\| = |x_1| + \dots + |x_J|$.

2.6 Connection with polling systems

It is worth emphasizing that we restrict the investigation in the next sections to the case of $R \geq 2$ queues in each group, and exclude the case $R = 1$ from the analysis. Indeed, the lingering effect that we intend to investigate only occurs when there are $R \geq 2$ queues in the same group.

The case $R = 1$ may be interpreted as a single-server two-class queueing system, where the server may switch from one class to the other after each service completion with a probability that depends on the queue length of the class that is currently being served. As a somewhat unusual feature, the queue length of the latter class increases by the random variable ζ when a switch occurs. This bears some resemblance with a

polling system, when ζ is viewed as the number of arrivals during a switch-over time. It is worth observing that this switching rule does not belong to the class of branching-type service disciplines which yield tractable joint queue length distributions in polling models. The only exception is the special case $\beta = +\infty$, which corresponds to the exhaustive service discipline in polling models (with a process-level heavy-traffic analysis in [4]) and the so-called random capture algorithm in [5]. In general, however, the analysis of the joint queue length process appears far from trivial, even when $\zeta = 0$, although the aggregate queue length distribution is then fairly easy to obtain. This connection with polling systems is actually at the heart of the proof of forthcoming Lemma 4.5.

In case $R \geq 2$, the system may in the same vein be interpreted as a set of R single-server two-queue polling systems, where servers are only allowed to switch between queues in a synchronized fashion. Such models with simultaneous service of several queues and synchronized switches are natural in applications, and indeed our model has already been studied in Boon et al. [1] to model traffic lights at intersections. The results in [1] are in some sense complementary to ours, since the authors perform a heavy-traffic analysis when the system's parameters are such that the lingering effect does not play a role. In particular, they study cases where the delay scales as $1/(1-\rho)$.

These models show similarity with multiple-server polling systems, which did receive some attention, but have largely defied analysis. This offers testimony of the mathematical complexity of an exact queueing analysis of the model under consideration, and provides justification for an asymptotic investigation.

3 Informal discussion of the lingering effect

In Sect. 5.1 we will prove that, when $\beta > 1$, an active queue only advertizes releases when it is close to being empty. In particular, once a queue gains possession of the resource, it holds onto it, even when some or all of the other queues in the same group are empty, and it would be more efficient for the queues in the other group to receive the resource. This causes a lingering effect as discussed in Sect. 1.2 and illustrated in Fig. 1 for a scenario with $R = 2$, $\beta = 2$.

It may appear that the two queues in the same group drain around the same time (as can indeed be shown to be the case on a “fluid scale”). When we zoom in, however, we see that there is actually a time period where one of the queues is already empty, while the other one clings to the resource and prevents the two queues in the other group from activating. In this section we give a heuristic explanation of how this inefficient use of the resource leads to a quadratic growth of the mean stationary delay in heavy traffic. This explanation is also aimed at developing intuition and introducing the structure of the proofs in the next section.

Now consider a regime in which active queues only advertize a release when they are close to empty, i.e., $\mathbf{A}(T^*) \approx \mathbf{0}$. Applying (2) with $f(\mathbf{q}) = \|\mathbf{q}^a\|$, we obtain

$$\mathbb{E}_{\mathbf{q}}(\|\mathbf{Q}^a(1)\|) = \|\mathbf{q}^i\| + \mathbb{E}_{\mathbf{q}}(\|\mathbf{S}(T^*)\|).$$

Thus in stationarity, we have

$$\begin{aligned}\mathbb{E}_\infty(\|\mathbf{Q}^a(0)\|) &= \mathbb{E}_\infty(\|\mathbf{Q}^a(1)\|) = \mathbb{E}_\infty(\|\mathbf{Q}^i(0)\|) + \mathbb{E}_\infty(\|\mathbf{S}(T^*)\|) \\ &= \mathbb{E}_\infty(\|\mathbf{A}(T^*)\|) + \mathbb{E}_\infty(\|\mathbf{S}(T^*)\|)\end{aligned}$$

the last equality following by applying (2) in stationarity and with $f(\mathbf{q}) = \|\mathbf{q}^i\|$. In particular, since $\mathbf{A}(T^*) \approx \mathbf{0}$, it follows that $\mathbb{E}_\infty(\|\mathbf{Q}^a(0)\|) \approx \mathbb{E}_\infty(\|\mathbf{S}(T^*)\|)$.

Since by definition $\mathbf{Q}^a(0) = \mathbf{A}(0)$ and $\|\mathbf{S}(\cdot)\|$ is a random walk with drift $R\mathbb{E}\xi$ independent of T^* , we obtain $\mathbb{E}_\infty(\|\mathbf{A}(0)\|) \approx R\mathbb{E}(\xi)\mathbb{E}_\infty(T^*)$ and so by symmetry,

$$\mathbb{E}_\infty(A_1(0)) \approx \mathbb{E}(\xi)\mathbb{E}_\infty(T^*). \quad (4)$$

The goal is now to relate T^* to $A_1(0)$. Remember that active queues only advertize a release when they are close to empty; moreover, active queues are stable and so once all active queues are close to 0, it is only a matter of constant time for them to simultaneously advertize a release. This suggests that the switching time should occur around the largest time at which an active queue empties, i.e., this suggests the approximation $T^* \approx \tau_{\max}$ (recall that τ_{\max} and the τ_r 's have been defined in Sect. 2.5). The law of large numbers combined with the central limit theorem show that $\tau_r \approx A_r(0)/(1 - \mathbb{E}\xi) + A_r(0)^{1/2}$ (where we neglect multiplicative constants, possibly random, appearing in front of first- or second-order terms and that do not influence the order of magnitude of the final result), which leads to the approximation $\tau_{\max} \approx |\mathbf{A}(0)|/(1 - \mathbb{E}\xi) + |\mathbf{A}(0)|^{1/2}$. Since under \mathbb{P}_∞ queues are symmetric, we have $|\mathbf{A}(0)| \approx A_1(0) + A_1(0)^{1/2}$ which finally leads to $T^* \approx A_1(0)/(1 - \mathbb{E}\xi) + A_1(0)^{1/2}$, i.e., going back to (4),

$$\mathbb{E}_\infty(A_1(0)) \approx \frac{\mathbb{E}\xi}{1 - \mathbb{E}\xi} \mathbb{E}_\infty(A_1(0)) + \mathbb{E}_\infty[A_1(0)^{1/2}].$$

Thus upon a concentration-like result of the kind $\mathbb{E}_\infty[A_1(0)^{1/2}] \approx [\mathbb{E}_\infty(A_1(0))]^{1/2}$ it is reasonable to expect

$$\left(1 - \frac{\mathbb{E}\xi}{1 - \mathbb{E}\xi}\right) \mathbb{E}_\infty(A_1(0)) \approx [\mathbb{E}_\infty(A_1(0))]^{1/2}.$$

Since $1 - \mathbb{E}(\xi)/(1 - \mathbb{E}\xi) \approx 1 - \rho$ this shows that $\mathbb{E}_\infty(A_1(0))$, and hence $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$, should grow as $1/(1 - \rho)^2$. While admittedly crude, the above heuristic arguments provide the correct estimates, and serve as a useful guide for a rigorous proof in Sect. 4.

As reflected in the above computations, the square factor really stems from the relation $T^* \approx \tau_{(1)} + |\mathbf{A}(0)|^{1/2}$, i.e., T^* occurs somehow long after $\tau_{(1)}$, the time at which it would be optimal to switch in order to avoid inefficient use of the resource. But it is difficult to make the system switch exactly at $\tau_{(1)}$ in a distributed fashion, and here the penalty incurred is a square root. Interestingly, the penalty may seem negligible but this small inefficiency has a significant impact in heavy traffic.

When $\beta = +\infty$ the above heuristic arguments can be made rigorous, and a complete proof is provided in the next section. When $\beta \in (1, \infty)$, we will prove in Sect. 5.1 that active queues indeed advertize a release only when they are close to being empty, thus justifying the above intuition. More formally, we will show in Proposition 5.1 and its Corollary 5.5 that the random variable $\mathbf{A}(T^*)$ converges weakly to a finite random variable as the initial state blows up. We therefore conjecture that the lingering effect will make the mean stationary delay scale like $1/(1-\rho)^2$ when $\beta > 1$, but a proof of that result may involve significantly more work than in the case $\beta = +\infty$. We leave this issue open for future research, and present in Sect. 5.2 extensive simulation results which support this conjecture.

4 Full investigation of the infinitely persistent case

In the infinitely persistent case $\beta = +\infty$, the system's performance and in particular the impact of the lingering effect can be analyzed rigorously. The main result of this section is given by the following theorem, which shows that the mean stationary delay grows quadratically in $1/(1-\rho)$ as $\rho \uparrow 1$.

Theorem 4.1 *If $\beta = +\infty$, then \mathbf{Q} is positive recurrent for $\rho < 1$ and*

$$0 < \liminf_{\rho \uparrow 1} \left[(1 - \rho)^2 \mathbb{E}_\infty (\|\mathbf{Q}(0)\|) \right] \leq \limsup_{\rho \uparrow 1} \left[(1 - \rho)^2 \mathbb{E}_\infty (\|\mathbf{Q}(0)\|) \right] < +\infty. \quad (5)$$

The rest of this section is devoted to the proof of Theorem 4.1, and so from now on we assume that $\beta = +\infty$. Remember that in this case we have $T^* = \inf\{k \geq 0 : \mathbf{A}(k) = \mathbf{0}\}$ and $\mathbf{A}(T^*) = \mathbf{0}$. Moreover, in this case active queues only advertize a release when they are empty, in which case they do not incur the additional arrivals given by ζ . In particular, active queues are independent random walks with step size distribution $\xi - 1$ and reflected at the origin. Since $\xi \geq 0$, these random walks belong to the class of skip-free random walks, i.e., random walks that only decrease by -1 . In particular, it is well-known and not difficult to show that $\mathbb{E}_a(\tau_1) = a/(1 - \mathbb{E}\xi)$ for any $a \in \mathbb{N}$, a fact that will be used several times in this section.

We prove Theorem 4.1 via a series of intermediate results that justify the various approximations of the previous section. We first prove that $T^* \approx \tau_{\max}$, i.e., the time at which the R independent random walks A_r simultaneously hit 0 is close to the first time at which each process has visited 0 at least once.

Lemma 4.2 *Let $\rho_0 < 2$: then $\sup_{\mathbf{a}, \rho} \mathbb{E}_{\mathbf{a}}(T^* - \tau_{\max})$ is finite, where the supremum is taken over $\mathbf{a} \in \mathbb{N}^R$ and $\rho \leq \rho_0$.*

Proof Fix any $\rho < 2$, so that the active queues are stable, and T^* is almost surely finite. Define the sequences $(\tau_{r,k}, k \geq 0)$ and $(\sigma_{\max,k}, k \geq 0)$ by induction on k as follows: for $k = 0$ set $\tau_{r,0} = \sigma_{\max,0} = 0$ and for $k \geq 0$,

$$\tau_{r,k+1} = \min\{i \geq \sigma_{\max,k} : A_r(i) = 0\} \text{ and } \sigma_{\max,k+1} = \max_{r \in \mathcal{R}} \tau_{r,k+1}.$$

Remember that (since $\beta = +\infty$) T^* is the first time at which all the queues are simultaneously empty, thus $T^* \geq \tau_{\max} = \sigma_{\max,1}$. If $T^* > \sigma_{\max,1}$, this means that at time $\sigma_{\max,1}$ at least one queue was not empty, and so T^* will be at least as large as the first time after time $\sigma_{\max,1}$ at which every queue will have visited 0 at least once, which is by definition $\sigma_{\max,2}$. Iterating this argument, we see that the sequence $(\sigma_{\max,k}, k \geq 1)$ is stationary (in the sense that $\sigma_{\max,k} = \sigma_{\max,k^*}$ for some finite k^* and all $k \geq k^*$), and that $T^* = \sigma_{\max,k^*}$. This can be rewritten as $T^* = \sum_{k \geq 1} \tau_{\max,k}$ with $\tau_{\max,k} = \sigma_{\max,k} - \sigma_{\max,k-1}$ for $k \geq 1$ or, since $\tau_{\max,1} = \tau_{\max}$, as $T^* - \tau_{\max} = \sum_{k \geq 2} \tau_{\max,k}$. We now proceed to derive a stochastic upper bound on $\sum_{k \geq 2} \tau_{\max,k}$ by using a coupling argument.

Let $(X'_{r,k}, k \geq 1, r \in \mathcal{R})$ and $(Y_{r,k}, k \geq 1, r \in \mathcal{R})$ be two sequences of i.i.d. random variables with common distribution the stationary distribution, say X , of the A_r 's; we assume moreover that these two sequences are independent from one another and also independent from the active queues \mathbf{A} . For each $k \geq 1$ let $A'_{r,k}$ be a stationary version of A_r such that $A'_{r,k}(0) = X'_{r,k}$ and $A'_{r,k}(i) \geq A_r(i + \tau_{r,k})$ for every $i \geq 0$. This can be done with, e.g., the usual coupling between two processes with different initial states and the same stochastic primitives; in this case, the ordering between the two processes comes from the fact that $A_r(\tau_{r,k}) = 0$ by definition of $\tau_{r,k}$.

Consider now, for $r \in \mathcal{R}$ and $k \geq 1$, the random variables $X''_{r,k}$ defined by:

$$X''_{r,k} = \begin{cases} A'_{r,k}(\sigma_{\max,k} - \tau_{r,k}) & \text{if } \sigma_{\max,k} > \tau_{r,k}, \\ Y_{r,k} & \text{otherwise.} \end{cases}$$

A moment's thought reveals that for each $k \geq 1$ the random variables $(X''_{r,k}, r \in \mathcal{R})$ are i.i.d. with common distribution X . Indeed, consider the process $A'_{r,k}$ at the random time $\sigma_{\max,k} - \tau_{r,k}$.

If $\sigma_{\max,k} > \tau_{r,k}$, the r th active queue A_r has hit 0 somewhere in the past at time $\tau_{r,k}$: at this time we forgot about its state by using $X'_{r,k}$ to build $A'_{r,k}$, and the remaining time to go $\sigma_{\max,k} - \tau_{r,k}$ depends only on the behavior of the other queues and is thus independent from $A'_{r,k}$. Thus in this case, considering $A'_{r,k}(\sigma_{\max,k} - \tau_{r,k})$ amounts to consider a stationary process sampled at an independent random time, and so this random variable is indeed distributed according to X and is independent from the other queues.

If on the other hand $\sigma_{\max,k} = \tau_{r,k}$, this means that the r th active queue is one of the last queue to hit 0, and again at this time we forget about this by using $Y_{r,k}$ which is distributed according to X independently from everything else.

We now use again the same coupling as before, between processes sharing the same stochastic primitives, to build for each $r \in \mathcal{R}$ and $k \geq 1$ a process $A''_{r,k}$ starting at $X''_{r,k}$ and sharing the same stochastic primitives as $A'_{r,k}$. This leads to a collection of processes $(A''_{r,k}, r \in \mathcal{R}, k \geq 1)$ such that: (1) for each $k \geq 1$, $\mathbf{A}''_k = (A''_{r,k}, r \in \mathcal{R})$ is a stationary version of \mathbf{A} and (2) for every $k \geq 1, i \geq 0$, and $r \in \mathcal{R}$, it holds that $A_{r,k}(\sigma_{\max,k} + i) \leq A''_{r,k}(i)$. In particular, $\tau_{\max,k} \leq \tau''_{\max,k}$ where $\tau''_{\max,k} = \inf\{i \geq 0 : \mathbf{A}''_k(i) = \mathbf{0}\}$, and so $T^* - \tau_{\max} \leq \sum_{k=2}^{G-1} \tau''_{\max,k}$, with the convention $\sum_2^1 = 0$ and

where $G = \inf\{k \geq 2 : \tau''_{\max,k} = 0\}$, which provides the desired stochastic upper bound.

To conclude the proof, it remains to show that $\mathbb{E}_{\mathbf{a}}(\sum_{k=2}^{G-1} \tau''_{\max,k})$ is bounded away from $+\infty$ uniformly in $\rho \leq \rho_0$ and $\mathbf{a} \in \mathbb{N}^R$. Let $\mathbf{X} = (X'_{1,r}, r \in \mathcal{R})$: then by construction of the $\tau''_{\max,k}$'s and by definition of G , $G - 2$ is a geometric random variable with parameter $\mathbb{P}(\mathbf{X} \neq \mathbf{0})$ and conditionally on G , $(\tau''_{\max,k}, 2 \leq k \leq G - 1)$ are i.i.d. with common distribution τ_{\max} under $\mathbb{P}_{\mathbf{X}}(\cdot \mid \mathbf{X} \neq \mathbf{0})$. In particular,

$$\mathbb{E}_{\mathbf{a}}\left(\sum_{k=2}^{G-1} \tau''_{\max,k}\right) = \mathbb{E}(G - 2)\mathbb{E}_{\mathbf{X}}(\tau_{\max} \mid \mathbf{X} \neq \mathbf{0}) = \frac{\mathbb{E}_{\mathbf{X}}(\tau_{\max})}{\mathbb{P}(\mathbf{X} = \mathbf{0})}.$$

Since $\tau_{\max} \leq \tau_1 + \dots + \tau_R$, $\mathbb{E}_a(\tau_1) = a/(1 - \mathbb{E}\xi)$ and \mathbf{X} is a vector of R i.i.d. random variables with common distribution X , we finally obtain the bound

$$\mathbb{E}_{\mathbf{a}}\left(\sum_{k=2}^{G-1} \tau''_{\max,k}\right) \leq \frac{R}{1 - \mathbb{E}\xi} \frac{\mathbb{E}(X)}{\mathbb{P}(X = 0)^R}.$$

This upper bound does not depend on \mathbf{a} , and so it remains to show that its supremum over $\rho \leq \rho_0 < 2$ is finite. Since X is the stationary distribution of A_1 , X is equal in distribution to $X + \xi - \mathbb{1}_{\{X > 0\}}$. In particular, taking the mean we obtain $\mathbb{P}(X = 0) = 1 - \mathbb{E}\xi$ and so it remains to show that $\sup_{\rho \leq \rho_0} \mathbb{E}X < +\infty$. To compute $\mathbb{E}X$, we can start from the equality

$$\mathbb{E}\left(e^{-\lambda X}\right) = \mathbb{E}\left(e^{-\lambda(X+\xi-\mathbb{1}_{\{X>0\}})}\right),$$

which leads to

$$\phi(\lambda) = \frac{\varphi(\lambda)(1 - e^{-\lambda})\varphi'(0)}{\varphi(\lambda) - 1},$$

where $\phi(\lambda) = \mathbb{E}(e^{-\lambda X})$ and $\varphi(\lambda) = \mathbb{E}(e^{-\lambda(\xi-1)})$, and we have used $\varphi'(0) = 1 - \mathbb{E}\xi = \mathbb{P}(X = 0)$. Using $\mathbb{E}X = -\phi'(0)$, we can compute after some algebra

$$\mathbb{E}X = \frac{\varphi''(0) - \varphi'(0)}{2\varphi'(0)} = \frac{\mathbb{E}\xi^2 + 2 - 3\mathbb{E}\xi}{2(1 - \mathbb{E}\xi)}.$$

Since $\sup_{\rho \leq \rho_0} \mathbb{E}\xi^2$ is finite by assumption (see the discussion at the end of Sect. 2.2), the result is finally proved. \square

The previous lemma justifies the approximation $T^* \approx \tau_{\max}$, and to further control τ_{\max} , we will use that $(\tau_r, r \in \mathcal{R})$ under $\mathbb{P}_{\mathbf{a}}$ is equal in distribution to $(V_r(a_r), r \in \mathcal{R})$, where $(V_r, r \in \mathcal{R})$ are i.i.d. random walks started at 0 and with step size distribution δ , equal in distribution to τ_1 under \mathbb{P}_1 (i.e., δ is the time needed for a random walk with step size distribution $\xi - 1$ to go from 1 to 0). Then δ has finite mean $1/(1 - \mathbb{E}\xi)$ and

also finite variance, which we denote by v . To control the maximum of random walks, we will use the following result.

Lemma 4.3 *Let $(W_r, r \in \mathcal{R})$ be i.i.d. random walks started at 0 with step size distribution having finite mean m and finite variance w . Then for any $\mathbf{x} = (x_r, r \in \mathcal{R}) \in \mathbb{N}^R$ it holds that*

$$\mathbb{E} \left(\max_{r \in \mathcal{R}} W_r(x_r) \right) \leq m|\mathbf{x}| + R(w|\mathbf{x}|)^{1/2}. \quad (6)$$

Furthermore, if the step size distribution of the W_r 's is non-negative, then

$$\mathbb{E} \left[\left(\max_{r \in \mathcal{R}} W_r(x_r) \right)^{1/2} \right] \geq (m|\mathbf{x}|)^{1/2} - (w/m^{3/2})|\mathbf{x}|^{-1/2}. \quad (7)$$

Proof Fix some $\mathbf{x} \in \mathbb{N}^R$ and for $r \in \mathcal{R}$ let $Y_r = (W_r(x_r) - mx_r)/(wx_r)^{1/2}$, so that $\mathbb{E}Y_r = 0$ and $\mathbb{E}(|Y_r|) \leq (\mathbb{E}(Y_r^2))^{1/2} = 1$. To prove the upper bound, we start from the equality $\max_r W_r(x_r) = \max_r (mx_r + (wx_r)^{1/2}Y_r)$ which implies

$$\max_{r \in \mathcal{R}} W_r(x_r) \leq \max_{r \in \mathcal{R}} \left(mx_r + (wx_r)^{1/2}|Y_r| \right) \leq m|\mathbf{x}| + (w|\mathbf{x}|)^{1/2}\|\mathbf{Y}\|.$$

Averaging on both sides yields (6), since $\mathbb{E}(\|\mathbf{Y}\|) = \sum_r \mathbb{E}(|Y_r|) \leq R$. Let us now prove the lower bound, so assume that the step size distribution of the W_r 's is non-negative. Since $\mathbb{E}(\max_r W_r(x_r)^{1/2}) \geq \max_r \mathbb{E}(W_r(x_r)^{1/2})$ and

$$\max_{r \in \mathcal{R}} \left((mx_r)^{1/2} - (w/m^{3/2})x_r^{-1/2} \right) \geq (m|\mathbf{x}|)^{1/2} - (w/m^{3/2})|\mathbf{x}|^{-1/2},$$

it is enough to show that $\mathbb{E}(W_r(x_r)^{1/2}) \geq (mx_r)^{1/2} - (w/m^{3/2})x_r^{-1/2}$ for each fixed $r \in \mathcal{R}$. Let

$$f(y) = \frac{1 + y/2 - (1 + y)^{1/2}}{y^2}, \quad y \geq -1,$$

so that, defining $y = (w/(m^2x_r))^{1/2}Y_r$ which satisfies $y \geq -1$ since $W_r(x_r) \geq 0$, we obtain

$$W_r(x_r)^{1/2} = \left((wx_r)^{1/2}Y_r + mx_r \right)^{1/2} = (mx_r)^{1/2} \left(1 + y/2 - y^2 f(y) \right).$$

Since $\mathbb{E}y = 0$, taking expectation on both sides leads to

$$\mathbb{E}(W_r(x_r)^{1/2}) = (mx_r)^{1/2} - (mx_r)^{1/2}\mathbb{E}(y^2 f(y)).$$

Since $\sup f = 1/2 \leq 1$ and $\mathbb{E}(y^2) = w/(m^2x_r)$, we get the result. \square

We now prove that the process is \mathbf{Q} stable for $\rho < 1$; moreover, we will need the fact that the mean stationary number of packets is finite.

Proposition 4.4 *If $\rho < 1$, then \mathbf{Q} is positive recurrent and $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|) < +\infty$.*

Proof Because $\beta = +\infty$, we have $\mathbf{Q}^i(k) = 0$ for $k \geq 1$: we therefore assume without loss of generality that $\mathbf{Q}^i(k) = 0$ for $k \geq 0$, and we only have to show that \mathbf{Q}^a is positive recurrent and has finite stationary mean. To prove Proposition 4.4, it is enough to prove that

$$\lim_{K \rightarrow +\infty} \sup_{\mathbf{q}^a: |\mathbf{q}^a| \geq K} \left(\frac{\mathbb{E}_{\mathbf{q}} [|\mathbf{Q}^a(2)| - |\mathbf{q}^a|]}{|\mathbf{q}^a|} \right) < 0, \quad (8)$$

where here and in the sequel, $\mathbf{q} = (\mathbf{q}^a, \mathbf{0})$. Indeed, this shows that $|\cdot|$ is a Lyapunov function, which implies positive recurrence of \mathbf{Q}^a using for instance the Foster-Lyapunov criterion. In the terminology of [7], it also shows that $|\cdot|$ is a geometric Lyapunov function, and Theorem 5 in [7] states that (8) implies that $\mathbb{E}_\infty[|\mathbf{Q}^a(0)|]$, and in particular $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$, is finite. Thus we only have to prove (8). By (3) and the fact that $\mathbf{A}(T^*) = \mathbf{0}$, we obtain $\mathbb{E}_{\mathbf{q}} [|\mathbf{Q}^a(1)|] = \mathbb{E}_{\mathbf{q}^a} (|\mathbf{S}(T^*)|)$. Since \mathbf{S} and T^* are independent, (6) gives that

$$\mathbb{E}_{\mathbf{a}}(|\mathbf{S}(T^*)|) \leq \mathbb{E}(\xi)\mathbb{E}_{\mathbf{a}}(T^*) + Rv^{1/2}\mathbb{E}_{\mathbf{a}}((T^*)^{1/2}),$$

for any $\mathbf{a} \in \mathbb{N}^R$ (recall that v is the variance of ξ , assumed to be finite). Thus after rearranging the terms and using Jensen's inequality, we end up with the bound

$$\mathbb{E}_{\mathbf{q}} [|\mathbf{Q}^a(1)| - |\mathbf{q}^a|] \leq -(1 - \rho)\mathbb{E}_{\mathbf{q}^a}(T^*) + \Psi(\mathbf{q}^a) + Rv^{1/2}[\mathbb{E}_{\mathbf{a}}(T^*)]^{1/2}, \quad (9)$$

where $\Psi(\mathbf{a}) = \mathbb{E}_{\mathbf{a}}((1 - \mathbb{E}\xi)T^* - |\mathbf{A}(0)|)$. We now argue that $\Psi(\mathbf{a}) \leq c[\mathbb{E}_{\mathbf{a}}(T^*)]^{1/2}$ for some finite constant c independent of \mathbf{a} . By Lemma 4.2, we have $\mathbb{E}_{\mathbf{a}}(T^*) \leq \mathbb{E}_{\mathbf{a}}(\tau_{\max}) + c'$ for some finite constant c' independent of \mathbf{a} . Further, since τ_{\max} is equal in distribution to $\max_r V_r(a_r)$, (6) gives (recall that v is the variance of the step size distribution of the V_r 's, which has mean $1/(1 - \mathbb{E}\xi)$)

$$\mathbb{E}_{\mathbf{a}}(T^*) \leq \frac{|\mathbf{a}|}{1 - \mathbb{E}\xi} + Rv^{1/2}|\mathbf{a}|^{1/2} + c',$$

which can be rewritten as $\Psi(\mathbf{a}) \leq R(1 - \mathbb{E}\xi)v^{1/2}|\mathbf{a}|^{1/2} + c'(1 - \mathbb{E}\xi)$. Since $T^* \geq \tau_r$ and $\mathbb{E}_{a_r}\tau_r = a_r/(1 - \mathbb{E}\xi)$, we obtain $|\mathbf{a}| \leq (1 - \mathbb{E}\xi)\mathbb{E}_{\mathbf{a}}(T^*)$ which, together with the previous inequality, implies the existence of the desired constant c such that $\Psi(\mathbf{a}) \leq c[\mathbb{E}_{\mathbf{a}}(T^*)]^{1/2}$. Defining

$$\Gamma(\mathbf{a}) = (1 - \rho)\mathbb{E}_{\mathbf{a}}(T^*) - (c + Rv^{1/2})[\mathbb{E}_{\mathbf{a}}(T^*)]^{1/2}, \quad (10)$$

we can rewrite (9) as $\mathbb{E}_{\mathbf{q}} [|\mathbf{Q}^a(1)| - |\mathbf{q}^a|] \leq -\Gamma(\mathbf{q}^a)$. Using the Markov property and (2), this gives

$$\mathbb{E}_{\mathbf{q}} [|\mathbf{Q}^a(2)| - |\mathbf{q}^a|] \leq -\mathbb{E}_{\mathbf{q}^a} (\Gamma(\mathbf{q}^a) + \Gamma(\mathbf{S}(T^*))). \quad (11)$$

When $|\mathbf{q}^a|$ is large, at least one of the R coordinates of \mathbf{q}^a must be large. Since $\mathbb{E}_{\mathbf{a}}(T^*) \geq a_r/(1 - \mathbb{E}\xi)$, it is not hard to show that

$$\lim_{K \rightarrow +\infty} \inf_{\mathbf{q}^a: |\mathbf{q}^a| \geq K} \frac{\mathbb{E}_{\mathbf{q}^a} (\Gamma(\mathbf{q}^a) + \Gamma(\mathbf{S}(T^*)))}{|\mathbf{q}^a|} > 0,$$

which completes the proof of the result. \square

We need a last technical lemma before proving Theorem 4.1.

Lemma 4.5 *Let \Rightarrow denote weak convergence under \mathbb{P}_{∞} as $\rho \uparrow 1$. Then for any $r \in \mathcal{R}$, $Q_r^a(0) \Rightarrow +\infty$. In particular, $\tau_{\max} \Rightarrow +\infty$ and*

$$\liminf_{\rho \uparrow 1} \mathbb{E}_{\infty} \left(|\mathbf{Q}^a(0)|^{1/2} - (Q_1^a(0))^{1/2} \right) > 0.$$

Proof Fix some $r \in \mathcal{R}$, and let us first prove that $Q_r^a(0) \Rightarrow +\infty$. The proof relies on a coupling argument between the system *in the normal time scale* (i.e., not embedded at switching epochs) and a polling system. So let $\tilde{Q}^a(t)$ and $\tilde{Q}^i(t)$ be the number of packets in the r th active and inactive queues, respectively, at the beginning of the t th time slot. Furthermore, assume without loss of generality that $\mathbf{Q}^i(0) = \mathbf{0}$.

We build a process $(\underline{Q}^a, \underline{Q}^i)$ that will be closely related to a polling system and which will provide a lower bound on $(\tilde{Q}^a, \tilde{Q}^i)$. We explain how to build $(\underline{Q}^a, \underline{Q}^i)$ on $\{0, \dots, T^*\}$, the construction can be repeated by induction between any two switching times:

- At time 0 we have $\underline{Q}^a(0) \leq \tilde{Q}^a(0)$ and $\underline{Q}^i(0) = 0$;
- The two processes $(\tilde{Q}^a, \tilde{Q}^i)$ and $(\underline{Q}^a, \underline{Q}^i)$ share the same stochastic primitives until \underline{Q}^a hits 0 at time $\underline{\tau} = \inf\{t \geq 0 : \underline{Q}^a(t) = 0\}$. Since $\underline{Q}^a(0) \leq \tilde{Q}^a(0)$, we have in particular $\underline{\tau} \leq T^*$;
- On $[\underline{\tau}, \dots, T^* - 1]$, the two processes \underline{Q}^a and \underline{Q}^i are frozen, i.e., $\underline{Q}^a(t) = \underline{Q}^a(\underline{\tau})$ and $\underline{Q}^i(t) = \underline{Q}^i(\underline{\tau})$ for $\underline{\tau} \leq t < T^*$;
- At time T^* the active and inactive queues are switched for both $(\tilde{Q}^a, \tilde{Q}^i)$ (by definition of T^*) and $(\underline{Q}^a, \underline{Q}^i)$ (by construction).

Since \underline{Q}^i and \tilde{Q}^i share the same stochastic primitives up to time $\underline{\tau}$, we have $\underline{Q}^i(\underline{\tau}) = \tilde{Q}^i(\underline{\tau})$; since \underline{Q}^i is frozen afterward while \tilde{Q}^i continues to increase, and the active and inactive queues are switched at time T^* , we obtain $\underline{Q}^a(T^*) \leq \tilde{Q}^a(T^*)$. Since moreover $\mathbf{Q}^i(1) = \mathbf{0}$, it follows by induction that $\underline{Q}^a(k) \leq \tilde{Q}^a(k)$ for every $k \geq 0$. Let $\underline{\tau}(k)$ be the k th time at which \underline{Q}^a hits 0 and $T^*(k)$ be the k th switching time: then

$$Q_r^a(k) = \tilde{Q}^a(T^*(k)) \geq \underline{Q}^a(T^*(k)) = \underline{Q}^i(\underline{\tau}(k)) \quad (12)$$

(the last equality is only valid if $\tau(k) < T^*(k)$; if $\tau(k) = T^*(k)$ then $\underline{Q}^i(\tau(k))$ is to be understood as the value of the inactive queue just before the switch occurs). Finally, we note that by construction, the process $(\underline{Q}^i(\tau(k)), k \geq 0)$ is equal in distribution to the process $(P(\gamma(k)), k \geq 0)$ where:

- $P(t)$ for $t \geq 0$ is the total number of packets in the k th time slot of a two-queue polling system, operating under the exhaustive service discipline, with zero switch-over time, where time is discrete and arrivals in each time slot in each queue are distributed according to ξ ;
- $\gamma(k)$ is the k th time at which a queue empties.

Next, we note that the process $(P(\gamma(k)), k \geq 0)$ is a branching process with immigration at 0, the offspring distribution η of the branching process being equal in distribution to $S_1(\tau_1)$ under \mathbb{P}_1 , i.e., η is the number of packets arriving in the inactive queue during the time needed to make the state of the active queue decrease by one. This property can be seen directly from the system's dynamics; at a higher level, it also comes from the fact that the exhaustive service discipline satisfies the well-known branching property [18]. In any case, we can adapt the proof of [6] to show that the stationary distribution of $(P(\gamma(k)), k \geq 0)$ converges weakly to $+\infty$ as $\rho \uparrow 1$. In view of (12) this shows that $Q_r^a(0) \Rightarrow +\infty$, which concludes the proof of the first claim of the lemma. The second claim is then immediate, and so it remains to prove the third and last claim. Using $a^{1/2} - b^{1/2} = (a - b)/(a^{1/2} + b^{1/2})$, we can rewrite

$$|\mathbf{Q}^a(0)|^{1/2} - (Q_1^a(0))^{1/2} = \frac{|\mathbf{Q}^a(0)| - Q_1^a(0)}{(Q_1^a(0) + 1)^{1/2}} \times \frac{(Q_1^a(0) + 1)^{1/2}}{|\mathbf{Q}^a(0)|^{1/2} + (Q_1^a(0))^{1/2}},$$

which by (3), is equal in distribution to

$$\begin{aligned} & \frac{|\mathbf{S}(T^*)| - S_1(T^*)}{(S_1(T^*) + 1)^{1/2}} \times \frac{(S_1(T^*) + 1)^{1/2}}{|\mathbf{S}(T^*)|^{1/2} + (S_1(T^*))^{1/2}} \\ &= \max_{r \in \mathcal{R}} \left(\frac{S_r(T^*) - S_1(T^*)}{(S_1(T^*) + 1)^{1/2}} \right) \times \frac{(S_1(T^*) + 1)^{1/2}}{|\mathbf{S}(T^*)|^{1/2} + (S_1(T^*))^{1/2}}. \end{aligned}$$

Since \mathbf{S} and T^* are independent and $T^* \Rightarrow +\infty$ (as a result of the second claim and $T^* \geq \tau_{\max}$), the law of large numbers implies that

$$\frac{(S_1(T^*) + 1)^{1/2}}{|\mathbf{S}(T^*)|^{1/2} + (S_1(T^*))^{1/2}} \Rightarrow \frac{1}{2},$$

while the central limit theorem implies that

$$\frac{S_r(T^*) - S_1(T^*)}{(S_1(T^*) + 1)^{1/2}} = \left(\frac{S_r(T^*) - T^* \mathbb{E}\xi}{(T^*)^{1/2}} - \frac{S_1(T^*) - T^* \mathbb{E}\xi}{(T^*)^{1/2}} \right) \times \left(\frac{S_1(T^*) + 1}{T^*} \right)^{-1/2}$$

converges weakly to the difference of two independent normal random variables. Gathering these results, we obtain the weak convergence of $|\mathbf{Q}^a(0)|^{1/2} - (Q_1^a(0))^{1/2}$

toward a non-negative random variable with strictly positive mean. Invoking Fatou's lemma completes the proof. \square

Remark Although the stationary distribution of the process P introduced in the previous proof scales like $1/(1-\rho)$ as $\rho \uparrow 1$, using similar techniques as in [6] one can show that the stationary distribution $P(\gamma(\infty))$ of the embedded process $(P(\gamma(k)), k \geq 0)$ scales like $1/(1-\rho)^U$ with U a random variable uniformly distributed on $(0, 1)$, in the sense that $\log P(\gamma(\infty))/\log(1/(1-\rho))$ converges weakly to U as $\rho \uparrow 1$.

Proof of Theorem 4.1 Since $\beta = +\infty$, we have $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|) = \mathbb{E}_\infty(\|\mathbf{A}(0)\|) = R\mathbb{E}_\infty(A_1(0))$ and so we only need to prove the two inequalities

$$\limsup_{\rho \uparrow 1} \left[(1-\rho)^2 \mathbb{E}_\infty(A_1(0)) \right] < +\infty \quad \text{and} \quad \liminf_{\rho \uparrow 1} \left\{ (1-\rho) \mathbb{E}_\infty[A_1(0)^{1/2}] \right\} > 0. \quad (13)$$

Indeed, the first one directly implies the upper bound in (5) while the second one implies the corresponding lower bound by Jensen's inequality.

Proof of the first inequality in (13) Starting from (4), using that $\mathbb{E}_a(\tau_1) = a/(1-\mathbb{E}\xi)$, subtracting on both sides $\mathbb{E}(\xi)\mathbb{E}_\infty(\tau_1)$ (for this precise operation we need the finiteness of the stationary first moment proved in Proposition 4.4, to avoid an indetermination of the kind $\infty - \infty$) and dividing by $\mathbb{E}\xi$, we end up with

$$\frac{1}{\mathbb{E}\xi} \left(1 - \frac{\mathbb{E}\xi}{1 - \mathbb{E}\xi} \right) \mathbb{E}_\infty(A_1(0)) = \mathbb{E}_\infty(T^* - \tau_1).$$

Then, adding and subtracting τ_{\max} in the right hand side, and using that $\mathbb{E}\xi = \rho/2$, we obtain

$$g_\rho(1-\rho)\mathbb{E}_\infty(A_1(0)) = \mathbb{E}_\infty(\tau_{\max} - \tau_1) + \mathbb{E}_\infty(T^* - \tau_{\max}),$$

with $g_\rho = 4/(\rho(2-\rho))$. Thus in view of Lemmas 4.2 and 4.5, to prove the first inequality in (13) we only have to show that

$$\limsup_{\rho \uparrow 1} \left(\frac{\mathbb{E}_\infty(\tau_{\max} - \tau_1)}{[\mathbb{E}_\infty(A_1(0))]^{1/2}} \right) < +\infty. \quad (14)$$

Applying (6) to τ_{\max} under \mathbb{P}_a (equal in distribution to $\max_r V_r(a_r)$), we obtain the inequality $\mathbb{E}_a(\tau_{\max} - \tau_1) \leq (|a| - a_1)/(1 - \mathbb{E}\xi) + Rv^{1/2}|a|^{1/2}$. Integrating over the stationary distribution of \mathbf{Q} and using Jensen's inequality, we obtain

$$\mathbb{E}_\infty(\tau_{\max} - \tau_1) \leq \frac{1}{1 - \mathbb{E}\xi} \mathbb{E}_\infty(|\mathbf{A}(0)| - A_1(0)) + Rv^{1/2} [\mathbb{E}_\infty(A_1(0))]^{1/2}.$$

In particular, to prove (14) it is enough to show that

$$\limsup_{\rho \uparrow 1} \left(\frac{\mathbb{E}_\infty (|\mathbf{A}(0)| - A_1(0))}{[\mathbb{E}_\infty (A_1(0))]^{1/2}} \right) < +\infty. \quad (15)$$

We have already shown in the proof of Proposition 4.4 that

$$\mathbb{E}_\infty (|\mathbf{A}(0)|) = \mathbb{E}_\infty (|\mathbf{S}(T^*)|) \leq \mathbb{E}(\xi)\mathbb{E}_\infty(T^*) + Rv^{1/2} [\mathbb{E}_\infty(T^*)]^{1/2},$$

and since $\mathbb{E}_\infty(A_1(0)) = \mathbb{E}(\xi)\mathbb{E}_\infty(T^*)$ this gives

$$\mathbb{E}_\infty (|\mathbf{A}(0)| - A_1(0)) \leq Rv^{1/2} [\mathbb{E}_\infty(A_1(0))/\mathbb{E}\xi]^{1/2}.$$

This proves (15) which completes the proof of the first inequality in (13).

Proof of the second inequality in (13) We have

$$\mathbb{E}_\infty [(Q_1^a(0))^{1/2}] = \mathbb{E}_\infty [(S_1(T^*))^{1/2}] \geq \mathbb{E}_\infty [(S_1(\tau_{\max}))^{1/2}].$$

Applying (7) (with $R = 1$) to $S_1(\tau_{\max})$ by using the independence between S_1 and τ_{\max} , we obtain

$$\mathbb{E}_\infty [(Q_1^a(0))^{1/2}] \geq (\mathbb{E}\xi)^{1/2} \mathbb{E}_\infty [(\tau_{\max})^{1/2}] - c \mathbb{E}_\infty [(\tau_{\max})^{-1/2}],$$

for some finite constant c independent of ρ . Applying again (7) to τ_{\max} we obtain

$$\begin{aligned} \mathbb{E}_\infty [(Q_1^a(0))^{1/2}] &\geq \left(\frac{\mathbb{E}\xi}{1 - \mathbb{E}\xi} \right)^{1/2} \mathbb{E}_\infty (|\mathbf{Q}^a(0)|^{1/2}) \\ &\quad - c' \mathbb{E}_\infty (|\mathbf{Q}^a(0)|^{-1/2}) - c' \mathbb{E}_\infty [(\tau_{\max})^{-1/2}] \end{aligned}$$

for some constant c' , still independent from ρ . Subtracting $(\mathbb{E}\xi/(1 - \mathbb{E}\xi))^{1/2} \mathbb{E}_\infty (Q_1^a(0)^{1/2})$ on both sides we finally end up with

$$\begin{aligned} h_\rho(1 - \rho) \mathbb{E}_\infty [(Q_1^a(0))^{1/2}] &\geq \left(\frac{\mathbb{E}\xi}{1 - \mathbb{E}\xi} \right)^{1/2} \mathbb{E}_\infty (|\mathbf{Q}^a(0)|^{1/2} - (Q_1^a(0))^{1/2}) \\ &\quad - c' \mathbb{E}_\infty (|\mathbf{Q}^a(0)|^{-1/2}) - c' \mathbb{E}_\infty [(\tau_{\max})^{-1/2}], \end{aligned}$$

with $h_\rho = (1 - (\rho/(2 - \rho))^{1/2})/(1 - \rho) \rightarrow 1$ as $\rho \rightarrow 1$. We can then invoke the results of Lemma 4.5 to conclude the proof of the second inequality in (13), which concludes the proof of the theorem. \square

5 Extension to the strongly persistent case

In this section we present various arguments supporting the conjecture formulated at the end of Sect. 3 that Theorem 4.1 should remain valid in the strongly persistent case $\beta \in (1, \infty)$. On the one hand, we prove in Proposition 5.1 that active queues only advertize a release when they are close to empty, which is the main theoretical result of this section. We also provide a corollary to Proposition 5.1 (Corollary 5.5 in Sect. 5.1.3) that brings closer the two cases $\beta = +\infty$ and $\beta > 1$. On the other hand, we present extensive simulation results in Sect. 5.2 that support our conjecture. Finally, based on these theoretical and simulation results, we elaborate in Sect. 5.3 on possible interesting heavy-tailed phenomena that may arise for $1 < \beta < 2$.

In the rest of this section we let T_1 be the time at which A_1 advertizes a release for the first time.

Proposition 5.1 *If $\rho < 2$ and $\beta > 1$, then $(A_1(T_1), T_1 - \tau_1)$ under \mathbb{P}_{a_1} converges weakly as $a_1 \rightarrow +\infty$ to a non-degenerate random variable.*

By non-degenerate, we mean a random variable $(X, Y) \in \mathbb{N} \times \mathbb{Z}$ such that both X and Y are non-deterministic and almost surely finite (an explicit expression for the weak limit of $(A_1(T_1), T_1 - \tau_1)$ is given in Lemma 5.3).

5.1 Proof of Proposition 5.1

Fix in the rest of this subsection some $\rho < 2$. The proof of Proposition 5.1 proceeds in two steps: in the first step we show that the proof of Proposition 5.1 reduces to proving a simpler property of some particular random walk [see (16)]. In the second step we prove that this property holds when $\beta > 1$.

5.1.1 First step

Since before time T_1 , the active queue A_1 does by definition not advertize any release, it is enough to prove Proposition 5.1 in the case $\zeta = 0$, which we therefore assume in this subsection and the following one. In particular, A_1 is a random walk with step size distribution $\xi - 1$ reflected at 0. The reduction of the proof of Proposition 5.1 to proving (16) relies on the following coupling of the processes A_1 for all possible initial states $a \geq 0$.

Let V and W^\uparrow be two independent processes with the following distribution. Let V be a version of A_1 under \mathbb{P}_0 , i.e., $(V(k), k \geq 0)$ is a random walk started at 0, with step size distribution $\xi - 1$ and reflected at 0.

Let W^\uparrow be a random walk started at 0, with step size distribution $1 - \xi$ and conditioned on never visiting 0 after time 0: since $\mathbb{E}(1 - \xi) > 0$ this conditioning is well-defined. Let moreover $\kappa(a) = \max\{k \geq 0 : W^\uparrow(k) = a\}$ be the time of the last visit to $a \in \{0, 1, \dots, \infty\}$ (to be understood as $\kappa(a) = +\infty$ for $a = +\infty$), so that $\kappa(a)$ is almost surely finite if a is finite. Let finally W_a^\uparrow be the process W^\uparrow stopped at $\kappa(a)$, i.e., $W_a^\uparrow(k) = W^\uparrow(k)$ if $k \leq \kappa(a)$ and $W_a^\uparrow(k) = W^\uparrow(\kappa(a)) = a$ if $k \geq \kappa(a)$.

Lemma 5.2 *Extend A_1 on \mathbb{Z} by setting $A_1(k) = A_1(0)$ for $k \leq 0$, and define the two processes $A^+ = (A_1(\tau_1 + k), k \geq 0)$ and $A^- = (A_1(\tau_1 - k), k \geq 0)$. Then for any finite $a \geq 0$, (A^+, A^-) under \mathbb{P}_a is equal in distribution to (V, W_a^\uparrow) .*

In particular, as $a \rightarrow +\infty$, (A^+, A^-) under \mathbb{P}_a converges weakly to (V, W^\uparrow) .

Proof That A^+ is equal in distribution to V and is independent from A^- follows from the strong Markov property at time τ_1 . That A^- is equal in distribution to W_a^\uparrow comes from duality. The weak convergence result then follows from the fact that $\kappa(a) \rightarrow +\infty$ almost surely as $a \rightarrow +\infty$, so that $(V, W_a^\uparrow) \rightarrow (V, W^\uparrow)$ almost surely as $a \rightarrow +\infty$. \square

Essentially, this representation of A_1 shifts the origin of time at τ_1 : A^+ looks at A_1 from time τ_1 forward in time, while A^- looks at A_1 from τ_1 backward in time. Moreover, this representation couples all the processes A_1 with different initial states on the same probability space, which yields a simple representation for the law of $(A_1(T_1), T_1 - \tau_1)$. Let in the sequel $Z = (Z(k), k \in \mathbb{Z})$ be the following process (indexed by \mathbb{Z}): $Z(k) = V(k)$ if $k \geq 0$ and $Z(k) = W^\uparrow(-k)$ if $k \leq 0$. The previous coupling immediately implies the following result.

Lemma 5.3 *Let $(U_k, k \in \mathbb{Z})$ be i.i.d., uniformly distributed in $[0, 1]$, independent from Z , and for $0 \leq a \leq +\infty$ and $k \in \mathbb{Z}$ let*

$$D_{a,k} = \begin{cases} 0 & \text{if } k < -\kappa(a), \\ \mathbb{1}_{\{U_k < \psi(Z(k))\}} & \text{else,} \end{cases}$$

and $T_a^Z = \inf \{k \in \mathbb{Z} : D_{a,k} = 1\}$. Then for any finite $a \geq 0$, $(A_1(T_1), T_1 - \tau_1)$ under \mathbb{P}_a is equal in distribution to $(Z(T_a^Z), T_a^Z)$ and in particular, it converges weakly as $a \rightarrow +\infty$ to $(Z(T_\infty^Z), T_\infty^Z)$ (to be understood as $Z(T_\infty^Z) = +\infty$ if $T_\infty^Z = -\infty$).

Proof The equality in law between $(A_1(T_1), T_1 - \tau_1)$ and $(Z(T_a^Z), T_a^Z)$ is clear from the construction, and not difficult to formalize. Moreover, T_a^Z is by construction decreasing in a and so its limit as $a \rightarrow +\infty$ exists. It is not hard to show that its limit is exactly T_∞^Z and by continuity we deduce that $Z(T_a^Z) \rightarrow Z(T_\infty^Z)$ as $a \rightarrow +\infty$, which implies the result. \square

Thus to prove Proposition 5.1, we only have to establish that $|T_\infty^Z|$ is (almost surely) finite. Since V is positive recurrent and starts at 0, it is clear that $\min \{k \geq 0 : D_{\infty,k} = 1\}$ is finite and so to prove that $|T_\infty^Z|$ is finite, we only have to demonstrate that the random variable $\inf \{k \leq 0 : D_{\infty,k} = 1\}$ is almost surely finite. Going back to the definition of $D_{\infty,k}$, we see that we have to prove that $\sup \{k \geq 0 : U_{-k} < \psi(W^\uparrow(k))\}$ is finite, which informally means that W^\uparrow advertizes a release only finitely many times.

So in the sequel, we consider $(U_k, k \geq 0)$ i.i.d. random variables, uniformly distributed in $[0, 1]$ and independent of W^\uparrow , and we define

$$N = \sum_{k \geq 0} \mathbb{1}_{\{U_k < \psi(W^\uparrow(k))\}},$$

which can be intuitively interpreted as the number of times W^\uparrow advertizes a release. The proof of Proposition 5.1 will thus be complete if we can prove that

$$\mathbb{P}(N < +\infty) = 1. \tag{16}$$

5.1.2 Second step

We now assume that $\beta > 1$ (the results of the previous subsection indeed hold for any function ψ) and we prove that $\mathbb{P}(N \geq n) \rightarrow 0$ as $n \rightarrow +\infty$, which will prove (16). By definition,

$$\mathbb{P}(N = 0) = \mathbb{P}\left(U_k > \psi(W^\uparrow(k)), k \geq 0\right),$$

and since W^\uparrow and the U_k 's are independent this gives

$$\mathbb{P}(N = 0) = \mathbb{E}\left[\prod_{k \geq 0} (1 - \psi(W^\uparrow(k)))\right].$$

Let $a \geq 0$: introducing $\varphi(a) = -\log(1 - \psi(a))$ and $L^\uparrow(a) = \sum_{k \geq 0} \mathbb{1}_{\{W^\uparrow(k)=a\}}$, the local time at level a , we obtain

$$\mathbb{P}(N = 0) = \mathbb{E}\left[\exp\left(-\sum_{a \geq 0} \varphi(a)L^\uparrow(a)\right)\right]. \tag{17}$$

Lemma 5.4 *The quantity $\sup_{a \geq 0} \mathbb{E}(L^\uparrow(a))$ is finite. In particular, $\mathbb{P}(N = 0) > 0$.*

Proof Let W^- be a random walk with step size distribution $\xi - 1$, started at 0 and independent from W^\uparrow , and for $k \in \mathbb{Z}$ define $W^*(k) = W^\uparrow(k)$ if $k \geq 0$ and $W^*(k) = W^-(-k)$ if $k \leq 0$. Thus, defining $L^*(a) = \sum_{k \in \mathbb{Z}} \mathbb{1}_{\{W^*(k)=a\}}$, we have the obvious inequality $L^\uparrow(a) \leq L^*(a)$, and so we only have to prove that $\sup_{a \in \mathbb{Z}} \mathbb{E}(L^*(a))$ is finite.

It is clear that L^* stays the same if W^* is shifted in time, and that shifting L^* in time amounts to shifting W^* in space. Moreover, for any $w \in \mathbb{Z}$ the process $(W^*(k) + w, k \in \mathbb{Z})$ shifted at the time of its last visit to 0 is equal in distribution to W^* . Combining these facts, we see that L^* is a stationary sequence and in particular, $\sup_{a \in \mathbb{Z}} \mathbb{E}(L^*(a)) = \mathbb{E}(L^*(0))$. By the strong Markov property, $L^*(0)$ is a geometric random variable with parameter the probability that a random walk started at -1 and with step size distribution $\xi - 1$ never visits 0. Since $\mathbb{E}\xi < 1$, this probability is strictly positive and so $\mathbb{E}(L^*(0))$ is finite. This proves the finiteness of $\sup_a \mathbb{E}(L^\uparrow(a))$.

As for $\mathbb{P}(N = 0)$, we have

$$\mathbb{E}\left(\sum_{a \geq 0} \varphi(a)L^\uparrow(a)\right) \leq \sup_a \mathbb{E}(L^\uparrow(a)) \sum_{a \geq 0} \varphi(a),$$

and since $\varphi(a) \sim a^{-\beta}$ as $a \rightarrow +\infty$ and $\beta > 1$, the sum $\sum_a \varphi(a)$ is finite. In view of the last display, this implies that the random variable $\sum_a \varphi(a) L^\uparrow(a)$ is almost surely finite. This proves $\mathbb{P}(N = 0) > 0$ because of (17) and concludes the proof of the lemma. \square

We now prove that $\mathbb{P}(N \geq n) \rightarrow 0$ as $n \rightarrow +\infty$. Let W be a random walk with step size distribution $1 - \xi$ and $I = \inf_{k \geq 1} W(k)$. Then, W^\uparrow is by definition equal in distribution to W under $\mathbb{P}_0(\cdot \mid I \geq 1)$ (where the subscript refers to the initial state of W). Let moreover B_n be the time at which W advertizes a release for the n th time (which is well defined in the event $\{I \geq 1\}$), so that

$$\mathbb{P}(N = n) = b \mathbb{P}_0(B_n < +\infty, B_{n+1} = +\infty, I \geq 1),$$

with $b = 1/\mathbb{P}_0(I \geq 1)$. Writing the event $\{I \geq 1\}$ as the union between the two events $\{\inf_{1 \leq k \leq B_n} W(k) \geq 1\}$ and $\{\inf_{k > B_n} W(k) \geq 1\}$, the strong Markov property at time B_n entails

$$\mathbb{P}(N = n) = b \mathbb{E}_0 \left(p(W(B_n)); B_n < +\infty, \inf_{1 \leq k \leq B_n} W(k) \geq 1 \right),$$

with $p(w) = \mathbb{P}_w(N = 0, I \geq 1)$. Coupling W under \mathbb{P}_w with a version of W under \mathbb{P}_0 that stays below it, we see that $p(w)$ is increasing in w and so

$$\mathbb{P}(N = n) \geq bp(0) \mathbb{P}_0 \left(B_n < +\infty, \inf_{1 \leq k \leq B_n} W(k) \geq 1 \right) \geq bp(0) \mathbb{P}_0(B_n < +\infty, I \geq 1).$$

This last lower bound is equal to $p(0) \mathbb{P}_0(B_n < +\infty \mid I \geq 1)$ which by definition is equal to $p(0) \mathbb{P}(N \geq n)$. Since $p(0) = \mathbb{P}(N = 0)/b$ is > 0 by Lemma 5.4, dividing by $p(0)$ leads to

$$\mathbb{P}(N \geq n) \leq \frac{b \mathbb{P}(N = n)}{\mathbb{P}(N = 0)}.$$

Since $\mathbb{P}(N = n) \rightarrow 0$, this finally proves (16) and hence Proposition 5.1.

5.1.3 Corollary to Proposition 5.1

From Proposition 5.1, which is concerned with the behavior of one active queue, one can deduce the behavior of $\mathbf{A}(T^*)$ and T^* . The intuition behind the following result is clear: Proposition 5.1 shows that an active queue only advertizes releases when it is of order 1. Thus when considering several active queues and the initial state of at least one of them blows up, by the time the last queue to advertize a release does so, all the other queues are already in their stationary regime. From there on, it thus only takes an additional (random) time of order 1 for all the queues to advertize a release simultaneously. We therefore omit the proof of the following result, which consists of a direct translation of this intuition and can be formalized using similar arguments

as in the proof of Lemma 4.2 (it is actually easier, since we make no claim about the mean behavior of the random variables involved).

Corollary 5.5 *Assume that $\rho < 2$ and that $\beta > 1$, and consider any sequence of initial states $\mathbf{a}_n = (a_{n,r})$ such that $\max_r a_{n,r} \rightarrow +\infty$: then $(\mathbf{A}(T^*), T^* - \tau_{\max})$ under $\mathbb{P}_{\mathbf{a}_n}$ converges weakly as $n \rightarrow +\infty$ to a non-degenerate random variable.*

5.2 Simulation results

In this subsection we present simulation results which complement the theoretical result of the previous subsection and support our claim that Theorem 4.1 remains valid for $\beta > 1$. Based on simulation, we discuss stability issues and the asymptotic behavior of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$ in the heavy-traffic regime $\rho \uparrow 1$.

Proposition 4.4 and Theorem 4.1 show that the stability of the system and the order of magnitude of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$ are not affected by the precise value of R or by the precise distribution of ξ (as long as $R \geq 2$ and $\mathbb{E}(\xi^2) < +\infty$). Proposition 5.1 suggests a similar result, and also that the precise distribution of ζ will not matter as far as stability, and the order of magnitude of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$ are concerned. We thus performed our simulation with ξ a geometric random variable with parameter $2/(2 + \rho)$ and $\mathbb{P}(\zeta = 1) = 1$. We ran simulations using different distributions for ξ and ζ as well, including extreme cases such as distributions with infinite third moment (even infinite second moment for ζ), that yielded qualitatively similar results.

5.2.1 Simulation results for stability

It is straightforward to prove that \mathbf{Q} is transient if $\rho > 1$, irrespective of the value of $\beta \in (0, \infty]$. Indeed, if $\tilde{Q}(t)$ is the number of packets at the beginning of the t th time slot, \tilde{Q} is lower bounded by a random walk with drift $R(\rho - 1)$. Thus when $\rho > 1$, we have $\tilde{Q}(t) \rightarrow +\infty$ as $t \rightarrow +\infty$ and since $(\|\mathbf{Q}(k)\|, k \geq 0)$ is a subsequence of $(\tilde{Q}(t), t \geq 0)$ this proves the transience of \mathbf{Q} in this case.

For $\rho < 1$ we present a heuristic argument which suggests that \mathbf{Q} is stable. If the active queues are in state $\mathbf{a} = (a_r)$ with $a_r > 0$ for every r , the variation of the mean number of packets in the system over the next time slot is equal to

$$-R \left(1 - \rho - z \sum_{r=1}^R \psi(a_r) \right)$$

(remember that $z = \mathbb{E}\zeta$). Since $\psi(a) \rightarrow 0$ as $a \rightarrow +\infty$, the drift is negative, close to $-R(1 - \rho)$, when each a_r is large enough. The problem in formalizing this argument is twofold: first, in order to prove stability, one must be able to control every possible initial configuration, not only those where every a_r is large (note however that because of symmetry, it is natural to expect all a_r 's to be large or small simultaneously); second, this argument considers the system on the normal time scale, whereas we are interested in the system embedded at switching times.

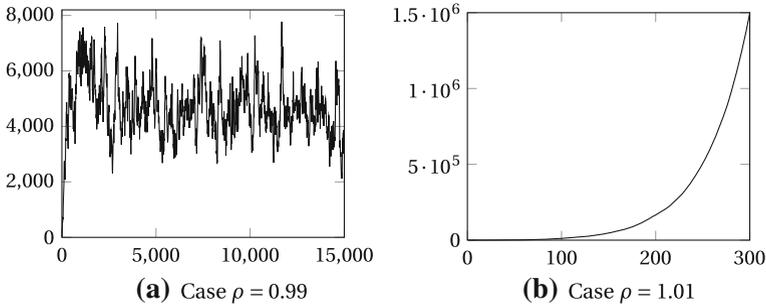


Fig. 2 Total number of packets at switching times: $\|\mathbf{Q}(k)\|$ versus k for $R = \beta = 2$ and **a** $\rho = 0.99$; **b** $\rho = 1.01$

In addition to this heuristic argument, simulation results strongly support stability of the system for $\rho < 1$. Indeed, Fig. 2 shows the evolution of $\|\mathbf{Q}(k)\|$, starting at $\mathbf{Q}(0) = \mathbf{0}$, in the case $R = \beta = 2$ and for two values of ρ : $\rho = 0.99$ and $\rho = 1.01$. When $\rho = 0.99$, Fig. 2a shows that $\|\mathbf{Q}(k)\|$ fluctuates between 2,000 and 8,000 for values of k up to 15,000. In contrast, Fig. 2b shows for $\rho = 1.01$ that $\|\mathbf{Q}(k)\|$ increases until we stop the simulation when $1.5 \cdot 10^6$ packets are present in the system, which only takes about 300 switching times. Note that for a transient system we would expect that, when the queues are large, the total queue size grows by a constant amount on average in every time slot, and that the time between two consecutive switching times increases. This explains the super-linear growth of $\|\mathbf{Q}(k)\|$ as we consider the system at switching times. In fact, the reasoning in Sect. 3 suggests that when $\|\mathbf{Q}(k)\|$ is large we have $\|\mathbf{Q}(k+1)\|/\|\mathbf{Q}(k)\| \approx \mathbb{E}\xi/(1 - \mathbb{E}\xi)$. This gives a heuristic explanation for the exponential growth (at rate $\rho/(2 - \rho) > 1$) observed in Fig. 2b. All in all, this is in line with the known transience of \mathbf{Q} for $\rho = 1.01$.

In general, it may be difficult to distinguish between positive recurrent and transient systems based on simulation results. Here however, \mathbf{Q} obeys two clearly distinguishable types of behavior: stochastic fluctuations on a long time interval when $\rho < 1$ (Fig. 2a), and almost deterministic exponential growth when $\rho > 1$ (Fig. 2b). This phase transition is a further indication of the stability and instability of the system for $\rho < 1$ and $\rho > 1$, respectively.

5.2.2 Simulation results for the asymptotic behavior of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$

Having discussed the stability of \mathbf{Q} when $\rho < 1$, we now discuss the asymptotic behavior of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$. To quantify its growth, it is convenient to introduce the notion of *scaling exponent*.

Definition 5.6 We call scaling exponent the number α such that

$$\lim_{\rho \uparrow 1} \left(\frac{\log \mathbb{E}_\infty(\|\mathbf{Q}(0)\|)}{\log(1/(1 - \rho))} \right) \stackrel{\text{def.}}{=} \alpha. \quad (18)$$

Our interest in the scaling exponent comes from an expected polynomial growth of $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$ in heavy traffic. In general, we expect as $\rho \uparrow 1$ a behavior of the kind

$$\mathbb{E}_\infty(\|\mathbf{Q}(0)\|) \approx \frac{C}{(1-\rho)^\alpha} \quad (19)$$

for some finite constant $C > 0$. Note that such a behavior would be stronger than (18), but that it is consistent with the result of Theorem 4.1. We have proved in the case $\beta = +\infty$ that the scaling exponent exists, and is equal to 2. In general we will assume that it exists and will be interested in discussing its value.

The scaling exponent depends on the four model parameters R, ξ, ζ and β . However, as explained in the beginning of Sect. 5.2, we will mostly be interested in its dependence on β and thus write $\alpha(\beta)$ when the other three parameters are kept fixed. In fact, our results suggest that α does not depend on these other parameters. As explained in Sect. 3, our main claim is that the lingering effect makes $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|)$ grow as $1/(1-\rho)^2$ when $\beta > 1$. Formally, this amounts to conjecture that the scaling exponent (exists and) satisfies $\alpha(\beta) = 2$ for $\beta > 1$.

In addition to the heuristic arguments presented in Sect. 3, we now present some simulation results supporting this claim. From the results in Fig. 2 we find by averaging over time $\mathbb{E}_\infty(\|\mathbf{Q}(0)\|) \approx 4.700$ for $\rho = 0.99$, corresponding in view of the definition (18) of the scaling exponent to the estimate $\alpha(2) \approx \log 4.700 / \log 100 \approx 1.84$. To facilitate the discussion we define in the sequel

$$F(\rho, \beta) = \frac{\log \mathbb{E}_\infty(\|\mathbf{Q}(0)\|)}{\log(1/(1-\rho))}, \quad \beta \in (0, \infty], \quad \rho < 1. \quad (20)$$

Then, the scaling exponent $\alpha(\beta)$ is defined in (18) via a limiting procedure, namely $\alpha(\beta) = F(1-, \beta) = \lim_{\rho \uparrow 1} F(\rho, \beta)$. Using the results of Fig. 2a to estimate $\alpha(2)$ amounts for using the approximation $F(1-, 2) \approx F(0.99, 2)$. In order to check whether this approximation is valid, we performed the same simulation for different values of ρ . In Fig. 3 the value of $F(\rho, 2)$ is plotted for $\rho \in (0.87, 0.999)$. We plotted $F(\rho, 2)$ against $\log(1/(1-\rho))$ in order to “dilate” time around the value $\rho = 1$ that we are interested in and also because it is natural to regress $F(\rho, 2)$, as function of ρ , against $\log(1/(1-\rho))$ as we discuss now.

Figure 3 shows that the limit $F(1-, 2)$ seems to exist, but that $F(\rho, 2)$ is still significantly increasing for $\rho = 0.999$. Thus $F(0.999, 2)$, and in particular $F(0.99, 2)$, cannot be used as an accurate estimate of $\alpha(2)$. It is numerically difficult to run a simulation for even higher values of ρ , and so to circumvent this problem, we use the simulation results displayed in Fig. 3 to find the asymptotic value of $F(\rho, 2)$. To do so, we use the approximation (19) to infer the form of $F(\rho, \beta)$, namely

$$F(\rho, \beta) \approx \alpha(\beta) + \frac{\log C}{\log(1/(1-\rho))}, \quad (21)$$

which suggests, as mentioned above, to regress $F(\rho, 2)$ against $a + b/x$ in the scale $x = \log(1/(1-\rho))$. We performed this regression for the curve displayed in Fig. 3 and

Fig. 3 Approximating $\alpha(2)$ for $R = 2$. The *solid line* represents $F(\rho, 2)$ versus $\log(1/(1-\rho))$ obtained by simulation from the definition (20) of F ; the *dashed line* represents the best regression of $F(\rho, 2)$ against $x \mapsto a + b/x$ in the scale $x = \log(1/(1-\rho))$ [see (21)]

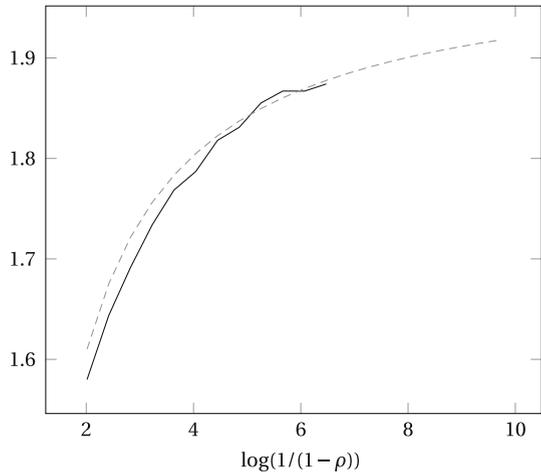
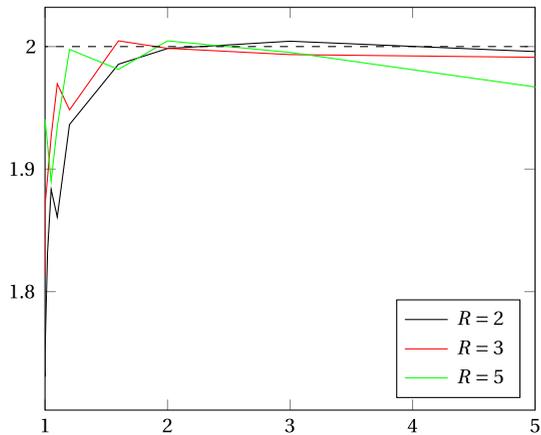


Fig. 4 $\alpha(\beta)$ (obtained by simulation) versus β for $R \in \{2, 3, 5\}$



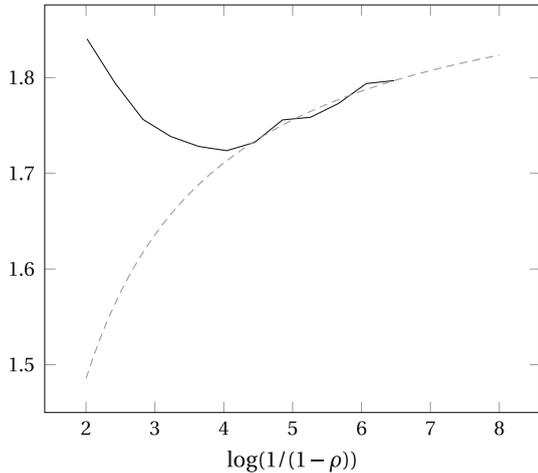
found an optimal value $a = 1.9984$, which gives an estimate for the scaling exponent $\alpha(2) \approx 1.9984$ in very good agreement with our conjecture.

Applying the same approach, we can find an estimate for $\alpha(\beta)$ for any value of β . The results for $R = 2, 3$ and 5 are given in Fig. 4, and confirm that R does not seem to influence $\alpha(\beta)$. Further, the approximation $\alpha(\beta) \approx 2$ appears to be very good for any $\beta > 1.2$, namely, the estimated value of $\alpha(\beta)$ is at most 3% away from 2 for any $\beta > 1.2$ and any $R = 2, 3, 5$.

5.3 More on the case $\beta > 1, \beta \approx 1$

The previous simulation results show a rather fuzzy behavior of $\alpha(\beta)$ for β close to 1. Indeed, the curves shown in Fig. 4 seem stable for $\beta > 1.2$ but much less so as β gets closer to 1. We believe that this may be due to interesting heavy-tailed phenomena.

Fig. 5 Approximating $\alpha(1.2)$ for $R = 2$. The *solid line* represents $F(\rho, 2)$ versus $\log(1/(1-\rho))$; the *dashed line* represents the best regression obtained by regressing $F(\rho, 2)$ past its infimum as in (21)



The shape of the function $\rho \mapsto F(\rho, \beta)$ depicted in Fig. 3 is typical for large values of β , say $\beta > 1.2$. In particular, this function is increasing which makes the regression of F against $1/\log(1/(1-\rho))$, such as in (21), reasonable. However, as β gets closer to one, the shape of this function changes. For instance, Fig. 5 shows simulation results for $F(\rho, 1.2)$ which are representative of $F(\rho, \beta)$ for small β , say $1 < \beta < 1.2$. Noticeably, the function $F(\rho, 1.2)$ is not monotone in ρ and so the approximation (21) cannot be valid for every ρ . Rather, we find that $F(\rho, \beta)$ decreases and then increases, and that the regression against $1/\log(1/(1-\rho))$ is only accurate past the minimum.

Regressing the curve obtained in Fig. 5 past the minimum leads to the approximation $\alpha(1.2) \approx 1.94$, which is still very much in line with the conjecture $\alpha(1.2) = 2$. However, the point where the minimum of the function $\rho \mapsto F(\rho, \beta)$ is attained shifts to the right when β gets closer to 1, leaving us with less points against which to regress. In particular, for $\beta < 1.2$ one would need to simulate the system at loads higher than 0.999 to get accurate results.

We suspect that this numerical instability is caused by heavy-tailed phenomena that seem to appear for $\beta < 2$. More precisely, inspecting the proofs of Proposition 4.4 and Theorem 4.1, one sees that Corollary 5.5 is not strong enough for the proofs of the case $\beta = +\infty$ to go through directly. Indeed, instead of controlling the behavior of $\mathbf{A}(T^*)$ and T^* in distribution, one needs (at least, with the proposed proof strategy) to control their mean behavior. Let us do a small computation: let T_∞^Z be the random variable introduced in Lemma 5.3, which is the weak limit of $T_1 - \tau_1$, and let B^\uparrow be the first time the process W^\uparrow advertizes a release. Then

$$\mathbb{E}(|T_\infty^Z|; T_\infty^Z \leq 0) = \mathbb{E}(B^\uparrow; B^\uparrow < +\infty) = \sum_{k \geq 0} k \mathbb{P}(B^\uparrow = k),$$

and as before, we have

$$\mathbb{P}(B^\uparrow = k) = \mathbb{E} \left[\psi(W^\uparrow(k)) \prod_{i < k} (1 - \psi(W^\uparrow(i))) \right].$$

The arguments of Sect. 5.1 show that the random variable $\prod_{i \geq 0} (1 - \psi(W^\uparrow(i)))$ is finite. In view of the last expression, it therefore seems that the tail behavior of B^\uparrow is dictated by $\psi(W^\uparrow(k))$ as $k \rightarrow +\infty$ and since $W^\uparrow(k) \approx k(1 - \mathbb{E}\xi)$ by the law of large numbers, we should obtain $\mathbb{P}(B^\uparrow = k) \approx k^{-\beta}$ as $k \rightarrow +\infty$. When $\beta < 2$, this suggests that B^\uparrow , and in particular $|T_\infty^Z|$, has infinite mean, although both random variables are almost surely finite for $\beta > 1$. It is possible to use this result to show that the (almost surely finite) weak limits (as the initial state grows to $+\infty$) of $\mathbf{A}(T^*)$ and $T^* - \tau_{\max}$ have infinite first moment. This potentially invalidates the back-of-the-envelope computations of Sect. 3, and so more care is needed. For instance, simulation experiments for $\beta = 1.2$ and $R = 2$ suggest that $\mathbb{E}_{a_1}(A_1(T^*))$ grows as $a_1^{0.4}$ as $a_1 \rightarrow +\infty$, whereas $A_1(T^*)$ converges weakly to a finite random variable.

6 Concluding remarks and suggestions for further research

Motivated by the poor delay performance of “cautious” activation rules in queue-based schemes for distributed medium access control, we investigated more aggressive schemes. Our main contribution lies in highlighting a new effect that we called the lingering effect and in studying the performance ramifications of this effect for a special topology. In this section we explain and discuss various directions in which our framework could possibly be extended, and also relate our results to those of [9].

First of all, it would be straightforward to extend our results to the following asymmetric case: instead of having R queues in each group with identically distributed arrival processes across queues, we have two groups of R_1 and R_2 queues and the arrivals into the k th queue of group g have distribution $\xi_{g,k}$. We chose to study a symmetric scenario for technical reasons, since then there is no need to label queues individually. In this setting, the lingering effect will occur whenever, informally speaking, the two dominant queues of at least one of the two groups have the same arrival rate. For instance, the delay will scale like $1/(1-\rho)^2$ with $\rho = \max_k \mathbb{E}(\xi_{1,k}) + \max_k \mathbb{E}(\xi_{2,k})$ if the condition $\mathbb{E}(\xi_{1,1}) = \mathbb{E}(\xi_{1,2}) \geq \max_{g,k} \mathbb{E}(\xi_{g,k})$ is satisfied.

We believe that the insights provided by the complete bipartite constraint graph carry over to more general topologies. Note that for a general topology, our model needs to be amended, since one needs then to specify more precisely how queues become active. One can for instance think of queues going into back-off, and then trying to acquire the shared resource at some rate. We conjecture that whenever the constraint graph is not complete, and thus contains an independent set of several nodes, the lingering effect can rear its head, provided some algebraic condition between the arrival rates at the various queues is satisfied. It would be very interesting to be able to formulate a precise and formal conjecture reflecting this intuition, and most probably quite challenging to prove it.

To conclude, we would like to put our results in perspective by linking them to those of [9]. On the one hand, our results show that the delay performance may be greatly improved by considering more aggressive schemes (compared to maximally stable algorithms known earlier). On the other hand, Ghaderi et al. [9] have shown that, by doing so, one may lose a fraction of the stability region. But in the example analyzed in [9], one only loses a small fraction of the stability region, namely about

3 %. Thus, our work suggests a potentially interesting trade-off between stability and delay performance, namely that more aggressive schemes may lead to a small loss of stability, whereas the delay in this stability region is much better than the delay for a maximum stable algorithm.

Acknowledgments This work was supported by Microsoft Research through its PhD Scholarship Programme and a TOP Grant from NWO.

References

1. Boon, M.A.A., Adan, I.J.B.F., Winands, E.M.M., Down, D.G.: Delays at signalized intersections with exhaustive traffic control. *Probab. Eng. Inf. Sci.* **26**(3), 337–373 (2012)
2. Bouman, N., Borst, S., and Van Leeuwen, J.: Achievable delay performance in CSMA networks. In: *Proceedings of Allerton Conference*, pp. 384–391 (2011)
3. Chaporkar, P., Sarkar, S.: Stable scheduling policies for maximizing throughput in generalized constrained queueing. In: *Proceedings of Infocom* (2006)
4. Coffman Jr, E.G., Puhalskii, A.A., Reiman, M.I.: Polling systems with zero switchover times: a heavy-traffic averaging principle. *Ann. Appl. Probab.* **5**(3), 681–719 (1995)
5. Feuillet, M., Proutiere, A., Robert, P.: Random capture algorithms: fluid limits and stability. In: *Proceedings of ITA Workshop* (2010)
6. Foster, J.H.: A limit theorem for a branching process with state-dependent immigration. *Ann. Math. Stat.* **42**, 1773–1776 (1971)
7. Gamarnik, D., Zeevi, A.: Validity of heavy traffic steady-state approximation in generalized Jackson networks. *Ann. Appl. Probab.* **16**(1), 56–90 (2006)
8. Georgiadis, L., Neely, M.J., Tassiulas, L.: Resource allocation and cross-layer control in wireless networks. *Found. Trends Netw.* **1**, 1–144 (2006)
9. Ghaderi, J., Borst, S.C., and Whiting, P.A.: Backlog-based random-access in wireless networks: fluid limits and instability issues. In: *Proceedings of IEEE WiOpt Conference* (2012)
10. Jiang, L., Leconte, M., Ni, J., Srikant, R., Walrand, J.: Fast mixing of parallel glauber dynamics and low-delay CSMA scheduling. *IEEE Trans. Inf. Theory* **58**(10), 6541–6555 (2012)
11. Jiang, L., and Walrand, J.: A distributed CSMA algorithm for throughput and utility maximization in wireless networks. In: *Proceedings of Allerton Conference* (2008)
12. Jiang, L., Walrand, J.: Stability and delay of distributed scheduling algorithms for networks of conflicting queues. *Queueing Syst.* **72**(1), 161–187 (2012)
13. Jiang, L., Shah, D., Shin, J., Walrand, J.: Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inf. Theory* **56**(12), 6182–6207 (2010)
14. Marbach, P., and Eryilmaz, A.: A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In: *Proceedings of Allerton Conference* (2008)
15. Modiano, E., Shah, D., and Zussman, G.: Maximizing throughput in wireless networks via gossiping. In: *Proceedings of ACM SIGMETRICS/Performance Conference*, pp. 27–38 (2006)
16. Ni, J., Tan, B., and Srikant, R.: Q-CSMA: queue length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In: *Proceedings of IEEE INFOCOM Mini-Conference* (2010)
17. Rajagopalan, S., Shah, D., and Shin, J.: Network adiabatic theorem: an efficient randomized protocol for content resolution. In: *Proceedings of ACM SIGMETRICS/Performance Conference*, pp. 133–144 (2009)
18. Resing, J.A.C.: Polling systems and multitype branching processes. *Queueing Syst. Theory Appl.* **13**(4), 409–426 (1993)
19. Shah, D., Shin, J., and Tetali, P.: Medium access using queues. In: *Proceedings of IEEE FOCS*, pp. 698–707 (2011)
20. Shah, D., Tse, D.N.C., Tsitsiklis, J.N.: Hardness of low delay network scheduling. *IEEE Trans. Inf. Theory* **57**(12), 7810–7817 (2011)
21. Shah, D., Shin, J.: Randomized scheduling algorithm for queueing networks. *Ann. Appl. Probab.* **22**(1), 128–171 (2012)

22. Shah, D., Wischik, D.: Switched networks with maximum weight policies: fluid approximation and multiplicative state space collapse. *Ann. Appl. Probab.* **22**(1), 70–127 (2012)
23. Sharma, G., Mazumdar, R.R., and Shroff, N.B.: On the complexity of scheduling in wireless networks. In: *Proceedings of MobiCom*, pp. 227–238 (2006)
24. Sharma, G., Shroff, N.B., Mazumdar, R.R.: Joint congestion control and distributed scheduling for throughput guarantees in wireless networks. In: *Proceedings of Infocom* (2007)
25. Simatos, F., Bouman, N., and Borst, S.: Linger issues in distributed scheduling. In: *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer, SIGMETRICS '13*, pp. 141–152, New York, NY, ACM (2013)
26. Stolyar, A.L.: Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.* **14**, 1–53 (2004)
27. Subramanian, V., and Alanyali, M.: Delay performance of CSMA in networks with bounded degree conflict graphs. In: *Proceedings of IEEE ISIT*, pp. 2373–2377 (2011)
28. Tassiulas, L.: Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In: *Proceedings of Infocom*, pp. 533–539 (1998)
29. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Autom. Control.* **37**, 1936–1948 (1992)
30. Tassiulas, L., Ephremides, A.: Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inf. Theory* **39**, 466–478 (1993)
31. Wu, X., Srikant, R.: Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. In: *Proceedings of Infocom* (2006)