

Tail Asymptotics for Discrete Event Systems

Marc Lelarge

Received: 16 January 2007 / Accepted: 18 January 2008
© Springer Science + Business Media, LLC 2008

Abstract In the context of communication networks, the framework of stochastic event graphs allows a modeling of control mechanisms induced by the communication protocol and an analysis of its performances. We concentrate on the logarithmic tail asymptotics of the stationary response time for a class of networks that admit a representation as (max,plus)-linear systems in a random medium. We are able to derive analytic results when the distribution of the holding times are light-tailed. We show that the lack of independence may lead in dimension bigger than one to non-trivial effects in the asymptotics of the sojourn time. We also study in detail a simple queueing network with multipath routing.

Keywords Queueing theory · Large deviations

1 Introduction

In this paper, we study tail asymptotics of the form

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\theta^* < 0,$$

where the random variable Z corresponds to a “global” state variable associated to a (max,plus)-linear system. We only deal with light-tailed distributions, i.e. distribution functions that decay exponentially fast. The simplest example of random variable covered by our results is the stationary waiting time in a FIFO $GI/GI/1$ queue.

This work was partially done while the author was with Boole Centre for Research in Informatics, Science Foundation Ireland Research Grant No. SFI 04/RP1/I512.

M. Lelarge (✉)
ENS-INRIA, 45 rue d’Ulm, 75005 Paris, France
e-mail: marc.lelarge@ens.fr
URL: <http://www.di.ens.fr/~lelarge>

This case has been extensively studied in the literature and much finer estimates are available, see the complementary works Iglehart (1972) and Pakes (1975).

In recent years, there has been some interest in extending this result to networks of queues. Chang (1995) considers anintree network and uses large deviations techniques to show that the queue length distributions have an exponentially decaying tail. Ganesh and Anantharam (1996) obtains the decay rate of the tail distribution for two exponential server queues in series fed by renewal arrivals. In Bertsimas et al. (1998), the decay rate of the stationary waiting time and queue length distributions at each node are computed in an acyclic network in the context of quite general arrival and service processes. Literature on large deviations of queueing networks with feedback is rare and confined to the setting of networks described by finite-dimensional Markov processes, see Dupuis and Ellis (1995), Dupuis et al. (1991) and the recent works Ignatiouk-Robert (2000), Ignatiouk-Robert (2005). Moreover, these works concentrate on local large deviations and cannot handle the large deviations of the network in its stationary regime. The large deviations asymptotics of queueing systems are difficult to analyze because they are dynamical systems with discontinuities. To the best of our knowledge, there is no rigorous result on the large deviations of non-exponential networks with feedback in their stationary regime.

In this paper, we consider a class of networks that admit a (max,plus)-linear representation. This class contains the stochastic event graphs (which can be used to model window-based congestion control mechanism like TCP) and hence our results give the tail asymptotics of the steady state end-to-end response times of these networks. We should stress that the results of this paper are not restricted to this subclass and we give an example of a network with multipath routing that is covered by our framework.

From a mathematical point of view, we study Z the stationary solution of a (max,plus)-linear recursion. For example, in the case of Lindley's recursion, Z is the stationary workload of a single server queue. In a multi-dimensional setting, Z will be the biggest coordinate of the solution vector. For example, in the case of queues in tandem, Z will be the stationary end-to-end delay. The main contribution of this paper is to derive the tail asymptotics for Z using moment generating functions. Precise results concerning large deviations of products of random topical operators have been obtained in Toomey (2002). However very restrictive conditions are required on the coefficients of the matrix and only the irreducible case is studied in Toomey (2002). Here we do not assume these requirements to be fulfilled and we show that under mild assumptions on the matrix structure, an expression for the tail behavior of Z is explicitly given and can be computed (or approximated) in practical cases.

In the next section, we first give the general (max,plus) framework (Baccelli et al. 1992), with some examples of queueing networks. Then we give the stochastic assumptions and the tail asymptotics of the stationary solution of the (max,plus)-linear recursion is derived in Theorem 1 which clearly extends the case of the single server queue. Theorem 2 gives a more explicit form of the exponential decay rate.

In Section 3, we study two queueing applications. First we consider a system of two queues in tandem and show that when the service times at both queues are identical, then depending on the value of the intensity of the arrival process there is a phase transition in the behavior of the network reaching a large end-to-end delay (Proposition 2). Then we study in detail a simple example of queueing networks

with resequencing. Multipath routing has recently received some attention in the context of both wired and wireless communication networks. By sending data packets along different paths, multipath routing can potentially help balance the traffic load and reduce congestion levels in the network, thereby resulting in lower end-to-end delay. We show how our framework can model such mechanisms and give analytical insights.

Sections 4, 5 contain the proofs of Theorem 1 and Theorem 2 respectively. We give some further research directions in the conclusion.

2 Tails for discrete event systems

In this paper we consider open systems with a single input marked point process $N = \{(T_n, A_n, B_n)\}_{-\infty < n < \infty}$, where in a queueing context the sequence $\{T_n\}$ describes the customer arrival times in the network and $\{A_n, B_n\}$ carries the information related to the n -th customer (like its service time at the different stations, routing decisions). We give a precise description of the dynamic of the system in the next section and of the random variable Z for which we derive the logarithmic tail asymptotics.

2.1 (Max, plus)-linear systems

The (max, plus) semi-ring \mathbb{R}_{\max} is the set $\mathbb{R} \cup \{-\infty\}$, equipped with max, written additively (i.e., $a \oplus b = \max(a, b)$) and the usual sum, written multiplicatively (i.e., $a \otimes b = a + b$). The zero element is $-\infty$. For matrices of appropriate sizes, we define $(A \oplus B)^{(i,j)} = A^{(i,j)} \oplus B^{(i,j)} := \max(A^{(i,j)}, B^{(i,j)})$ and $(A \otimes B)^{(i,j)} = \bigoplus_k A^{(i,k)} \otimes B^{(k,j)} := \max_k (A^{(i,k)} + B^{(k,j)})$. By convention if A is a matrix and $c \in \mathbb{R}_{\max}$, then $(A \otimes c)^{(i,j)} := A^{(i,j)} \otimes c$.

Let s be an arbitrary fixed natural number. We assume that we are given with a sequence of matrices with non-negative coefficients: A_n of size $s \times s$ and B_n of size $s \times 1$. To the sequences $\{A_n\}_n, \{B_n\}_n$, and $\{T_n\}_n$, we associate the following (max, plus)-linear recurrence:

$$\mathcal{X}_{n+1} = A_{n+1} \otimes \mathcal{X}_n \oplus B_{n+1} \otimes T_{n+1}, \tag{1}$$

where $\{\mathcal{X}_n, n \in \mathbb{Z}\}$ is a sequence of state variables of dimension s . In Examples 1, 2, 3, we derive the explicit form of this recursion for the single server queue, queues in tandem, and a fork join system. We refer to these examples to get an interpretation of the various quantities.

The stationary solution of this equation is constructed as follows. We write

$$Y_{[m,n]} := \bigoplus_{m \leq k \leq n} D_{[k+1,n]} \otimes B_k \otimes T_k, \tag{2}$$

where for $k < n$, $D_{[k+1,n]} = \bigotimes_{j=k+1}^{n-1} A_j = A_n \otimes \dots \otimes A_{k+1}$ and $D_{[n+1,n]} = E$, the identity matrix (the matrix with all its diagonal elements equal to 1 and all its

non-diagonal elements equal to $-\infty$). It is readily checked that $Y_{[m,m]} = B_m \otimes T_m$, and for all $n \geq m$,

$$Y_{[m,n+1]} = A_{n+1} \otimes Y_{[m,n]} \oplus B_{n+1} \otimes T_{n+1}.$$

In view of Eq. 2, the sequence $\{Y_{[-n,0]}\}$ is non-decreasing in n , so that we can define the stationary solution of Eq. 1,

$$Y_{(-\infty,0]} := \lim_{n \rightarrow \infty} Y_{[-n,0]} \leq \infty.$$

We define the *stationary maximal dater* by

$$0 \leq Z := Z_{(-\infty,0]} = \bigoplus_{1 \leq i \leq s} Y_{(-\infty,0]}^{(i)} - T_0 \leq \infty. \tag{3}$$

The following expression for the stationary maximal dater shows the similitude with the traditional stationary workload of a single server queue:

$$Z = \sup_{n \leq 0} \{S_n - (T_0 - T_n)\}, \tag{4}$$

where the process $\{S_n\}_{n \leq 0}$ is defined by

$$S_n := \bigoplus_{1 \leq i \leq s} (D_{[n+1,0]} \otimes B_n)^{(i)}. \tag{5}$$

Example 1 Consider a FIFO single server queue where T_n is the arrival time of the n -th customer and σ_n is its service time. Equation 1 is then the standard Lindley's recursion,

$$\begin{aligned} \mathcal{X}_{n+1} &= \max(\mathcal{X}_n + \sigma_{n+1}, T_{n+1} + \sigma_{n+1}) \\ &= \sigma_{n+1} \otimes \mathcal{X}_n \oplus \sigma_{n+1} \otimes T_{n+1}. \end{aligned}$$

The interpretation of \mathcal{X}_n is the departure time of the n -th customer from the queue. Note that in this case, we have $S_n = \sum_{i=n}^0 \sigma_i$ and Z is the stationary workload.

Example 2 Consider now a system of two queues in tandem, where T_n is the arrival time of the n -th customer in the system and $\sigma_n^{(i)}$ is its service time at queue i , for $i = 1, 2$. Then Eq. 1 is given by

$$\begin{pmatrix} \mathcal{X}_{n+1}^{(1)} \\ \mathcal{X}_{n+1}^{(2)} \end{pmatrix} = \begin{pmatrix} \sigma_{n+1}^{(1)} & -\infty \\ \sigma_{n+1}^{(1\otimes 2)} & \sigma_{n+1}^{(2)} \end{pmatrix} \begin{pmatrix} \mathcal{X}_n^{(1)} \\ \mathcal{X}_n^{(2)} \end{pmatrix} \oplus \begin{pmatrix} \sigma_{n+1}^{(1)} \\ \sigma_{n+1}^{(1\otimes 2)} \end{pmatrix} T_{n+1},$$

where we used the shorthand notation $\sigma_{n+1}^{(1\otimes 2)} = \sigma_{n+1}^{(1)} \otimes \sigma_{n+1}^{(2)}$. In this case $\mathcal{X}_n^{(1)}$ is the departure time of the n -th customer from the first queue and $\mathcal{X}_n^{(2)}$ from the second queue. Hence Z is the stationary end-to-end delay of the network. Note that in this case, we have

$$S_n = \sup_{n \leq \ell \leq 0} \sum_{i=n}^{\ell} \sigma_i^{(1)} + \sum_{j=\ell}^0 \sigma_j^{(2)}. \tag{6}$$

Example 3 Consider the standard fork and join system as depicted (with Petri net formalism) in Fig. 1. In this model, each time a packet (say k) finishes its service $\sigma_k^{(1)}$ in node 1, there is one packet sent up and one packet sent down simultaneously. The ‘up’-packet (‘down’-packet) is then also the k -th packet for node 2 (for node 3 respectively). The k -th packet joins the queue of node 4 once both packets have left node 2 and 3 respectively. Each node is a standard $\cdot/G/1/\infty$ queue.

Let $\mathcal{X}_n^{(i)}$ denote the departure time of the n -th packet from node i . We have the following equations:


$$\begin{aligned} \mathcal{X}_{n+1}^{(1)} &= (T_{n+1} \oplus \mathcal{X}_n^{(1)}) \otimes \sigma_{n+1}^{(1)}, \\ \mathcal{X}_{n+1}^{(2)} &= (\mathcal{X}_{n+1}^{(1)} \oplus \mathcal{X}_n^{(2)}) \otimes \sigma_{n+1}^{(2)}, \\ \mathcal{X}_{n+1}^{(3)} &= (\mathcal{X}_{n+1}^{(1)} \oplus \mathcal{X}_n^{(3)}) \otimes \sigma_{n+1}^{(3)}, \\ \mathcal{X}_{n+1}^{(4)} &= (\mathcal{X}_{n+1}^{(2)} \oplus \mathcal{X}_{n+1}^{(3)} \oplus \mathcal{X}_n^{(4)}) \otimes 0. \end{aligned}$$

This system is linear in the (max, plus) semi-ring \mathbb{R}_{\max} , and we can write the recursion 1 with the following matrices:

$$A_n = \begin{pmatrix} \sigma_n^{(1)} & -\infty & -\infty & -\infty \\ \sigma_n^{(1\otimes 2)} & \sigma_n^{(2)} & -\infty & -\infty \\ \sigma_n^{(1\otimes 3)} & -\infty & \sigma_n^{(3)} & -\infty \\ \sigma_n^{(1\otimes 2\oplus 3)} & \sigma_n^{(2)} & \sigma_n^{(3)} & 0 \end{pmatrix}, \quad B_n = \begin{pmatrix} \sigma_n^{(1)} \\ \sigma_n^{(1\otimes 2)} \\ \sigma_n^{(1\otimes 3)} \\ \sigma_n^{(1\otimes 2\oplus 3)} \end{pmatrix},$$

where we used the shorthand notations, $\sigma_n^{(i\otimes j)} = \sigma_n^{(i)} \otimes \sigma_n^{(j)}$ and $\sigma_n^{(i\otimes j\oplus k)} = \sigma_n^{(i)} \otimes (\sigma_n^{(j)} \oplus \sigma_n^{(k)})$. In this case Z is the stationary end-to-end delay of the network.

Fig. 1 Fork and join model



We refer to Baccelli et al. (1992) for other examples of (max,plus)-linear networks (see also Lelarge 2005 for an example showing how to model window control mechanism).

2.2 Tail asymptotics for the stationary solution

First we need to give conditions that ensure the stability of the system, i.e. that the limit 3 is finite. Hence we introduce the following assumption in order to apply first-order Theorems of Section 7.3 of Baccelli et al. (1992):

Assumption (ST) (Structure of (A_n, B_n))

The random variables A_n, B_n have a fixed structure, i.e. each entry of A_n or B_n is either a.s. equal to $-\infty$ or non-negative for all n . And each diagonal entry of A_n is non-negative.

Assumption (IA) (independence assumption)

We suppose that the sequences $\{(A_n, B_n)\}_n$ and $\{\tau_n := T_{n+1} - T_n\}_n$ are mutually independent and each of them consists of i.i.d. random variables with finite means.

This assumption implies a law of large numbers for $\{S_{-n}\}$ as defined in Eq. 5, namely,

$$\frac{S_{-n}}{n} \rightarrow_{n \rightarrow \infty} \gamma \quad \text{both a.s. and in } L_1, \tag{7}$$

where γ is a constant referred to as the top Lyapunov exponent of the sequence $\{A_n\}$ see Theorems 7.27 and 7.36 in Baccelli et al. (1992).

Assumption (S) (stability)

We assume that $\gamma < \mathbb{E}[\tau_1] = a$.

We have that under (IA) and (S) the maximal dater Z defined in Eq. 3 is almost surely finite.

We denote by $\mathbf{0}$ the vector with all its entries equal to 0.

Assumption (SP) (Separability)

We assume that we have for all n ,

$$A_n \otimes \mathbf{0} = B_n \oplus \mathbf{0}.$$

This assumption ensures that for a solution \mathcal{X}_n of Eq. 1 with any initial condition: if $\mathcal{X}_n \leq \mathbf{0} \otimes T_{n+1}$ then the process $\mathcal{X}_{n+1}, \mathcal{X}_{n+2}, \dots$ does not depend on the past $\mathcal{X}_n, \mathcal{X}_{n-1}, \dots$. Note that this assumption is clearly satisfied in the examples described above since we have $A_n \otimes \mathbf{0} = B_n$. We refer to Section 2.2.4 of Lelarge (2005) for an example of network with $A_n \otimes \mathbf{0} = B_n \oplus \mathbf{0} \neq B_n$. In fact, Propositions 3 and 4 of Lelarge (2005) show that any FIFO event graph with a single input fits into our framework. However this condition allows also to deal with some type of networks with (random) routing as described in Section 3.2. This property of separability can be made precise in a larger framework than (max,plus)-linear networks: this is the class of monotone separable networks introduced in Baccelli and Foss (1995).

We now give the stochastic assumptions that ensure that the random variable Z is light tailed:

Assumption (LT) (Light-tailed)

Let

$$\eta = \sup \left\{ \theta > 0, \bigoplus_i \mathbb{E} \left[e^{\theta B_i^{(i)}} \right] < \infty \right\}.$$

We assume that $\eta > 0$.

We will always assume that Assumptions **(ST)**, **(IA)**, **(S)**, **(SP)**, **(LT)** hold. We are now in position to state our main result.

Theorem 1 *The following limit exists as an extended real number:*

$$\Lambda_S(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta S_n} \right]. \tag{8}$$

We have

$$\theta^* = \sup \{ \theta > 0, \Lambda_S(\theta) + \Lambda_T(-\theta) < 0 \} > 0, \tag{9}$$

where $\Lambda_T(\theta) = \log \mathbb{E} \left[e^{\theta \tau_1} \right]$ and the tail asymptotics of Z is given by,

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\theta^*.$$

In the case of the single server queue, we have clearly $\Lambda_S(\theta) = \log \mathbb{E}[\exp(\theta \sigma_1)]$ and Theorem 1 extends a standard result of queueing theory that goes back to the work Cramér (1938) and in a queueing context to Iglehart (1972). Note that usually authors assume in addition to our assumptions the so-called Cramér condition. This result for the single server queue without this extra condition also follows from Lelarge (2007b). We will give more comments about this theorem in Section 3.

2.3 More detailed results

In this section we give a more explicit form for θ^* . Without loss of generality, we may assume that the matrices A_n have the following block structure:

$$\begin{pmatrix} A_n(1, 1) & | & -\infty & | & -\infty & | & -\infty \\ - & - & - & - & - & - & - \\ A_n(2, 1) & | & A_n(2, 2) & | & -\infty & | & -\infty \\ - & - & - & - & - & - & - \\ & & \vdots & & \vdots & & \vdots \\ - & - & - & - & - & - & - \\ A_n(d, 1) & | & A_n(d, 2) & | & & | & A_n(d, d) \end{pmatrix},$$

where each $A_n(\ell, \ell)$ is an irreducible matrix.

Theorem 2 Associated to the irreducible matrices $\{A_n(\ell, \ell)\}$, we define the following function:

$$\Lambda_\ell(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta(A_n(\ell, \ell) \otimes \dots \otimes A_1(\ell, \ell))^{(u,v)}} \right],$$

where the limit exists in $\mathbb{R} \cup \{\infty\}$ and is independent of u, v . Then we have $\theta^* = \min\{\eta, \theta^\ell\}$ where the θ^ℓ 's are defined as follows

$$\theta^\ell = \sup\{\theta > 0, \Lambda_\ell(\theta) + \Lambda_T(-\theta) < 0\}.$$

In the case of a single server queue with exponentially distributed service times, we have $\eta > \theta^*$ and this property remains valid for a large class of distributions. However, we show in the next section that as soon as we consider a network (i.e. with at least 2 nodes) then the parameter η can play a role even with exponentially distributed service times.

We first give a framework where η cannot play any role. Given a vector $v = (v^{(1)}, \dots, v^{(K)})$, we call a (max, plus) expression \mathcal{P} a polynomial in v of unit maximum degree if it has the form

$$\mathcal{P} = \bigoplus_j \bigotimes_{k \in \mathcal{K}_j} v^{(k)},$$

where $\mathcal{K}_j \subset [1, K]$.

Corollary 1 If there exists a sequence of random variables $\{\sigma_n = (\sigma_n^{(1)}, \dots, \sigma_n^{(K)})\}_n$ such that

- (1) The components of σ_n are independent of each other;
- (2) For all i , there exists k such that $A_n^{(k,k)} = \sigma_n^{(i)}$;
- (3) Each entry of A_n that is not 0 or $-\infty$ is a polynomial (in \mathbb{R}_{\max}) in σ_n of unit maximal degree.

Then we have $\theta^* = \min\{\theta^\ell\}$.

In a queueing context, the sequence of matrices $\{A_n(\ell, \ell)\}$ corresponds to a specific “component” of the network. It is well-known that the stability of such a network is constrained by the “slowest” component (Baccelli et al. 1992). Here we see that in a large deviations regime, if each component is independent of each other, then the “bad” behavior of the network is due to a “bottleneck” component (which is not necessarily the same as the “slowest” component in average).

Remark 1 In the framework of last Corollary, the tail asymptotics for Z under heavy-tailed (more precisely subexponential) assumptions [i.e. when Assumption (LT) is not satisfied] has been derived in Baccelli et al. (2004). In this case, the exact asymptotics (i.e. not in the logarithmic scale) are derived and the Lyapunov exponents of the sub-matrices $\{A_n(\ell, \ell)\}$ appear. These exponents are known to be hard to compute (Blondel et al. 2000). Similarly in the light-tailed case we see that the asymptotics in the logarithmic scale is given by the functions $\Lambda_\ell(\theta)$. The computation of these functions is not easy, in particular when the network has some feedback mechanism. More formally, we will see that this function is convex and its right-derivative at zero is exactly the Lyapunov exponent of the sub-matrices $\{A_n(\ell, \ell)\}$,

so that knowing the function $\Lambda_\ell(\theta)$ allows to determine the Lyapunov exponent. In particular, at the level of generality considered in this paper, our result cannot be made in a more explicit form. We will see in the next section several examples for which the value of θ^* has a simple expression in terms of the parameters of the problem. It is interesting to note that, as opposed to the heavy-tailed case, exact tail asymptotics for Z under light-tailed assumptions seem to be out of reach in the general framework of (max,plus)-linear networks.

Example 4 Going back to the fork and join system described in Example 3, we see that the irreducible matrices are one-dimensional and we have for $\ell = 1, 2, 3$,

$$\Lambda_\ell(\theta) = \log \mathbb{E} \left[e^{\theta \sigma_1^{(\ell)}} \right].$$

Hence if θ^ℓ denotes the exponential rate of decay for the single server queue fed by $\{T_n, \sigma_n^{(\ell)}\}$, then we have $\theta^* = \min\{\theta^\ell\}$ in the case where each sequence of service times at each station are independent of each other.

3 Queueing applications

3.1 The impact of dependence

In view of Eq. 4, Z is the supremum of a random process with negative drift and to make the connection with the existing literature, we state the following result (for a proof we refer to Lelarge 2007b Corollary 2):

Proposition 1 *Under Assumptions (IA) and (S) and if*

- (1) *The sequence $\{S_{-n}/n\}$ satisfies a large deviation principle (LDP) with a good rate function I ;*
- (2) *There exists $\epsilon > 0$ such that $\Lambda_S(\theta^* + \epsilon) < \infty$,*

where θ^ is defined as in Eq. 9. Then we have*

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\theta^* = - \inf_{\alpha > 0} \frac{I(\alpha)}{\alpha}. \tag{10}$$

This kind of result has been extensively studied in the queueing literature and follows directly from the work Duffy et al. (2003). However, we see that considering the moment generating function instead of the rate function allows us to get a more general result than Eq. 10 since we do not require the assumption on the tail (which is essential for Eq. 10 to hold see Duffy et al. 2003). Indeed this assumption ensures that the tail asymptotics of $\mathbb{P}(Z > x) \geq \mathbb{P}(S_n - (T_0 - T_n) > n\alpha)$, where we take $x = n\alpha$, cannot be dominated by $\mathbb{P}(S_n - (T_0 - T_n) > n\alpha)$ for a single n value. In this case, Eq. 10 has a nice interpretation: the natural drift of the process $S_n - (T_0 - T_n)$ is $(\gamma - a)n$, where $\gamma - a < 0$. The quantity $I(\alpha)$ can be seen as the cost for changing the drift of this process to $\alpha > 0$. Now in order to reach level x , this drift has to last for a time x/α . Hence the total cost for reaching level x with drift α is $xI(\alpha)/\alpha$ and the process naturally chooses the drift with the minimal associated cost. This can be

made precise in some cases by a conditional limit theorem that characterizes the most likely path.

In this section we show that this interpretation might be misleading in a queueing context. We consider a very simple example: a system of two queues in tandem. We assume that the sequence $\{(\sigma_n^{(1)}, \sigma_n^{(2)})\}_n$ is a sequence of i.i.d. random variables with $\max\{\mathbb{E}[\sigma_1^{(1)}], \mathbb{E}[\sigma_1^{(2)}]\} < a$ and $\mathbb{E}[\exp \theta(\sigma_1^{(1)} + \sigma_1^{(2)})]$ finite in a neighborhood of the origin.

If the service times of station 1 and station 2 are independent of each other, then the most likely cause of a given customer suffering a large delay is that a large number of its immediate predecessors require service times in excess of their inter-arrival times at one of the stations. However in the case where the service times are the same at both stations, we show that depending on the intensity of the arrival process λ , two situations may occur:

- (1) If $\lambda < \lambda_c$, then the most likely reason that a given customer suffers a large delay is that its own service time is large;
- (2) If $\lambda > \lambda_c$, then the tail asymptotic of the end-to-end delay is the same as in the independent case.

Let consider first the case where $\sigma_n^{(1)}$ and $\sigma_n^{(2)}$ are independent. We are in the framework of Corollary 1. Hence if we denote by $\theta^{(i)}$ the exponential rate of decay for the tail asymptotics of the stationary workload of a single server queue with arrival times T_n and service times $\sigma_n^{(i)}$, then we have

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\min(\theta^{(1)}, \theta^{(2)}).$$

This result has been obtained in Ganesh (1998). In words, we can say that the large deviation of the end-to-end delay in a system of two queues in tandem with independent service times is dominated by the “worst” one.

Consider now the case where $\sigma_n^{(1)} = \sigma_n^{(2)}$ for all n and the sequence $\{\sigma_n^{(1)}\}_n$ is a sequence of i.i.d. random variables exponentially distributed with mean $1/\mu$. We assume also that the arrival process is Poisson with rate $\lambda < \mu$. Then a direct application of Theorem 2 gives,

Proposition 2 *In the previous framework, we have*

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\theta^*,$$

with

$$\begin{aligned} \lambda \leq \mu/2 &\Rightarrow \theta^* = \mu/2, \\ \lambda > \mu/2 &\Rightarrow \theta^* = \mu - \lambda. \end{aligned}$$

This proposition completes the result in Ganesh (1998). For small values of λ , the tail of the end-to-end delay is determined by the total service requirement of a single customer whereas when $\lambda > \mu/2$, it is the same as in the independent case.

This shows that the behavior of tandems differs from that of a single server queue. In particular Anantharam (1989) shows that for $GI/GI/1$ queues, the build-up of large delays can happen in one of two ways. If the service times have exponential

tails, then it involves a large number of customers (whose inter-arrival and service times differ from their mean values). This behavior is analogous to that of tandems where the service times are independent at each station or if the arrival intensity is sufficiently large. If the service times do not have exponential tails, then large delays are caused by the arrival of a single customer with large service requirement. In contrast, we see that a single customer can create large delays in tandems even under the assumption of exponential service times, if the arrival intensity is sufficiently low.

3.2 A case study: queueing network with resequencing

The aim of this section is to show that the results of this paper are not restricted to the class of event graphs and that our framework can deal with complex synchronization problems encountered in applications.

In many distributed applications (e.g., remote computations, database manipulations, or data transmission over a computer network), information integrity requires that data exchanges between different nodes of a system be performed in a specific order. However, due to random delays over different paths in a system, the packets or updates may arrive at the receiver in a different order than their chronological order. In such a case, a buffer (with infinite capacity) at the receiver has to store disordered packets temporarily. There is an extensive literature on resequencing problems and we refer the interested reader to the survey Baccelli and Makowski (1989), see also the more recent work Jean-Marie and Gün (1993).

We consider a simple queueing model of disordering, namely a set of two parallel single server queueing station $|GI|1$ with renewal arrivals under probabilistic state-independent routing. This model constitutes an ersatz of the very complex situation appearing in communication networks. While the details of any protocol have been eliminated, the essence of network behavior (i.e. disordering) is preserved. In the sequel, we shall thus consider the model described in Fig. 2.

We assume that packets arrive in the first queue according to a renewal process $\{T_n\}$. We will model the routing at node 1 by a Bernoulli routing: with probability p [resp. $(1 - p)$] the packet is sent up to node 2 (resp. down to node 3). Once packet k reaches the receiver, it leaves the system if all packets j with $j < k$ have already left the system. Otherwise it stays in the resequencing buffer, where it waits for the packets with number less than k .

This model is similar to the standard fork and join system described in Example 3. Even if the routing mechanism is not the same here, we now show how we can adapt Example 3. In order to model the desired routing mechanism we will use the idea of




Fig. 2 Resequencing problem

clones, i.e., packets that behave like real packets except that they never require any service time: their service time is null. Suppose that the real route of packet k is up. Then at the end of its service in the first node, a clone is sent to node 3. Since $\sigma_k^{(3)} = 0$, the departure time of packet k from node 3 is $\mathcal{X}_k^{(3)} = \max(\mathcal{X}_k^{(1)}, \mathcal{X}_{k-1}^{(3)})$. Similarly, if the real route of packet k is down, then a clone is sent up. In both cases the “real” packet k joins the queue of node 4 once “real” packet $k - 1$ has joined it (and not before). In particular packets are ordered when they leave node 4.

Let $\{\zeta_n = (\zeta_n^{(1)}, \dots, \zeta_n^{(3)})\}_n$ be a sequence of i.i.d. mutually independent random variables with finite mean and such that $\bigoplus_i \mathbb{E} \left[\exp(\theta \zeta_1^{(i)}) \right] < \infty$ for θ in a neighborhood of the origin. Let $\{r_n\}_{n \in \mathbb{Z}}$ be a sequence of i.i.d. random variables, independent of everything else, with values in $\{2, 3\}$. We write $\mathbb{P}(r_n = 2) = 1 - \mathbb{P}(r_n = 3) =: p$, and assume that $0 < p < 1$. In order to apply our idea of clones, we consider the (max,plus)-recursion defined in Example 3 and we define

$$\sigma_n^{(1)} := \zeta_n^{(1)}, \sigma_n^{(2)} := \zeta_n^{(2)} \mathbf{1}_{\{r_n=2\}}, \sigma_n^{(3)} := \zeta_n^{(3)} \mathbf{1}_{\{r_n=3\}}, \sigma_n^{(4)} := 0.$$

We see that our system satisfies the assumptions of Theorem 2. In this case Z is the end-to-end sojourn time of a packet (taking into account the resequencing delay). The following proposition is a direct application of Theorem 2. In particular, note that $\eta \geq \theta^\ell$, hence we have

Proposition 3 *In the previous framework, we have*

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) = -\theta^* < 0,$$

where $\theta^* = \min\{\theta^\ell\}$ and the θ^ℓ 's are defined as follows

$$\theta^\ell = \sup\{\theta > 0, \Lambda_\ell(\theta) + \Lambda_T(-\theta) < 0\},$$

with

$$\begin{aligned} \Lambda_1(\theta) &= \log \mathbb{E} \left[e^{\theta \zeta_1^{(1)}} \right], \\ \Lambda_2(\theta) &= \log \left(p \mathbb{E} \left[e^{\theta \zeta_1^{(2)}} \right] + 1 - p \right), \\ \Lambda_3(\theta) &= \log \left((1 - p) \mathbb{E} \left[e^{\theta \zeta_1^{(3)}} \right] + p \right). \end{aligned}$$

Then it is possible to make some optimizations. Assume to simplify that $\zeta_n^{(1)} = 0$ for all n and that the sequence $\{\zeta_n^{(2)}\}$ and $\{\zeta_n^{(3)}\}$ are independent sequences of i.i.d. random variables exponentially distributed with respective mean $1/\mu_2$ and $1/\mu_3$. We assume also that the arrival process is Poisson with rate $\lambda > \max(\mu_2, \mu_3)$, i.e. we cannot send all packets to one node. In this case, we have to load-balance the traffic in order to have a stable system. We have to chose p such that $\lambda < \min(\mu_2/p, \mu_3/(1 - p))$. In order to do so efficiently, we would like to maximize θ^* . Here, we have $\theta_2 = \mu_2 - \lambda p$ and $\theta_3 = \mu_3 - \lambda(1 - p)$, hence the optimal value of p is given by

$$p = \frac{1}{2} \left(\frac{\mu_2 - \mu_3}{\lambda} + 1 \right) \text{ and then, } \theta^* = \frac{\mu_2 + \mu_3 - \lambda}{2}.$$

In particular, in the symmetric case $\mu_1 = \mu_2$, we find $p = 1/2$ which is in accordance with standard results of resequencing literature (Jean-Marie and Gün 1993). Also under heavy-tailed assumptions on the distribution of the $\zeta^{(i)}$'s, results are available in Lelarge (2007a) and optimal values for the routing probabilities are computed when minimizing the resequencing delay or the size of the resequencing queue size.

This model is certainly an oversimplification of the reality but it is intended to be of pedagogical interest. We should stress that it could incorporate more sophistication to enable to take into account general distributions, more than 2 paths or window control mechanism (where at any time, the number of packets in the network is upper bounded by the window size). In these cases, computations are much more complex and numerical approximations have to be made.

4 Proof of Theorem 1

We first prove the existence of the moment generating function Λ_S given by the limit 8. Then we prove that θ^* defined by Eq. 9 is positive and then we derive the tail asymptotics for Z .

Note that under our assumptions we have for $n \geq 0$,

$$S_n =_d \bigoplus_{1 \leq i \leq s} (D_{[1,n]} \otimes B_0)^{(i)}, \tag{11}$$

where the equality is in distribution. Hence for the simplicity of notation, we will “inverse time” and index the processes by non-negative indexes. Hence S_n is now given by the right-hand term of Eq. 11 and more generally we define for $u \leq v$,

$$S_{[u,v]} = \bigoplus_{1 \leq i \leq s} (D_{[u+1,v]} \otimes B_u)^{(i)}.$$

4.1 Computation of the moment generating function

Note that thanks to Assumption (ST), we have with i^* defined as the argument of the maximum in Eq. 11: $S_n = (D_{[1,n]} \otimes B_0)^{(i^*)}$, then

$$S_{n+1} \geq A_{n+1}^{(i^*, i^*)} + S_n \geq S_n.$$

In particular the process $\{S_n\}$ is non-decreasing and we have the following expression:

$$S_{[u,v]} = \bigoplus_{1 \leq i \leq s} \bigoplus_{u \leq k \leq v} (D_{[k+1,v]} \otimes B_k)^{(i)}.$$

The following lemma shows that the process has a subadditive property.

Lemma 1 *We have for $n, m \geq 0$,*

$$S_{n+m} \leq S_n + S_{[n+1, n+m]}.$$

Proof We have by definition,

$$\bigoplus_{0 \leq k \leq n} D_{[k+1,n]} \otimes B_k \leq \mathbf{0} \otimes S_n,$$

hence by monotonicity, we have

$$\begin{aligned} A_{n+1} \otimes \bigoplus_{0 \leq k \leq n} D_{[k+1,n]} \otimes B_k &\leq A_{n+1} \otimes \mathbf{0} \otimes S_n, \\ \bigoplus_{0 \leq k \leq n} D_{[k+1,n+1]} \otimes B_k &\leq (B_{n+1} \oplus \mathbf{0}) \otimes S_n, \end{aligned}$$

iterating we get:

$$\bigoplus_{0 \leq k \leq n} D_{[k+1,n+m]} \otimes B_k \leq \left(\bigoplus_{n+1 \leq j \leq n+m} D_{[j+1,n+m]} \otimes B_j \oplus \mathbf{0} \right) \otimes S_n.$$

Hence we have

$$\bigoplus_{0 \leq k \leq n+m} D_{[k+1,n+m]} \otimes B_k \leq \left(\bigoplus_{n+1 \leq j \leq n+m} D_{[j+1,n+m]} \otimes B_j \oplus \mathbf{0} \right) \otimes S_n,$$

from which the lemma follows since $S_{[n+1,n+m]} \geq 0$. □

Lemma 2 *The following limit*

$$\Lambda_S(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} [e^{\theta S_n}],$$

exists in $\mathbb{R} \cup \{+\infty\}$ for all $\theta \geq 0$. $\Lambda_S(\cdot)$ is a proper convex function which is finite on the interval $[0, \eta)$. Moreover for all n and for all $\theta < \eta$, we have $\mathbb{E}[\exp \theta S_n] < \infty$.

Proof In view of Lemma 1, we have clearly the following subadditive property, for $\theta \geq 0$,

$$\log \mathbb{E} [e^{\theta S_{m+n}}] \leq \log \mathbb{E} [e^{\theta S_n}] + \log \mathbb{E} [e^{\theta S_{m-1}}],$$

and the existence of Λ_S follows and moreover, we have

$$\Lambda_S(\theta) = \inf_{n \geq 1} \frac{1}{n} \log \mathbb{E} [e^{\theta S_n}]. \tag{12}$$

The last part of the lemma follows from

$$\bigoplus_i B_0^{(i)} \leq S_n \leq \sum_{k=0}^n \bigoplus_i B_k^{(i)},$$

thanks to Assumption **(SP)**. Then we have

$$\left(\bigoplus_i \mathbb{E} [e^{\theta B_0^{(i)}}] \right) \leq \mathbb{E} [e^{\theta S_n}] \leq \left(\bigotimes_i \mathbb{E} [e^{\theta B_0^{(i)}}] \right)^{n+1},$$

hence we see that $\Lambda_S(\theta)$ is finite for $\theta \in [0, \eta)$ and infinite for $\theta > \eta$. Then the fact that Λ_S is a proper convex function follows from Lemma 2.3.9 of Dembo and Zeitouni (1998). □

4.2 Computation of θ^*

Lemma 3 *We have $\theta^* > 0$ and*

$$\begin{aligned} \Lambda_S(\theta) + \Lambda_T(-\theta) &< 0 \text{ if } \theta \in (0, \theta^*), \\ \Lambda_S(\theta) + \Lambda_T(-\theta) &> 0 \text{ if } \theta > \theta^*. \end{aligned}$$

Proof In view of Assumption (S) and Eq. 7, we can choose n such that $\mathbb{E}[S_n] < na$. Then the function $\theta \mapsto \log \mathbb{E}[\exp(\theta S_n)]$ is convex, continuous and differentiable on $[0, \eta)$. Hence we have

$$\frac{1}{n} \log \mathbb{E}[e^{\theta S_n}] + \log \mathbb{E}[e^{-\theta \tau_1}] = \theta \left(\frac{\mathbb{E}[S_n]}{n} - a \right) + o(\theta),$$

which is less than zero for sufficiently small $\theta > 0$. By Eq. 12, we have for such θ ,

$$\Lambda_S(\theta) + \Lambda_T(-\theta) \leq \frac{1}{n} \log \mathbb{E}[e^{\theta S_n}] + \log \mathbb{E}[e^{-\theta \tau_1}] < 0.$$

Hence θ^* is positive and the last part of the lemma follows from the convexity of $\Lambda_S(\theta) + \Lambda_T(-\theta)$. □

4.3 Tail asymptotics for Z

Lemma 4 *We have,*

$$\limsup_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) \leq -\theta^*.$$

Proof First note that by Lemma 2, we have $\eta \geq \theta^*$ and for all $\theta < \eta$, we have $\mathbb{E}[\exp(\theta S_n)] < \infty$. We denote $S_n^\tau = (T_0 - T_{-n})$ for $n \geq 0$. For any $0 < \theta < \theta^*$, we have thanks to Chernoff's inequality,

$$\begin{aligned} \mathbb{P}(Z > x) &= \mathbb{P}(\cup_n \{S_n - S_n^\tau > x\}) \\ &\leq \sum_n \mathbb{P}(S_n - S_n^\tau > x) \\ &\leq e^{-\theta x} \sum_n \mathbb{E}[e^{\theta S_n}] e^{n\Lambda_T(-\theta)}, \end{aligned}$$

where the last series converge because we proved in Lemma 3 that for $\theta < \theta^*$, we have,

$$\frac{1}{n} \log \mathbb{E}[e^{\theta S_n}] + \log \mathbb{E}[e^{-\theta \tau_1}] \rightarrow \Lambda_S(\theta) + \Lambda_T(-\theta) < 0. \quad \square$$

Lemma 5 *We have,*

$$\liminf_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) \geq -\theta^*.$$

Proof We denote $\Lambda(\theta) = \Lambda_S(\theta) + \Lambda_T(-\theta)$ and $S_n^r = (T_0 - T_{-n})$ as in the previous proof. We only need to consider the case $\theta^* < \infty$. We consider first the case where there exists $\theta > \theta^*$ such that $\Lambda(\theta) < \infty$. First note that the function $\theta \mapsto \Lambda(\theta)$ is convex, hence the left-hand derivatives $\Lambda'(\theta-)$ and the right-hand derivatives $\Lambda'(\theta+)$ exist for all $\theta > 0$. Moreover, we have $\Lambda'(\theta-) \leq \Lambda'(\theta+)$ and the function $\theta \mapsto \frac{1}{2}(\Lambda'(\theta-) + \Lambda'(\theta+))$ is non-decreasing, hence $\Lambda'(\theta) = \Lambda'(\theta-) = \Lambda'(\theta+)$ except for $\theta \in \Delta$, where Δ is at most countable. Since $\Lambda(\theta) < \infty$ for $\theta > \theta^*$, we have $\Lambda(\theta^*) = 0$ and $\Lambda'(\theta^*+) > 0$. To prove this, assume that $\Lambda'(\theta^*+) = 0$. Take $\theta < \theta^*$, thanks to Lemma 3, we have $\Lambda(\theta) < 0$. Choose $\epsilon > 0$ such that $0 < \Lambda(\theta^* + \epsilon) < \epsilon|\Lambda(\theta)|$. We have

$$\frac{\Lambda(\theta^* + \epsilon)}{\epsilon} < \frac{-\Lambda(\theta)}{\theta^* - \theta},$$

which contradicts the convexity of $\Lambda(\theta)$. Hence, we can find $t \leq \theta^* + \epsilon$ such that

$$0 < \Lambda(t), \quad t \notin \Delta.$$

Note that these conditions imply $t > \theta^*$ and $\Lambda'(t) \geq \Lambda'(\theta^*+) > 0$.

Thanks to Gärtner–Ellis theorem (Theorem 2.3.6 in Dembo and Zeitouni 1998), we have

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(S_n - S_n^r > n\alpha) \geq - \inf_{x \in \mathcal{F}, x > \alpha} \Lambda^*(x), \tag{13}$$

where \mathcal{F} is the set of exposed points of Λ^* and $\Lambda^*(x) = \sup_{\theta \geq 0} (\theta x - \Lambda(\theta))$. Note that from the monotonicity of $\theta x - \Lambda(\theta)$ in x as θ is fixed, we deduce that Λ^* is non-decreasing. Moreover take $\alpha = \Lambda'(t)$, then $\Lambda^*(\alpha) = t\alpha - \Lambda(t)$ and $\alpha \in \mathcal{F}$ by Lemma 2.3.9 of Dembo and Zeitouni (1998).

Given $x > 0$, define $n = \lceil x/\alpha \rceil$. We have

$$\frac{1}{x} \log \mathbb{P}(Z > x) \geq \frac{1}{n\alpha} \log \mathbb{P}(S_n - S_n^r \geq n\alpha),$$

taking the limit in x and n (while $\alpha = \Lambda'(t)$ is fixed) gives thanks to Eq. 13,

$$\liminf_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(Z > x) \geq - \frac{t\alpha - \Lambda(t)}{\alpha} \geq -t \geq -\theta^* - \epsilon.$$

We consider now the case where for all $\theta > \theta^*$, we have $\Lambda(\theta) = \infty$, hence $\eta = \theta^*$. Take $K > 0$ and define $\tilde{S}_{[n,m]}^K = S_{[n,m]} \prod_{i=n}^m \mathbf{1}(\oplus_j B_i^{(j)} \leq K)$ and $\tilde{Z}^K = \sup_{n \geq 0} (\tilde{S}_{[-n,0]}^K - S_n^r)$. We have clearly $Z \geq \tilde{Z}^K$. It is easy to see that the proof of Lemma 2 is still valid (note that the subadditive property carries over to $\tilde{S}_{[n,m]}^K$) and the following limit exists

$$\tilde{\Lambda}_S^K(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta \tilde{S}_{[1,n]}^K} \right] = \inf_n \frac{1}{n} \log \mathbb{E} \left[e^{\theta \tilde{S}_{[1,n]}^K} \right].$$

Moreover thanks to the subadditive property of S (see Lemma 1), we have $\tilde{S}_{[1,n]}^K \leq \tilde{S}_{[1,1]}^K + \dots + \tilde{S}_{[n,n]}^K = \oplus_j B_1^{(j)} + \dots + \oplus_j B_n^{(j)}$. Hence we have $\mathbb{P}(\tilde{S}_{[1,n]}^K \leq nK) = 1$, so that $\tilde{\Lambda}_S^K(\theta) \leq \theta K$. Hence by the first part of the proof, we have

$$\liminf_{x \rightarrow \infty} \frac{1}{x} \log \mathbb{P}(\tilde{Z}^K > x) \geq -\tilde{\theta}^K,$$

with $\tilde{\theta}^K = \sup\{\theta > 0, \tilde{\Lambda}_S^K(\theta) + \Lambda_T(-\theta) < 0\}$. We now prove that $\tilde{\theta}^K \rightarrow \eta$ as K tends to infinity which will conclude the proof. Note that for any fixed $\theta \geq 0$, the function $\tilde{\Lambda}_S^K(\theta)$ is nondecreasing in K and $\lim_{K \rightarrow \infty} \tilde{\Lambda}_S^K(\theta) = \tilde{\Lambda}_S(\theta) \leq \Lambda_S(\theta)$. This directly implies that $\tilde{\theta}^K \geq \eta$. Take $\theta > \eta$, so that $\Lambda_S(\theta) = \infty$. If $\tilde{\Lambda}_S(\theta) < \infty$, then for all K , we have $\tilde{\Lambda}_S^K(\theta) \leq \tilde{\Lambda}_S(\theta) < \infty$. But, we have $\tilde{\Lambda}_S^K(\theta) = \inf_n \frac{1}{n} \log \mathbb{E} \left[e^{\theta \tilde{S}_{[1,n]}^K} \right]$, so that there exists n such that

$$\mathbb{E} \left[e^{\theta S_{[1,n]}} \right], \max(\oplus_j B_1^{(j)}, \dots, \oplus_j B_n^{(j)}) \leq K \leq e^{\tilde{\Lambda}_S^K(\theta)+1} \leq e^{\tilde{\Lambda}_S(\theta)+1},$$

but the left-hand side tends to infinity as $K \rightarrow \infty$. Hence we proved that for all $\theta > \eta$, we have $\tilde{\Lambda}_S^K(\theta) \rightarrow \infty$ as $K \rightarrow \infty$. This implies that $\tilde{\theta}^K \rightarrow \eta$ as $K \rightarrow \infty$. \square

5 Proof of Theorem 2 and Corollary 1

We begin with a general result showing the existence of the function Λ_ℓ . Let $\{M_n\}$ be an i.i.d. sequence of irreducible aperiodic (max,plus)-matrices with fixed structure. We denote

$$M_{[1,n]}^{(i,j)} = (M_n \otimes \dots \otimes M_1)^{(i,j)}.$$

Lemma 6 For $\theta \geq 0$, the following limit exists in $\mathbb{R} \cup \{+\infty\}$ and is independent of i and j ,

$$\Lambda_M(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta M_{[1,n]}^{(i,j)}} \right].$$

Proof We denote

$$\Lambda_M^{(i,j)}(\theta, n) = \log \mathbb{E} \left[e^{\theta M_{[1,n]}^{(i,j)}} \right].$$

We have for $\theta \geq 0$,

$$\begin{aligned} \Lambda_M^{(i,j)}(\theta, n+m) &= \log \mathbb{E} \left[e^{\theta M_{[1,n+m]}^{(i,j)}} \right] \\ &= \log \mathbb{E} \left[\max_k e^{\theta M_{[n+1,n+m]}^{(i,k)}} e^{\theta M_{[1,n]}^{(k,j)}} \right] \\ &\geq \max_k \left\{ \log \mathbb{E} \left[e^{\theta M_{[n+1,n+m]}^{(i,k)}} \right] + \log \mathbb{E} \left[e^{\theta M_{[1,n]}^{(k,j)}} \right] \right\} \\ &= \max_k \left\{ \Lambda_M^{(i,k)}(\theta, m) + \Lambda_M^{(k,j)}(\theta, n) \right\}. \end{aligned}$$

In particular for $j = i$, we have

$$\Lambda_M^{(i,i)}(\theta, n + m) \geq \Lambda_M^{(i,i)}(\theta, m) + \Lambda_M^{(i,i)}(\theta, n).$$

Moreover thanks to the fixed structure assumption and the aperiodicity, there exists N such that for $n \geq N$, we have $M_{[1,n]}^{(i,j)} > -\infty$ for all i and j , hence $\Lambda_M^{(i,j)}(\theta, n) > -\infty$ and we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \Lambda_M^{(i,i)}(\theta, n) = \sup_{n \geq N} \frac{1}{n} \Lambda_M^{(i,i)}(\theta, n) > -\infty.$$

For arbitrary i and j , choose $n, m \geq N$ and note that

$$\begin{aligned} \Lambda_M^{(i,j)}(\theta, n + m) &\geq \Lambda_M^{(i,i)}(\theta, n) + \Lambda_M^{(i,j)}(\theta, m), \\ \Lambda_M^{(i,i)}(\theta, n + m) &\geq \Lambda_M^{(i,j)}(\theta, n) + \Lambda_M^{(j,i)}(\theta, m), \end{aligned}$$

where all terms are in $\mathbb{R} \cup \{+\infty\}$. Letting $n \rightarrow \infty$ while keeping m fixed, it follows that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \Lambda_M^{(i,j)}(\theta, n) = \lim_{n \rightarrow \infty} \frac{1}{n} \Lambda_M^{(i,i)}(\theta, n). \quad \square$$

Note that Assumption **(ST)** ensures that the matrices $\{A_n(\ell, \ell)\}$ are irreducible aperiodic with fixed structure. We now extend previous lemma to the sequence $\{A_n\}$ of reducible matrices. To do so, we first associate a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to A_n , as in Section 2.3 of Baccelli et al. (1992). Set $\mathcal{V} := \{1, \dots, s\}$, which we abbreviate as $[1, s]$. An edge (i, j) belongs to \mathcal{E} if and only if $A_n^{(j,i)} \geq 0$. Two nodes of \mathcal{V} are said to belong to the same communication class if there is a directed path from the first to the second and another one from the second to the first. Let $\mathcal{C}_1, \dots, \mathcal{C}_d$ be the communication classes of \mathcal{G} and \prec the associated partial order, namely $\mathcal{C}_\ell \prec \mathcal{C}_m$ if there is a path from any vertex in \mathcal{C}_ℓ to any vertex in \mathcal{C}_m . Without loss of generality, we assume that $\mathcal{C}_\ell \prec \mathcal{C}_m$ implies $\ell \leq m$; this is a notationally convenient restriction on the numbering of the communication classes.

We use the following notation:

- For any coordinate $i \in \mathcal{V}$, its communication class is denoted by $[i]$,
- For any coordinate i , the subset of coordinates j such that $[j] \prec [i]$ is denoted by $[\leq i]$;
- For any coordinate i , the subset of coordinates j such that $[i] \prec [j]$ is denoted by $[i \leq]$;
- For any coordinate i and $j \in [i \leq]$, we write

$$[i \leq j] := [i \leq] \cap [\leq j].$$

We now extend previous lemma. We introduce first some notation.

$$\Lambda_\ell(\theta) = \Lambda_{[i]}(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta D_{[1,n]}^{(i,i)}} \right],$$

which does not depend on $i \in \mathcal{C}_\ell$ as shown above.

Lemma 7 For $\theta \in [0, \eta)$, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} \left[e^{\theta D_{[1,n]}^{(i,j)}} \right] = \sup_{k \in [i \leq j]} \Lambda_{[k]}(\theta),$$

where the supremum over the empty set is $-\infty$.

Proof If $[i \leq j] = \emptyset$, the result is obvious and if $j \in [i]$, the result follows from previous lemma. Hence we consider only the case: $[i \leq j] \supset [i]$. We denote

$$\Lambda^{(i,j)}(\theta, n) = \log \mathbb{E} \left[e^{\theta D_{[1,n]}^{(i,j)}} \right].$$

With the same argument as in previous lemma, we have for $\theta \geq 0$, and for any $[i] \ll \mathcal{C}_\ell \ll [j]$, there exists $k \in \mathcal{C}_\ell$ and $u, v \geq 1$,

$$\Lambda^{(i,j)}(\theta, n + u + v) \geq \Lambda^{(i,k)}(\theta, u) + \Lambda^{(k,k)}(\theta, n) + \Lambda^{(k,j)}(\theta, v),$$

where each term is finite since $\theta < \eta$. Hence we have by previous lemma:

$$\liminf_{n \rightarrow \infty} \frac{\Lambda^{(i,j)}(\theta, n)}{n} \geq \sup_{k \in [i \leq j]} \Lambda_{[k]}(\theta).$$

For the upper bound, note that there exists $u \geq 1$ such that $D_{[1,u]}^{(i,j)} > -\infty$ for all $i \leq j$. Consider first the case $d = 2, i \in \mathcal{C}_1$ and $j \in \mathcal{C}_2$, then we have

$$\begin{aligned} \Lambda^{(i,j)}(\theta, n) &\leq \log \mathbb{E} \left[\max_{a \in \mathcal{C}_2} e^{\theta D_{[1,u]}^{(i,a)}} e^{\theta D_{[u+1,n]}^{(a,j)}} + \max_{b \in \mathcal{C}_1} e^{\theta D_{[1,n-u]}^{(i,b)}} e^{\theta D_{[n-u+1,n]}^{(b,j)}} \right], \\ &\leq \log \left(\sum_{a \in \mathcal{C}_2} e^{\Lambda^{(i,a)}(\theta, u)} e^{\Lambda^{(a,j)}(\theta, n)} + \sum_{b \in \mathcal{C}_1} e^{\Lambda^{(i,b)}(\theta, n)} e^{\Lambda^{(b,j)}(\theta, u)} \right), \end{aligned}$$

Hence by Lemma 1.2.15 of Dembo and Zeitouni (1998), we have

$$\limsup_{n \rightarrow \infty} \frac{\Lambda^{(i,j)}(\theta, n)}{n} \leq \max(\Lambda_1(\theta), \Lambda_2(\theta)).$$

We have clearly by induction that

$$\limsup_{n \rightarrow \infty} \frac{\Lambda^{(i,j)}(\theta, n)}{n} \leq \sup_{k \in [i \leq j]} \Lambda_{[k]}(\theta),$$

which concludes the proof. □

We now compute $\Lambda_S(\theta)$ for a (max,plus)-linear system under the assumptions of Theorem 2.

Lemma 8 We have for $\theta \in [0, \eta)$

$$\Lambda_S(\theta) = \sup_{\ell} \Lambda_\ell(\theta).$$

Proof The lower bound follows directly from the following inequality: for all ℓ , we have for $i \in \mathcal{C}_\ell$,

$$\mathbb{E} \left[e^{\theta S_n} \right] \geq \mathbb{E} \left[e^{\theta D_{[1,n]}^{(i,i)}} \right].$$

We now derive the upper bound. Note that $A_k \otimes \mathbf{0} = B_k \oplus \mathbf{0} \geq B_k$, hence we have

$$\begin{aligned} S_n &= \bigoplus_{1 \leq i \leq s} \bigoplus_{0 \leq k \leq n} (D_{[k+1,n]} \otimes B_k)^{(i)} \\ &\leq \bigoplus_{1 \leq i, j \leq s} \bigoplus_{0 \leq k \leq n} D_{[k,n]}^{(i,j)}. \end{aligned}$$

Hence we have

$$\begin{aligned} \mathbb{E} \left[e^{\theta S_n} \right] &\leq \sum_{i,j} \sum_k \mathbb{E} \left[e^{\theta D_{[k,n]}^{(i,j)}} \right] \\ &\leq n \sum_{i,j} \mathbb{E} \left[e^{\theta D_{[0,n]}^{(i,j)}} \right], \end{aligned}$$

and the lemma follows directly from Lemma 1.2.15 of Dembo and Zeitouni (1998). □

Theorem 2 follows directly from the fact that $\Lambda_S(\theta) = \infty$ as soon as $\theta > \eta$ which follows from the lower bound $S_n \geq B_0^{(i)}$.

We now prove Corollary 1. The following lemma implies that $\Lambda_S(\theta) = \sup_\ell \Lambda_\ell(\theta)$ for all $\theta \geq 0$ and then Corollary 1 follows:

Lemma 9 *Under assumptions of Corollary 1, there exists $\ell \in [1, d]$ such that $\Lambda_\ell(\theta) = \infty$ for all $\theta > \eta$.*

Proof We only need to consider the case $\eta < \infty$. Take i such that $\sup\{\theta, \mathbb{E}[\exp \theta B_0^{(i)}] < \infty\} = \eta$. By the condition on the entries of A_n , we have

$$B_n^{(i)} \leq \bigoplus_j A_n^{(i,j)} \leq \sum_{k=1}^K \sigma_n^{(k)}.$$

Hence there exists k such that $\mathbb{E}[\exp(\theta \sigma_n^{(k)})] = \infty$ for $\theta > \eta$ and there exist j such that $A_n^{(i,j)} = \sigma_n^{(k)}$. Then ℓ defined by $\mathcal{C}_\ell = [j]$ satisfies the property claimed in the lemma. □

6 Conclusion

We have shown that the distribution of the stationary solution of a (max,plus) recursion has an exponentially decaying tail and we gave an analytical formula to compute the decay rate. We also analyzed a queueing network with multipath routing and showed on a simple example how our analysis could help in the design of the routing decision depending on the characteristic of the traffic.

We are currently working on some possible extensions of our work. Of particular interest would be a large deviations principle for the process $\{S_n/n\}$ introduced in Eq. 5. It would allow to give the most probable way for a large deviation of the maximal dater. Also, it should be possible to use the distributional Little's law (Haji and Newell 1971) to get asymptotics for the number of packets in the networks.

In general, the characterization of the decay rate is given by the moment generating function which is not easy to compute, especially in the case of feedback. One practical question of interest would be to find good ways to estimate this function from the statistics made on the traffic.

Acknowledgements The author would like to thank Peter Friz for pointing out a mistake in an earlier version of this work and the participants of Valuetools 2006 (where this work was presented) and especially Bruno Gaujal for a comment related to Remark 1.

References

- Anantharam V (1989) How large delays build up in a $GI/G/1$ queue. *Queueing Syst Theory Appl* 5(4):345–367
- Baccelli F, Cohen G, Olsder GJ, Quadrat J-P (1992) Synchronization and linearity. Wiley. Available at <http://www-rocq.inria.fr/metalau/cohen/SED/book-online.html>
- Baccelli F, Foss S (1995) On the saturation rule for the stability of queues. *J Appl Probab* 32:494–507
- Baccelli F, Foss S, Lelarge M (2004) Asymptotics of subexponential max plus networks: the stochastic event graph case. *Queueing Syst* 46:75–96
- Baccelli F, Makowski, A (1989) Queueing models for systems with synchronization constraints. *Proc IEEE* 77(1):138–161
- Bertsimas D, Paschalidis IC, Tsitsiklis JN (1998) On the large deviations behavior of acyclic networks of $G/G/1$ queues. *Ann Appl Probab* 8(4):1027–1069
- Blondel VD, Gaubert S, Tsitsiklis JN (2000) Approximating the spectral radius of sets of matrices in the max-algebra is NP-hard. *IEEE Trans Automat Contr* 45(9):1762–1765
- Chang CS (1995) Sample path large deviations and intree networks. *Queueing Syst Theory Appl* 20:7–36
- Cramér H (1938) Sur un nouveau théorème limite de la théorie des probabilités. *Actual Sci Ind* 736:5–23
- Dembo A, Zeitouni O (1998) Large deviations techniques and applications. Springer-Verlag
- Duffy K, Lewis JT, Sullivan WG (2003) Logarithmic asymptotics for the supremum of a stochastic process. *Ann Appl Probab* 13(2):430–445
- Dupuis P, Ellis RS (1995) The large deviation principle for a general class of queueing systems. I. *Trans Am Math Soc* 347(8):2689–2751
- Dupuis P, Ellis RS, Weiss A (1991) Large deviations for Markov processes with discontinuous statistics. I. General upper bounds. *Ann Probab* 19(3):1280–1297
- Ganesh A (1998) Large deviations of the sojourn time for queues in series. *Ann Oper Res* 79:3–26
- Ganesh A, Anantharam V (1996) Stationary tail probabilities in exponential server tandems with renewal arrivals. *Queueing Syst Theory Appl* 22(3-4):203–247
- Haji R, Newell GF (1971) A relation between stationary queue and waiting time distributions. *J Appl Probab* 8:617–620
- Iglehart DL (1972) Extreme values in the $GI/G/1$ queue. *Ann Math Stat* 43:627–635
- Ignatiouk-Robert I (2000) Large deviations of Jackson networks. *Ann Appl Prob* 3:962–1001
- Ignatiouk-Robert I (2005) Large deviations for processes with discontinuous statistics. *Ann Probab* 33(4):1479–1508
- Jean-Marie A, Gün L (1993) Parallel queues with resequencing. *J Assoc Comput Mach* 40(5):1188–1208
- Lelarge M (2005) Rare events in networks. PhD thesis, Ecole polytechnique Available at <http://www.di.ens.fr/~lelarge>
- Lelarge M (2007a) Packet reordering in networks with heavy-tailed delays. *Math Methods Oper Res.* doi:10.1007/s00186-007-0175-8

- Lelarge M (2007b) Tail asymptotics for monotone-separable networks. *J Appl Probab* 44(2):306–320.
- Pakes A (1975) On the tails of waiting-time distributions. *J Appl Probab* 12:555–564
- Toomey F (2002) Large deviations of products of random topical operators. *Ann Appl Probab* 12:317–333



Marc Lelarge graduated from Ecole Polytechnique (Palaiseau, France) in 1999 and qualified as an engineer at Ecole Nationale Supérieure des Télécommunications (Paris, France) in 2001. He received his PhD degree in Applied Mathematics from Ecole Polytechnique in 2005 and is research scientist at INRIA since 2006.