THIELE CENTRE
FOR APPLIED MATHEMATICS IN NATURAL SCIENCE

On Preemptive-Repeat LIFO Queues

Søren Asmussen and Peter W. Glynn

Research Report No. 02 | January 2017

On Preemptive-Repeat LIFO Queues

Søren Asmussen[1] and Peter W. Glynn[2]

[1]Department of Mathematics , Aarhus University, asmus@math.au.dk
[2]Huang Engineering Center 359A, Stanford University, glynn@stanford.edu

Abstract

In this paper, we study the basic properties of last-in first-out (LIFO) preemptive repeat single-server queues in which the server needs to start service from scratch whenever a preempted customer reaches the server. In particular, we study the question of when such queues are stable (in the sense that the equilibrium time-in-system is finite-valued with probability one), and show how moments of the equilibrium customer sojourn time can be computed when the system is stable. A complete analysis of stability is provided in the setting of Poisson arrivals and in that of the Markovian arrival process. The stability region depends upon the detailed structure of the interarrival- and service time distributions, and can not be expressed purely in terms of expected values. This is connected to the fact that such preemptive repeat queues are not work-conserving.

Keywords: Branching process, Iterative scheme, Markovian arrival process, Random walk, Sojourn time, Stability, Stochastic fixpoint equation

1 Introduction

In some systems in which queueing occurs, it can be attractive to always focus the service resource on the job that has waited least. This leads to a LIFO (last-in first-out) queue discipline. If LIFO is implemented preemptively, traditional analyses assume that when the server resumes work on the interrupted job, the server resumes the processing of the job from the point at which it was interrupted. This is a work-conserving queue discipline that is known as the "preemptive resume" discipline, and which has been extensively discussed in the literature; see for example [16], [19], [28], [13] for textbook treatments.

This paper is concerned with LIFO preemptive queues in which an interrupted job has to be restarted and processed from scratch, so that all the previous service effort expended on that job is lost. One modeling possibility is that every time the job is restarted, the remaining service is *identical* to its initial service requirement. We view such a "preemptive repeat identical", denoted PRI, model as a vehicle for studying job stacks in which jobs accumulate a complex state as they run. When applied in a computing environment, the state may consist of files that the job has opened, partial computations stored in temporary registers, as well as the program counter keeping track of the job's progress thus far in executing the code. When a job is preempted, it cannot be resumed where it left off unless its state is somehow recreated. When the state is complex, it often is easiest to just restart the job from scratch.

We also study LIFO "preemptive repeat different", denoted PRD, queues, in which each time the remaining service time for a restarted job is generated, it is generated independently from a common distribution. That is, the residual service times are *different* at each restart. We view this model as describing systems in which the service time requirements are deterministic and constant. The actual time required to finish the task is then determined by the service capacity available to that job at the time that service is rendered. If the server shares its effort across both the queue and other non-queueing tasks, the service capacity available to the queue may then vary randomly over time. In this case, it may be reasonable to model the time required to complete the interrupted job from the successive instants of restart as a sequence of independent and identically distributed (i.i.d.) random variables (r.v.'s).

Both preemptive resume and preemptive repeat models have been studied in the setting of priority queues; see Gaver (1962) and Avi-Itzhak (1963a,b) for early work on such non-work conserving queueing models. As noted earlier, our contribution here is to extend the theory for preemptive repeat models to the LIFO setting. Because every customer can be interrupted by future arriving customers, this leads to branching process structure that is not present in the context of priority queues.

As might be expected, the stability region for such queues depends on the detailed structure of the underlying arrival and service time distributions, and not exclusively through their associated rates; see Theorem 2.3. In Section 2, we develop necessary and sufficient conditions for stability of LIFO PRI queues, and compute the moments of the equilibrium sojourn time D. As in most of the literature,we assume Poisson arrivals, but much of the discussion is formulated in terms of the solution of a more general stochastic fixed point equation. As expected, one finds that the stability region for LIFO PRI systems is smaller than for LIFO preemptive resume, and $\mathbb{E}D$ larger. Section 3 extends the analysis to Markovian arrival processes, [25], [3, pp. 302–6] (examples are the Markov-modulated Poisson process and phase-type renewal processes, and Markovian arrival processes are dense in the space of point processes). The recursive structure here leads to an iterative computational scheme and numerical examples are presented. A discussion of LIFO PRD queues is given in Section 4. Section 5 concludes the paper with a number of remarks not directly in the mainstream of the paper.

Our notational convention PRI, PRD follows, e.g., [21] and [29], but to avoid confusion, note that the R stands for Repeat, not Resume! We also let T denote a generic interarrival time (in the setting of renewal arrivals) and F its distribution, with S denoting a generic service time. We assume that customer service requirements are always generated i.i.d. from the distribution of S. A peculiarity of the setting is that the M/G/1 case of Poisson arrivals admits a more natural interpretation but that the GI/D/1 case of deterministic service times $S \equiv \ell$ is often easier to analyze. On the methodological side, more than one approach is possible and we have sometimes presented several solution techniques. A basic tool is a stochastic fixed point equation presented in (2.1) that may be exploited in various ways. One is to relate the equation to to Galton-Watson family trees (the tree structure has in fact a prominent role in the matrix-analytic literature, see Remark 3.5 below for references). This structure comes in two versions, binary splitting (e.g. the proof of Proposition 2.2) and one where the children of a customer is the totality of customers that interrupted his service (e.g. the proof of Theorem 2.3). But there are also sometimes embedded random walks (RWs) such that stability of the systems under study is equivalent to recurrence of the RW.

2 The LIFO Preemptive-Repeat-Identical Queueing Model

We start by introducing the basic stochastic fixed point equation (SFPE). Let $D(s)$ be the sojourn time (or time-in-system) of a customer, conditional on its service requirement $S = s$. Then the family $\big(D(s) : 0 < s < \infty\big)$ satisfies the SFPE

$$D(s) \overset{\mathcal{D}}{=} T \wedge s + \mathbb{I}(T \le s)\big(D + D^*(s)\big) \tag{2.1}$$

where $D, D^*(s), T$ are independent, $D^*(s) \overset{\mathcal{D}}{=} D(s)$. and $D \overset{\mathcal{D}}{=} D(S)$ where S is independent of $(D(s) : s \ge 0)$. In the queueing context, one can think of D as the time-in-system of the preempting customer and of $D^*(s)$ as the remaining time-in-system of the customer himself when reentering service.

Remark 2.1. The SFPE (2.1) makes mathematical sense even for non-exponential interarrival time r.v.'s T. We will study it in its own right in this GI/G/1 setting, but it does not then admit a natural queueing interpretation. In particular, in the queueing context, it must be that when the customer resumes service, the residual time until the next arrival is the forward recurrence time in the associated renewal process and hence is not distributed as T in the non-Poisson setting. This feature is not represented in the SFPE (2.1). A similar complication, for the customer preempting the initial one (corresponding to the D in (2.1)), is also not represented in the SFPE. Partial results for the queueing model that extend the arrival process beyond the Poisson setting are given at the end of this section; these results do not directly apply the SFPE.

It follows immediately from (2.1) and $D = D(S)$ that either $\mathbb{P}\big(D(s) < \infty\big) < 1$ for all s or $\mathbb{P}\big(D(s) < \infty\big) = 1$ for all s. The following result shows that $D(s)$ is the minimal solution that is meaningful in the queueing context:

Proposition 2.2. *Assume that* $\mathbb{P}\big(D(s) < \infty\big) = 1$ *for all* s. *Let, for a fixed* s, $\widetilde{D}(s)$ *be any other non-negative solution of* (2.1). *Then, for all fixed* s $\widetilde{D}(s) \ge D(s)$ *in the sense of stochastic order.*

Proof. We provide an iterative argument. Set

$$D_1(s) = T \wedge s \tag{2.2}$$

and recursively define the distribution of $D_n(\,\cdot\,)$ via

$$D_{n+1}(s) = T \wedge s + I(T \le s)(D_n^*(S) + D_n(s)), \tag{2.3}$$

for $n \ge 1$, where $T, S, D_n^*(\,\cdot\,), D_n(\,\cdot\,)$ are independent of one another and $D_n^*(\,\cdot\,) \overset{\mathcal{D}}{=} D_n(\cdot)$. The recursion (2.3) can be interpreted in terms of a binary splitting branching process where the primary customer with service time s is the ancestor and, if preempted, is replaced by two children having service times S and s, and respectively. It follows then by induction that $D_{n+1}(s)$ is the total time in system for the primary customer when the depth of the family trees of the two children is at most n, that is, the depth Δ of the whole family tree is at most $n+1$. The assumption $\mathbb{P}\bigl(D(s) < \infty\bigr) = 1$ is equivalent to $\Delta < \infty$ a.s., and hence by monotone convergence $D_n(s) \xrightarrow{\mathcal{D}} D(s)$ as $n \to \infty$.

Now let $\widetilde{D}(s)$ be some other non-negative solution of (2.1). It is then evident that from the SFPE that $D_n(\,\cdot\,) \le \widetilde{D}(\,\cdot\,)$ for $n \ge 1$ in the sense of stochastic dominance. Let $n \to \infty$. □

We next identify the necessary and sufficient conditions guaranteeing that the queue empties infinitely often.

Theorem 2.3. *Suppose that the arrival process is Poisson with rate λ. Such a* LIFO PRI M/G/1 *system empties infinitely often if and only if*

$$\mathbb{E}\mathrm{e}^{\lambda S} \le 2\,. \tag{2.4}$$

More generally, if T has a general distribution F, then $D < \infty$ a.s. if and only if

$$\theta = \mathbb{E}\Bigl[\frac{1}{\overline{F}(S)}\Bigr] \le 2\,. \tag{2.5}$$

Remark 2.4. For the LIFO M/M/1 queue, (2.4) translates into

$$2 \ge \mathbb{E}\mathrm{e}^{\lambda S} = \int_0^\infty \mathrm{e}^{\lambda s}\,\mu \mathrm{e}^{-\mu s}\,\mathrm{d}s = \frac{\mu}{\mu - \lambda} = \frac{1}{1-\rho}$$

which is the same as $\rho \le 1/2$. Note that the quantity $\mathbb{E}\mathrm{e}^{\lambda S}$ also comes up in the setting of M/G/1 priority queues, cf. [16].

Remark 2.5. In general, one can view the finiteness of $\mathbb{E}\overline{F}(S)^{-1}$ as a condition asserting that the tail of S must be sufficiently lighter than the tail of T. In a similar vein, $\mathbb{E}\overline{F}(S)^{-1} \le 2$ means that the main part of the mass of S is located sufficiently close to the origin. Note also that if F has a power-like tail, $\overline{F}(t) \sim ct^{-\alpha}$ as $t \to \infty$, $\mathbb{E}\overline{F}(S)^{-1} < \infty$ implies that $\mathbb{E}S^\alpha < \infty$. This condition is much weaker than existence of exponential moments, and it can even hold in settings in which the r.v.'s have infinite mean. To further emphasize the interplay between the tails of S and T, note that if F, F_1 are Pareto with tails $(1+x)^{-\alpha}$, respectively $(1+x)^{-\alpha_1}$, with $\alpha_1 \le \alpha$, then there are more service time distributions for which D is finite-valued under F than for F_1. (Note that if $\mathbb{E}S^\alpha \le 2$, then by Jensen's inequality $\mathbb{E}S^{\alpha_1} \le \bigl[\mathbb{E}S^\alpha\bigr]^{\alpha_1/\alpha} \le 2$). Intuitively, an interarrival time distribution F that generates longer interarrival times generates more favorable intervals within which jobs can terminate.

Remark 2.6. An interesting and important element of Theorem 2.3 is that the time-in-system D is finite-valued when θ is exactly 2, so that D is finite-valued on the boundary of the stability region, and not just its interior. This implies that, unlike conventional heavy-traffic analysis, there is no heavy-traffic limit theorem in this setting for the distribution of D, in which some normalized aD converges weakly to a proper non-zero limit rv as θ tends to 2, for some suitably chosen factor a tending simultaneously to 0. This phenomenon also arises in the setting of preemptive-resume queues. Note that in the preemptive-resume context, it is known, however, that such heavy-traffic (distributional) limit theorems do hold for the number-in-system process, see Limic [23]. It may well be that a similar behavior occurs here, in that a suitably normalized version of the number-in-system process for PRI queues converges weakly to a non-degenerate limit when θ tends to 2.

Proof of Theorem 2.3. A busy period has finite duration when the number of customers served in that busy period is finite-valued. We now can think of customers within such a busy period as individuals in a Galton-Watson process. The root of the associated tree corresponds to the customer that initiates the busy period, and the children of a customer are then the ones that interrupted his service. Because of the recursive structure of the busy period, the parent customer has probability 0 of infinitely many descendants if and only if each child has that property. When each child's "tree" is finite, the total number of individuals served in the busy period is exactly the total number of progeny of the parent at the root of the Galton-Watson tree. Furthermore, the Galton-Watson tree has infinitely many progeny if and only if infinitely many customers arrive during the the busy period. So, stability corresponds precisely to the Galton-Watson tree being finite.

Turning now to the finiteness of the tree, observe that when a customer has service time $S = s$, the number N of children is geometric with parameter $\mathbb{P}(T \le s) = F(s)$, where $F(s) = 1 - \mathrm{e}^{-\lambda s}$. So, the offspring mean is

$$\mathbb{E}N = \mathbb{E}\Big[\frac{F(S)}{1 - F(S)}\Big] = \mathbb{E}\Big[\frac{1}{\overline{F}(S)}\Big] - 1\,. \tag{2.6}$$

But the customer will terminate service when all his followers in the branching tree have done so, i.e. when the branching tree is finite. The (well known) necessary and sufficient condition for this is $\mathbb{E}N \le 1$, cf. [14], which is the same as (2.5). Relation (2.4) is just the special case $\overline{F}(s) = \mathrm{e}^{-\lambda s}$. □

Corollary 2.7. *Consider the unstable case $\mathbb{E}\overline{F}(S)^{-1} > 2$ where $q = \mathbb{P}(D < \infty) < 1$. Then q is the minimal root in $[0, 1]$ of the fixpoint problem*

$$q = \mathbb{E}\Big[\frac{\overline{F}(S)}{1 - qF(S)}\Big]\,. \tag{2.7}$$

Proof. Define $q(s) = \mathbb{P}\big(D(s) < \infty\big)$. Then the SFPE (2.1) implies

$$q(s) = \overline{F}(s) + F(s)q \cdot q(s) \quad \text{i.e. } q(s) = \frac{\overline{F}(s)}{1 - qF(s)}\,.$$

Integrating s w.r.t. $G(\mathrm{d}s)$ gives (2.7). Letting $h(q)$ be the r.h.s. of (2.7), $h(q)$ is convex with $h(0) = \mathbb{E}\overline{F}(S)$, $h(1) = 1$ and $h'(1) = \mathbb{E}\big[F(S)/\overline{F}(S)\big]$ which is > 1 in the unstable case. From this the result follows by graphical inspection. □

Given that stability (in the sense of the system emptying infinitely often) has been settled, we turn next to studying the expected value of the key performance measure for this model, namely the sojourn time of a typical customer. Let D_k be the sojourn time for the k'th customer to enter the system. Note that the fact that customer k's sojourn time depends only on inter-arrival and service times subsequent to his arrival implies that the sequence $(D_k : k \geq 1)$ is a stationary sequence. If the system is stable, then we also know that the sequence is regenerative. Hence, we may conclude that the customer-average of the first n $f(D_k)$'s converges almost surely (a.s) to $\mathbb{E}f(D)$ for any non-negative function f, where D is a r.v. having the common distribution of the D_k's.

We turn to the computation of $m = \mathbb{E}D$.

Theorem 2.8. *In the* LIFO PRI *queue with Poisson arrivals having rate* λ*,*

$$m = \mathbb{E}D = \frac{\mathbb{E}\big[(T \wedge S)\mathrm{e}^{\lambda S}\big]}{2 - \mathbb{E}\mathrm{e}^{\lambda S}} = \frac{1}{\lambda}\Big[\frac{1}{2 - \mathbb{E}\mathrm{e}^{\lambda S}} - 1\Big]\,, \tag{2.8}$$

provided that $\mathbb{E}\mathrm{e}^{\lambda S} < 2$*. More generally, if* T *has a general distribution* F *and* (2.5) *holds, then*

$$m = \mathbb{E}D = \frac{\mathbb{E}\big[(T \wedge S)/\overline{F}(S)\big]}{2 - \mathbb{E}\big[1/\overline{F}(S)\big]} \tag{2.9}$$

Proof. We have $m = \mathbb{E}m(S)$, where

$$m(s) = \mathbb{E}[D \,|\, S = s] = \mathbb{E}[T \wedge s] + F(s)\big(m + m(s)\big)\,,$$

Taking expectations in (2.1), an argument similar to that above establishes that $m(s)$ is the smallest non-negative solution of this equation, so that

$$m(s) = \frac{\mathbb{E}[T \wedge s] + F(s)m}{\overline{F}(s)}\,.$$

Integrating both sides of this equation with respect to $G(\mathrm{d}s)$ yields the equality

$$m = \mathbb{E}\Big[\frac{T \wedge S}{\overline{F}(S)}\Big] + m\mathbb{E}\Big[\frac{F(S)}{\overline{F}(S)}\Big],$$

From this (2.9) and the first identity in (2.8) immediately follows. Now just note that in the Poisson case

$$\mathbb{E}\big[T\mathrm{e}^{\lambda S}\mathbb{I}(T > s) \,|\, S = s\big] = \mathrm{e}^{\lambda s}\mathbb{E}[T; T > s] = \mathrm{e}^{\lambda s}\mathrm{e}^{-\lambda s}\Big[s + \frac{1}{\lambda}\Big] = s + \frac{1}{\lambda}$$

and so

$$\begin{aligned}\mathbb{E}\big[(T \wedge S)\mathrm{e}^{\lambda S}\big] &= \mathbb{E}\big[T\mathrm{e}^{\lambda S}; T < S\big] + \mathbb{E}\big[S\mathrm{e}^{\lambda S}; T > S\big] \\ &= \mathbb{E}T \cdot \mathbb{E}\mathrm{e}^{\lambda S} - \mathbb{E}\big[T\mathrm{e}^{\lambda S}; T > S\big] + \mathbb{E}\big[S\mathrm{e}^{\lambda S}\mathrm{e}^{-\lambda S}\big] \\ &= \frac{1}{\lambda}\mathbb{E}\mathrm{e}^{\lambda S} - \mathbb{E}S - \frac{1}{\lambda} + \mathbb{E}S = \frac{1}{\lambda}[\mathbb{E}\mathrm{e}^{\lambda S} - 1]\,.\end{aligned}$$

□

The theory of branching processes provides an alternative argument:

Proof. Let X_n be the size of the n'th generation in the Galton-Watson process. and $P = 1+X_1+X_2+\cdots$ the total progeny. Then $\mathbb{E}Z_n = (\mathbb{E}N)^n$ and so $\mathbb{E}P = 1/(1-\mathbb{E}N)$. Let further Z be the time spent in service by an individual and let $Z(s)$ have the distribution of Z given $S = s$. The Z of an individual is his service time and the sum of the interrupted service attempts, the number of which is geometric$\bigl(F(s)\bigr)$ given $S = s$. This yields

$$\begin{aligned}\mathbb{E}Z(s) &= s + \frac{F(s)}{\overline{F}(s)}\mathbb{E}[T\,|\,T<s] = \frac{s\overline{F}(s)+\mathbb{E}[T;\,T<s]}{\overline{F}(s)}\\ &= \mathbb{E}\Bigl[\frac{s\mathbb{I}(T>s)+T\mathbb{I}(T\le s)}{\overline{F}(s)}\Bigr] = \mathbb{E}\Bigl[\frac{T\wedge s}{\overline{F}(s)}\Bigr]\end{aligned}$$

and so $\mathbb{E}Z = \mathbb{E}\bigl[(T\wedge S)/\overline{F}(S)\bigr]$. But the Z's of different individuals in generation n are i.i.d. given the past which easily gives $\mathbb{E}D = \mathbb{E}P\cdot\mathbb{E}Z$. Now just insert the derived expressions for $\mathbb{E}P, \mathbb{E}Z$ together with (2.6) for $\mathbb{E}N$. □

Remark 2.9. As noted earlier, the LIFO PRI queue is regenerative when it is stable. Because the associated regenerative cycle has a time duration that is equal in distribution to D plus an independent exponential interarrival time, it follows that the number-in-system process is positive recurrent if and only if $\mathbb{E}D < \infty$. Hence, it is evident that when $\mathbb{E}\mathrm{e}^{\lambda S} < 2$, time-averages of indicator functionals of the number-in-system converge a.s. to their associated equilibrium expected values; see [3, VI.3]. This is to be contrasted against our earlier comment in which we pointed out that customer averages of indicator functionals of the time-in-system converge a.s. to their equilibrium values, even when $\mathbb{E}\mathrm{e}^{\lambda S} = 2$.

Remark 2.10. For the LIFO PRI M/M/1 queue with arrival rate λ and service rate μ, it is easily seen that

$$m(s) = \frac{1}{\lambda}\mathrm{e}^{\lambda s}(1-\mathrm{e}^{-\lambda s})(1+\lambda m), \tag{2.10}$$

$$m = \frac{1}{\mu-2\lambda}. \tag{2.11}$$

In this special setting, there is an instructive alternative argument that one can exploit. Define $m(t,s)$ as the expected time to departure of a customer with $S = s$ and who has currently been in service for t units of time. Then, assuming differentiability in t, we find that

$$\begin{aligned} m(t,s) &= \bigl(h + m(t-h,s)\bigr)(1-\lambda h) + (m + m(s,s))\lambda h + \mathrm{o}(h)\\ &= \bigl(h + m(t,s) - m'(t,s)h\bigr)(1-\lambda h) + (m + m(s,s))\lambda h + \mathrm{o}(h)\,,\\ m'(t,s) &= 1 + \lambda m + \lambda m(s,s) - \lambda m(t,s), \end{aligned} \tag{2.12}$$

where the prime denotes differentiation w.r.t. t. For fixed s, (2.12) is a standard linear ordinary differential equation with solution $\alpha(s)+\beta(s)\mathrm{e}^{-\lambda t}$. The obvious boundary condition $m(0,s) = 0$ gives $\beta(s) = -\alpha(s)$ and substituting back in (2.12) gives after easy algebra that $\alpha(s) = \mathrm{e}^{\lambda s}(1+\lambda m)/\lambda$ from which (2.10) follows. We then get

$$\begin{aligned} m = \mathbb{E}m(S,S) &= \int_0^\infty \frac{1}{\lambda}\mathrm{e}^{\lambda s}(1-\mathrm{e}^{-\lambda s})(1+\lambda m)\mu\mathrm{e}^{-\mu s}\,\mathrm{d}s \\ &= \frac{\mu(1+\lambda m)}{\lambda}\int_0^\infty (\mathrm{e}^{(\lambda-\mu)s}-\mathrm{e}^{-\mu s})\,\mathrm{d}s \\ &= \frac{\mu(1+\lambda m)}{\lambda}\Big(\frac{1}{\mu-\lambda}-\frac{1}{\mu}\Big) = \frac{1-\lambda m}{\mu-\lambda}\,. \end{aligned}$$

Solving for m gives (2.11).

We turn next to the computation of higher-order moments for this model. For $n \geq 1$, let $m_n(s) = \mathbb{E}D^n(s)$ and $m_n = \mathbb{E}D^n$. Then, (2.1) implies that $m_n(s)$ is the minimal non-negative solution of

$$\begin{aligned} m_n(s) = \mathbb{E}(T\wedge s)^n + \mathbb{P}(T\leq s)\mathbb{E}D^n + \mathbb{P}(T\leq s)\mathbb{E}D(s)^n \\ + \sum_{0\leq i_j<n, 1\leq j\leq 3}\binom{n}{i_1, i_2, i_3}\mathbb{E}T^{i_1}I(T\leq s)m_{i_2}m_{i_3}(s) \end{aligned}$$

from which the moments can be computed recursively in n. In particular, we can compute the heavy-traffic behavior of, for example, the variance of D from the above recursion, by considering a sequence of systems in which the appropriate moments of S and T converge, with θ simultaneously increasing to 2. A straightforward but tedious calculation then establishes that

$$\mathbb{V}\mathrm{ar}(D) \sim 2(2-\theta)^{-3}\mathbb{E}(\overline{F}(S)^{-1}-1)^2(ET\wedge S)/\overline{F}(S))^2$$

as θ increases to 2. Note that the standard deviation of D is scaling at least as fast as $(2-\theta)^{-3/2}$, while the mean is scaling as $(2-\theta)^{-1}$, which (of course) is typically inconsistent with the weak convergence of aD's to a limiting distribution for some sequence of scaling constants a. This provides additional support for Remark 2.6. We note that similar scaling behavior for the moments of time-in-system has been observed in the setting of LIFO preemptive resume queues; see Abate & Whitt [1].

We conclude this section with some partial results for the case of renewal (but non-Poisson) arrival traffic, cf. Remark 2.1.

Proposition 2.11. *The* LIFO PRI GI/D/1 *queueing system with* $S \equiv \ell$ *empties infinitely often a.s. if and only if*

$$\mathbb{E}\lfloor T/\ell\rfloor \geq 1\,. \tag{2.13}$$

Proof. Starting from a non-empty state, the queue length at arrival epochs evolves as a random walk $\xi_0, \xi_1, \xi_2, \ldots$ with increments distributed as $\Delta = 1 - \lfloor T/\ell\rfloor$ until the time of emptiness. But the PRI GI/D/1 system is stable if and only if this RW is recurrent which occurs if and only if $\mathbb{E}\Delta \leq 0$. This is precisely the stated condition. □

Proposition 2.12. *Assume the interarrival distribution F has an increasing failure rate* (IFR)*. Then, $\mathbb{E}\overline{F}(S)^{-1} \le 2$ is a necessary condition for the* LIFO PRI G/G/1 *queue to empty infinitely often a.s.*

Proof. Consider a modified queue which operates identically to the G/G/1 queue under consideration, except at service restart times. At the instant of each customer restart, the modified system independently generates from the distribution F the interarrival time of the next customer to arrive, rather than using the residual interarrival time associated with the renewal process. The IFR assumption implies that such a modified system generates customer arrivals more slowly than for the original system, so that the modified system must be stable whenever the original system is. The modified system has the property that, as for the Poisson case, the total number of progeny of the customer initializing the busy period is described by a Galton-Watson branching process. Following the same argument as in the Poisson case shows that the total number of progeny for this tree is finite a.s. when $\mathbb{E}\overline{F}(S)^{-1} \le 2$, proving the necessity for the original system. □

Note: The family tree in the original system is not a Galton-Watson tree, but this is not needed for the argument either!

A similar argument gives:

Proposition 2.13. *Assume the interarrival distribution F has a decreasing failure rate* (DFR)*. Then, $\mathbb{E}\overline{F}(S)^{-1} \le 2$ is a sufficient condition for the* LIFO PRI G/G/1 *queue to empty infinitely often a.s.*

3 Preemptive-Repeat-Identical Queues with Markovian Arrivals

We now consider the PRI queue (not the fixpoint problem (2.1)!!) with a Markovian arrival process. The state space for the underlying Markov process $J = \big\{J(t)\big\}_{t\ge 0}$ is $\{1,\dots,q\}$, the generator is $\boldsymbol{\Lambda}$ and as usual, the matrix giving the rate of transitions without arrivals is denoted by $\boldsymbol{C}$, that of transitions with an arrival by $\boldsymbol{D}$ (the ij'th element is the rate of an arrival in state i with Markov state j just after the transition. Thus $\boldsymbol{\Lambda} = \boldsymbol{C} + \boldsymbol{D}$ (assumed irreducible). For simplicity, the service time distribution G is assumed to be the same for all Markov states i, j.

Denote by τ the time of the first arrival. Then shall use repeatedly that $\mathbb{P}\big(\tau > s, J(s) = j\big)$, $\mathbb{P}\big(\tau \in \mathrm{d}t, J(t) = j\big)$ and $\mathbb{P}\big(J(\tau) = j\big)$ are the ij'th elements of the matrices $\mathrm{e}^{\boldsymbol{C}}$, resp. $\mathrm{e}^{\boldsymbol{C}t}\boldsymbol{D}\,\mathrm{d}t$, resp.

$$\boldsymbol{Q} = \int_0^\infty \mathrm{e}^{\boldsymbol{C}t}\boldsymbol{D}\,\mathrm{d}t = -\boldsymbol{C}^{-1}\boldsymbol{D}\,.$$

In particular, by irreducibility $\boldsymbol{Q}$ is a stochastic matrix, i.e. $\boldsymbol{Q}\boldsymbol{e} = \boldsymbol{e}$ where $\boldsymbol{e}$ is the column vector of ones..

Define $p_{ij} = \mathbb{P}_i\big(J(D) = j\big)$ as the probability that a customer starting service when the arrival process is in state i will eventually terminate service and will do so in state j (the period between these events of course include the busy periods initiated by interruptions of service by an arriving customer). Then stability is equivalent to the matrix $\boldsymbol{P} = (p_{ij})_{i,j=1,\dots p}$ being stochastic rather than properly substochastic. The following results gives an algorithm for computing $\boldsymbol{P}$ and thereby identifying the stability region. Numerical examples and further discussion is given below.

Theorem 3.1. *The matrix $\boldsymbol{P}$ is the minimal substochastic solution of the fixpoint equation $\boldsymbol{P} = \Psi(\boldsymbol{P})$ where*

$$\Psi(\boldsymbol{R}) = \int_0^\infty \big(\boldsymbol{I} - (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}\boldsymbol{R}\big)^{-1}\mathrm{e}^{\boldsymbol{C}s}\, G(\mathrm{d}s) \tag{3.1}$$

$$= \int_0^\infty \sum_{n=0}^{\infty}\big((\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}\boldsymbol{R}\big)^{n}\mathrm{e}^{\boldsymbol{C}s}\, G(\mathrm{d}s)\,. \tag{3.2}$$

If further $\boldsymbol{P}_k$ is defined by

$$\boldsymbol{P}_0 = \int_0^\infty \mathrm{e}^{\boldsymbol{C}s}\, G(\mathrm{d}s)\,, \quad \boldsymbol{P}_k = \Psi(\boldsymbol{P}_{k-1}),\ k \ge 1, \tag{3.3}$$

then $\boldsymbol{P}_k \uparrow \boldsymbol{P}$ as $k \to \infty$.

Proof. Let $\boldsymbol{P}(s)$ be defined as $\boldsymbol{P}$ but for a customer with service time $S = s$. Then

$$\boldsymbol{P} = \int_0^\infty \boldsymbol{P}(s)\, G(\mathrm{d}s), \tag{3.4}$$

$$\boldsymbol{P}(s) = (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}\boldsymbol{P}\boldsymbol{P}(s) + \mathrm{e}^{\boldsymbol{C}s}. \tag{3.5}$$

Indeed, (3.4) is clear and the second term on the r.h.s. of (3.5) is the contribution from the event that the customer terminates service without being preempted. In the first term,

$$(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q} = \int_0^s \mathrm{e}^{\boldsymbol{C}t}\boldsymbol{D}\,\mathrm{d}t \tag{3.6}$$

is the matrix giving the transition probabilities from the state at start of service until the first preemption. A preemption is followed by the busy period of the interrupting customer after which the customer himself resumes service. The two set of state changes of the arrival have transition matrices $\boldsymbol{P}$, resp. $\boldsymbol{P}(s)$, and from these observations (3.5) follows.

Since preemption does not occur w.p. 1, the first term on the r.h.s. of (3.6) is properly substochastic so that (3.5) can be solved for $\boldsymbol{P}(s)$ to give

$$\boldsymbol{P}(s) = \big(\boldsymbol{I} - (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}\boldsymbol{P}\big)^{-1}\mathrm{e}^{\boldsymbol{C}s} \tag{3.7}$$

and combining with (3.4) gives (3.1); (3.2) then just follows by standard power series formulas.

For the rest of the proof, define a sub-busy period (SBP) as the busy period of any of the preempting customers. The depth D of the family tree of the customer is then $D = 0$ if there are no SBP's, $D = 1$ if there is preemption but none of the preempting customers are themselves preempted and so on. It follows immediately that $\boldsymbol{P}_0$ is the contribution to $\boldsymbol{P}$ from the event $D = 0$. In the expression

$$\boldsymbol{P}_1 = \int_0^\infty \sum_{n=0}^{\infty}\big((\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}\boldsymbol{P}_0\big)^{n}\mathrm{e}^{\boldsymbol{C}s}\, G(\mathrm{d}s)$$

the matrix $\big((\mathrm{e}^{\boldsymbol{C}s} - \boldsymbol{I})\boldsymbol{C}^{-1}\boldsymbol{D}\boldsymbol{P}_0\big)^n$ is then the transition matrix for state changes from start of service until termination of the n'th SBP when any of SBP's $1, \ldots, n$ have depth 0. This shows that $\boldsymbol{P}_0$ is the contribution to $\boldsymbol{P}$ from the event $D \le 1$. Continuing in this manner gives that $\boldsymbol{P}_k$ is the contribution to $\boldsymbol{P}$ from the event $D \le k$. This shows that the sequence $\boldsymbol{P}_k$ is increasing with and that the limit includes all contributions from the event $D < \infty$ so the limit must be $\boldsymbol{P}$.

Finally, let $\boldsymbol{R}$ be any other substochastic solution (note that we need to assume that $\boldsymbol{R}$ is substochastic rather than just non-negative since otherwise the existence of the inverse in (3.1) may be a problem). Then $\boldsymbol{R}$ is at least the $n = 0$ term in (3.2) which equals $\boldsymbol{P}_0$. Applying Ψ and using induction then gives $\boldsymbol{R} \ge \boldsymbol{P}_k$ for all k and hence $\boldsymbol{R} \ge \boldsymbol{P}$. □

For deterministic service times, the calculations in the proof of Theorem 3.1 lead to an explicit stability criterion:

Proposition 3.2. *The* PRI MAP/D/1 *queue with $S \equiv \ell$ is stable if and only if*

$$\frac{\boldsymbol{\pi}\boldsymbol{C}\boldsymbol{e}}{\boldsymbol{\pi}\boldsymbol{C}(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})^{-1}\boldsymbol{e}} \le \frac{1}{2} \tag{3.8}$$

where $\boldsymbol{\pi}$ is the stationary distribution of $\boldsymbol{Q}$ written as row vector, i.e. $\boldsymbol{\pi}(\boldsymbol{C}+\boldsymbol{D}) = \boldsymbol{0}$.

Proof. In this case, $\boldsymbol{P} = \boldsymbol{P}(\ell)$ and (3.5) takes the form $\boldsymbol{0} = \boldsymbol{A}_1\boldsymbol{P}^2 + \boldsymbol{A}_0\boldsymbol{P} + \boldsymbol{A}_{-1}$ where

$$\boldsymbol{A}_1 = (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})(-\boldsymbol{C}^{-1}\boldsymbol{D}),\ \boldsymbol{A}_0 = -\boldsymbol{I},\ \boldsymbol{A}_{-1} = \mathrm{e}^{\boldsymbol{C}\ell}\,.$$

This is the equation for the $\boldsymbol{G}$-matrix in a continuous-time QBD (quasi-birth-death process) where the matrix of rates of state changes of the underlying Markov associated with upward jumps is $\boldsymbol{A}_1$ and the similar matrix for downward jumps is $\boldsymbol{A}_{-1}$, cf. [22], [3, XI.3]. To see that the $\boldsymbol{A}_1$ are legitimate QBD parameters, we must check that $\boldsymbol{A}_1, \boldsymbol{A}_{-1}$ are non-negative and that $\widetilde{\boldsymbol{Q}}\boldsymbol{e} = \boldsymbol{0}$ where $\widetilde{\boldsymbol{Q}} = \boldsymbol{A}_1 + \boldsymbol{A}_0 + \boldsymbol{A}_{-1}$. But non-negativity of $\boldsymbol{A}_1$ follows from non-negativity of $\boldsymbol{D}$ and the discussion around (3.6), that of $\boldsymbol{A}_{-1}$ is obvious since $\mathrm{e}^{\boldsymbol{C}\ell}$ is substochastic, and finally

$$\widetilde{\boldsymbol{Q}}\boldsymbol{e} = (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})\boldsymbol{e} - \boldsymbol{e} + \mathrm{e}^{\boldsymbol{C}\ell} = \boldsymbol{0}$$

since $-\boldsymbol{C}^{-1}\boldsymbol{D}$ is a transition matrix and hence $-\boldsymbol{C}^{-1}\boldsymbol{D}\boldsymbol{e} = \boldsymbol{e}$.

It follows by general results on the $\boldsymbol{G}$-matrix that $\boldsymbol{P}$ is stochastic if and only if the mean drift of the QBD is non-positive, i.e.

$$\widetilde{\boldsymbol{\pi}}\boldsymbol{A}_1\boldsymbol{e} \le \widetilde{\boldsymbol{\pi}}\boldsymbol{A}_{-1}\boldsymbol{e} \tag{3.9}$$

where $\widetilde{\boldsymbol{\pi}}$ is the stationary distribution of $\widetilde{\boldsymbol{P}}$. But $\widetilde{\boldsymbol{\pi}}$ is proportional to

$$\boldsymbol{\xi} = \boldsymbol{\pi}\big(-\boldsymbol{C}(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})^{-1}\big) = \boldsymbol{\pi}\big(-(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})^{-1}\big)\boldsymbol{C}$$

because $\boldsymbol{\xi}\boldsymbol{A}_1 = \boldsymbol{D} = -\boldsymbol{\pi}C$, $\widetilde{\boldsymbol{\xi}}\boldsymbol{A}_0 = -\boldsymbol{\xi}$, $\widetilde{\boldsymbol{\xi}}\boldsymbol{A}_{-1} = \boldsymbol{C} + \xi$ and so $\boldsymbol{\xi}\widetilde{\boldsymbol{P}} = \boldsymbol{0}$. These formulas also give $\boldsymbol{\xi}(\boldsymbol{A}_1 - \boldsymbol{A}_{-1}) = -2\boldsymbol{\pi}\boldsymbol{C}\boldsymbol{e} - \boldsymbol{\pi}\boldsymbol{C}(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}\ell})^{-1}$ and so (3.9) is indeed equivalent to (3.8). □

Example 3.3. For PRI M/D/1, we have $p = 1$, $\boldsymbol{e} = \boldsymbol{\pi} = 1$, $\boldsymbol{C} = -\lambda$ and condition (3.8) means $1 - \mathrm{e}^{-\lambda\ell} \le 1/2$, i.e. $\mathrm{e}^{\lambda\ell} \le 2$ which is the same as we found in Section 2 (as should be in the Poisson case).

Example 3.4. Consider PRI PH/D/1, with phase generator $\boldsymbol{T}$ and initial vector $\boldsymbol{\alpha}$ of the interarrival distribution. Here $\boldsymbol{C} = \boldsymbol{T}$, $\boldsymbol{D} = -\boldsymbol{T}\boldsymbol{e}\boldsymbol{\alpha}$ and it is known that $\boldsymbol{\pi} = -\boldsymbol{\alpha}\boldsymbol{T}^{-1}/-\boldsymbol{\alpha}\boldsymbol{T}^{-1}\boldsymbol{e}$. Thus $\boldsymbol{\pi}\boldsymbol{C}$ is proportional to $\boldsymbol{\alpha}$, and since $\alpha\mathbf{e} = \mathbf{1}$, condition (3.8) becomes

$$\frac{1}{2} \ge \frac{1}{\boldsymbol{\alpha}(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{T}\ell})^{-1}\boldsymbol{e}}\,, \quad \text{i.e. } \boldsymbol{\alpha}(\boldsymbol{I} - \mathrm{e}^{\boldsymbol{T}\ell})^{-1}\boldsymbol{e} \ge 2 \tag{3.10}$$

whereas our condition for the fixpoint problem (2.1) is $1/\overline{F}(\ell) \le 2$ which means $\boldsymbol{\alpha}\mathrm{e}^{\boldsymbol{T}\ell}\boldsymbol{e} \ge 1/2$.

Remark 3.5. For a general service time distribution, the $\boldsymbol{G}$-matrix approach of Theorem 3.2 leads into tree structures such as those considered in [22, Ch. 14] and references there, Van Houdt & Blondia [27] and He & Alfa [15].The closest of these references is [15], but only PRD is considered there. It is remarked in [22, p. 292] (see also [15, p. 283] and Walraevens et al. [29, p. 239] that no stability conditions in terms of model primitives are known.

We turn to the computation of the expected values $m_i = \mathbb{E}_i D$. Once $\boldsymbol{P}$ has been computed, the following scheme gives a procedure of similar complexity as a single step in the iteration for $\boldsymbol{P}$. Define $\boldsymbol{m}$ as the column vector with i'th entry m_i and

$$\boldsymbol{a}(s) = \boldsymbol{C}^{-1}(\mathrm{e}^{\boldsymbol{C}s} - \boldsymbol{I})\boldsymbol{e}, \quad \boldsymbol{A}_1(s) = (\boldsymbol{I} - \mathrm{e}^{\boldsymbol{C}s})\boldsymbol{Q}, \quad \boldsymbol{A}_2(s) = \bigl(\boldsymbol{I} - \boldsymbol{A}_1(s)\boldsymbol{P}\bigr)^{-1},$$
$$\boldsymbol{A}_{21} = \int_0^\infty \boldsymbol{A}_2(s)\boldsymbol{A}_1(s)\, G(\mathrm{d}s)\,, \qquad \boldsymbol{A}_{2a} = \int_0^\infty \boldsymbol{A}_2(s)\boldsymbol{a}(s)\, G(\mathrm{d}s)\,.$$

Theorem 3.6. *In the stable case,* $\boldsymbol{m} = (\boldsymbol{I} - \boldsymbol{A}_{21})^{-1}\boldsymbol{A}_{2a}$.

Proof. Let $m_{ij}(s) = \mathbb{E}_i\bigl[D;\, J(D) = j \bigm| S = s\bigr]$. Then $m_{ij}(s)$ can be split into an "instantaneous" part defined as the expected time until either service is completed without preemption or the first preemption occurs, and a "continuation" part defined as the expected time after the first preemption. The instantaneous part is $\mathbb{E}_i[s \wedge T;\, J(D) = j]$, and using similar arguments as those leading to (3.7) for the continuation part then gives

$$m_{ij}(s) = \mathbb{E}_i[s \wedge T;\, J(D) = j] + \sum_{k,\ell} \boldsymbol{e}_i^\top \boldsymbol{A}_1(s)\boldsymbol{e}_k\bigl[m_{kl}p_{\ell j}(s) + p_{k\ell}m_{\ell j}(s)\bigr]. \tag{3.11}$$

Indeed, $\sum_\ell m_{kl}p_{\ell j}(s)$ is the contribution from the busy period initited by the first interrupting arrival given $J(\tau) = k$ while $\sum_\ell m_{kl}p_{\ell j}(s)$ corresponds to the period after the customer himself resumes work.

Now

$$\begin{aligned}
\mathbb{E}_i[s \wedge T] &= s\mathbb{P}_i(T > s) + \mathbb{E}_i[T;\, T \le s] \\
&= s\boldsymbol{e}_i^\top \mathrm{e}^{\boldsymbol{C}s}\boldsymbol{e} + \int_0^s t\, \boldsymbol{e}_i^\top \mathrm{e}^{\boldsymbol{C}t}\boldsymbol{D}\boldsymbol{e}\,\mathrm{d}t = s\boldsymbol{e}_i^\top \mathrm{e}^{\boldsymbol{C}s}\boldsymbol{e} - \int_0^s t\boldsymbol{e}_i^\top \mathrm{e}^{\boldsymbol{C}t}\boldsymbol{C}\boldsymbol{e}\,\mathrm{d}t \\
&= s\boldsymbol{e}_i^\top \mathrm{e}^{\boldsymbol{C}s}\boldsymbol{e} - \boldsymbol{e}_i^\top\bigl\{s\mathrm{e}^{\boldsymbol{C}s} - \boldsymbol{C}^{-1}\mathrm{e}^{\boldsymbol{C}s} + \boldsymbol{C}^{-1}\bigr\}\boldsymbol{e} = \boldsymbol{e}_i^\top \boldsymbol{C}^{-1}(\mathrm{e}^{\boldsymbol{C}s} - \boldsymbol{I})\boldsymbol{e}\,,
\end{aligned}$$

where the third step used $\boldsymbol{De} = -\boldsymbol{CQe} = -\boldsymbol{Ce}$. This is just the i'th element $a_i(s)$ of $\boldsymbol{a}(s)$, and since $\boldsymbol{P}(s)$ has row sums 1 in the stable case, it follows by summing (3.11) over j and rewriting in matrix notation that

$$\begin{aligned}\boldsymbol{m}(s) &= \boldsymbol{a}(s)\boldsymbol{e} + \boldsymbol{A}_1(s)\boldsymbol{m} + \boldsymbol{A}_1(s)\boldsymbol{P}\boldsymbol{m}(s)\\ &= \boldsymbol{A}_2(s)\big[\boldsymbol{A}_0(s)\boldsymbol{e} + \boldsymbol{A}_1(s)\boldsymbol{m}\big].\end{aligned}$$

Integrating w.r.t. $G(\mathrm{d}s)$ gives $\boldsymbol{m} = \boldsymbol{A}_{0a} + \boldsymbol{A}_{12}\boldsymbol{m}$. From this the result follows. $\square$

Example 1. Consider PRI M/M/1 with arrival rate λ and service rate μ. Here all matrices are numbers and $\boldsymbol{D} = -\boldsymbol{C} = \lambda$. One gets $\boldsymbol{a}(s) = (1-\mathrm{e}^{-\lambda s})/\lambda$, $\boldsymbol{A}_1(s) = 1-\mathrm{e}^{-\lambda s}$, $\boldsymbol{A}_2(s) = \mathrm{e}^{\lambda s}$, $\boldsymbol{A}_{21} = \lambda/(\mu-\lambda)$, $\boldsymbol{A}_{2a} = 1/(\mu-\lambda)$ and so $m = 1/(2\mu-\lambda)$ in accordance with Proposition 2.3 (see e.g. Remark 2.10).

3.1 Numerical examples

For numerical illustration of the above results we considered the stability problem for PRI queues with Erlangian or hyperexponential interarrival times and exponential service times, denoted PRI E_q/M/1 and H_2/M/1. The critical point for stability was taken as the common point at which the deviations of $\boldsymbol{P}$ from being stochastic becomes marked and at which $\mathbb{E}D$ explodes to $+\infty$. Thereby Theorem 3.1 and Theorem 3.6 serve as a double-check., For these and other algorithmic details, see Section 5 of the Appendix.

Remark 3.7. For systems so complicated as PRI MAP/G/1, it is not clear that the stability condition has a so simple form as $\rho \le \rho_0$. It is not even clear that if the system is stable for a given service time distribution G, then it will also be so for a stochastically smaller $\widetilde{G}$ (see Remark 4.3 below for a related counterexample). In the examples of an exponential G to follow, the numerical calculations showed, however, that this was certainly the case.

Example 2. We first considered PRI E_q/M/1. The first entrance for each q in the following table is the critical value ρ_0 for the fixpoint problem (2.1) and the second the one ρ_q for PRI E_q/M/1.

$q=2$	$q=3$	$q=4$
0.44 0.36	0.35 0.32	0.29 0.29

Since E_q is IFR, this is in good agreement with Proposition 2.12 which predicts that $\rho_0 \ge \rho_q$.

Example 3. For an DFR example, we considered PRI H_2/M/1 with interarrival density

$$\theta\mu_1\mathrm{e}^{-\mu_1 t} + (1-\theta)\mu_2\mathrm{e}^{-\mu_2 t}$$

where we take $\mu_1 < \mu_2$ for uniqueness; thus θ is the weight of the component with the heaviest tail. One of the three degrees of freedom is just the scaling which is unimportant for the critical value of ρ and to report the results, we chose the remaining two as θ and the s.c.v. $\eta = \mathbb{V}\mathrm{ar}(T)/[\mathbb{E}T]^2$. It is easy to see that when

$\mu_1 < \mu_2$, the range of η for a given θ is the open interval from 1 to $2/\theta - 1$, and in the following table we use the values

$$\eta_1 = 1 + 0.1(2/\theta - 1), \eta_2 = 1 + 0.3(2/\theta - 1), \ldots, \eta_5 = 1 + 0.9(2/\theta - 1).$$

Thus the highest s.c.v. is in the upper right corner and equals $1+0.9(2/(1/8)-1) = 14.6$, corresponding to $\theta = 1/8$, μ_2/μ_1 =149.9. The values for a given pair (θ, η) have the same meaning as in Example 2.

θ	η_1	η_2	η_3	η_4	η_5
1/8	0.43 0.58	0.31 0.66	0.21 0.72	0.12 0.78	0.04 0.84
3/8	0.49 0.53	0.43 0.58	0.36 0.62	0.25 0.66	0.11 0.71
5/8	0.50 0.52	0.48 0.54	0.46 0.56	0.43 0.58	0.37 0.60
7/8	0.50 0.50	0.50 0.51	0.50 0.52	0.49 0.53	0.49 0.53

The table shows that the stability region for PRI $H_2/M/1$ is larger than for the fixpoint problem (2.1), as should be according to Proposition 2.13. It is also seen that the stability region is increasing in the s.c.v. η but decreasing in θ. The difference in the upper right corner of the table is in fact quite remarkable!

4 The LIFO Preemptive-Repeat-Different Queueing Model

As in our analysis of the "identical" model, we are interested in the equilibrium sojourn time (or time-in-system) D. Note that if we assume the arrival process is Poisson, then analogously to (2.1), D must be the minimal solution (in the sense of stochastic dominance) to the SFPE

$$D \stackrel{\mathcal{D}}{=} T \wedge S + \mathbb{I}(T \leq S)\big(D_1 + D_2\big), \tag{4.1}$$

where D_1, D_2, S, T are independent and D_1, D_2 are copies of the rv D. This type of SFPE has been studied in the literature on such equations, and the minimal non-negative solution D that is meaningful from a queueing viewpoint is known in that literature as the *exogenous solution.*

Theorem 4.1. *The* LIFO PRD *queue with Poisson*(λ) *arrivals empties infinitely often a.s. if and only if*

$$\mathbb{P}(S < T) = \mathbb{E}e^{-\lambda S} \geq 1/2\,. \tag{4.2}$$

More generally, in the case of a general distribution F for the rv T, the minimal solution D of (4.1) *is finite a.s. if and only if* $\mathbb{P}(S < T) \geq 1/2$.

Note that condition (4.2) implies (2.5) so that the stability region for LIFO PRD is larger than for LIFO PRI. This follows by Jensen's inequality applied to $1/x$: $\mathbb{E}[1/\overline{F}(S)] \geq 1/\mathbb{E}\overline{F}(S)$.

Proof. The family tree of the customer (defined in the obvious manner) is the family tree of a Galton-Watson process where an individual gets two children w.p. $\mathbb{P}(T \leq S)$ and dies without children w.p. $\mathbb{P}(T > S)$. The tree is infinite if and only if the offspring mean $2\mathbb{P}(T \leq S)$ is at most 1. This is precisely the stated condition. □

For non-Poisson arrivals, one has, as for PRI queues, that the SFPE no longer applies to the analysis of the queue. In particular, (4.1) no longer applies. However, the stability analysis of non-Poisson PRD queues is much more easily resolved than for PRI queues.

Theorem 4.2. *The* LIFO PRD *queueing system empties infinitely often a.s. if and only if*

$$\mathbb{E}U_G(T) \geq 2 \tag{4.3}$$

where $U_G = \sum_0^\infty G^{*n}$ *is the renewal function of* G.

Note: We use here the usual convention that 0 is counted as a renewal.

Proof. Let $\eta_0, \eta_1, \eta_2, \ldots$ be the number of customers in systems at the arrival epochs. Then $\{\eta_n\}$ is a random walk and the PrRD-Q system is stable if and only if this RW is recurrent which occurs if and only if $\mathbb{E}\Delta \leq 0$ where Δ is the RW increment. But let $S_1, S_2, \ldots$ be i.i.d. service times. Then

$$\mathbb{E}\Delta = 1 - \mathbb{E}\big[\#\{n : S_1 + \cdots + S_n \leq T\}\big] = 1 - \big(U_G(T) - 1\big)\,. \qquad \square$$

Remark 4.3. For GI/M/1, $U_G(t) = 1 + \mu t$ and so condition (4.3) becomes $1 + \mu\mathbb{E}T \geq 2$ which is the same as $\rho \leq 1$, the stability condition for the standard GI/M/1 queue (as is readily seen should be the case!). For GI/D/1 with $S \equiv \ell$, we have $U_G(t) = 1 + \lfloor t/\ell \rfloor$, and so the condition becomes $\mathbb{E}\lfloor T/\ell \rfloor \geq 1$ as in Proposition 2.11. These conclusions for GI/M/1 and GI/D/1 are indeed clear after some reflection! To see that condition (4.3) is the same as (4.2) in M/G/1, let $\widehat{G}[\alpha] = \mathbb{E}\mathrm{e}^{-\alpha S}$ denote the Laplace transform of G. Then

$$\widehat{G}[\lambda] = \int_0^\infty \mathrm{e}^{-\lambda s}\, G(\mathrm{d}s) = \int_0^\infty \mathbb{P}(T > s)\, G(\mathrm{d}s) = \mathbb{P}(T > S)\,,$$

$$\begin{aligned}
\mathbb{E}U_G(T) &= \mathbb{E}\sum_{n=0}^\infty G^{*n}(T) = \sum_{n=0}^\infty \mathbb{P}(S_1 + \cdots + S_n \leq T) \\
&= \sum_{n=0}^\infty \mathbb{E}\exp\{-\lambda(S_1 + \cdots + S_n)\} = \sum_{n=0}^\infty \widehat{G}[\lambda]^n \\
&= \frac{1}{1 - \widehat{G}[\lambda]} = \frac{1}{1 - \mathbb{P}(T > S)}
\end{aligned}$$

which readily implies the desired conclusion.

Beyond exponential and one-point distributions, phase-type distributions are the main class of distributions for which the renewal function is known. Suppose for example that G is Erlang(2) with density $s\mathrm{e}^{-s}$. Then ([3], Exercise 5.1, p. 91) $U_G(t) = \ 3/4 + t/2 + \mathrm{e}^{-2t}/4$, and so condition (4.3) becomes

$$2\mathbb{E}T + \mathbb{E}\mathrm{e}^{-2T} \geq 5. \tag{4.4}$$

One could easily believe that if F_1, F_2 are distributions of T such that F_1 is larger than F_2 in stochastic order, then stability for F_2 would imply stability of F_2 when the service time distribution is the same. However, since $\mathbb{E}T$ increases as F increases in stochastic order but $\mathbb{E}\mathrm{e}^{-2T}$ decreases, (4.4) easily provides counterexamples to this belief.

As for the PRI model, it is easy to compute the moments of D.

Proposition 4.4. *In the* LIFO PRD *queue with a Poisson*(λ) *arrival stream,*

$$m = \mathbb{E}D = \frac{1}{\lambda}\Big[\frac{1}{2\mathbb{E}\mathrm{e}^{-\lambda S} - 1} - 1\Big]. \tag{4.5}$$

Proof. Note that m is the finite-valued solution of

$$m = \mathbb{E}[T \wedge S] + \mathbb{P}(S > T)[m + m]$$

and that

$$\begin{aligned}
\mathbb{E}[T \wedge S] &= \int_0^\infty \mathbb{P}(T \wedge S > s)\,\mathrm{d}s = \int_0^\infty \mathbb{P}(T \wedge S > s)\,\mathrm{d}s \\
&= \int_0^\infty \mathbb{P}(T > s)\mathbb{P}(S > s)\,\mathrm{d}s = \int_0^\infty \mathrm{e}^{-\lambda s}\mathbb{P}(S > s)\,\mathrm{d}s \\
&= \frac{1}{\lambda}\big[1 - \mathbb{E}\mathrm{e}^{-\lambda s}\big].
\end{aligned}$$

The result then follows by easy algebra. □

Higher-order moments can now be computed recursively, as for the PRI model. As there, we see that the distribution of D does not converge weakly to a non-degenerate rv in the heavy-traffic setting, regardless of the normalization used.

5 Remarks

We collect here some further observations and remarks concerning our preemptive repeat models.

1. The preemptive models of this paper are somewhat related to the so-called RESTART models on tasks that need to be restarted after a failure. Their basic SFPE is $D(s) \stackrel{\mathcal{D}}{=} T \wedge s + \mathbb{I}(T \leq s)D^*(s)$ which obviously is related to (2.1) but simpler. Such models do not have the queueing structure of this paper, because the first task is not disturbed by other arrivals, as in our model. They had been investigated already at the time of Jaiswal's book [16] under the name *break-down systems*, and had a revival around 1990 motivated from computer reliability issues, see e.g. Trivedi et al. [20, 21, 8] (various other preemptive schemes are studied there as well). Recently, some key issues concerning probabilities of long delays have been considered in Asmussen et. al. [4, 5] and Jelenković et al. [17, 18]. Further related studies are in [11] and [2].
2. While the stability characterizations for our classes of LIFO queueing models are more complex than for the LIFO preemptive resume discipline, they share the interesting property that the performance of LIFO systems, as measured through typical time-in-system for a customer, typically degrades gracefully as the system moves towards and even through criticality. In particular, for LIFO PRI queues, the time-in-system r.v. converges weakly to a proper limit as θ increases to 2 (the "critical" point). In fact, even slightly beyond criticality, the

great majority of customers will experience finite-valued times-in-system, as measured through the probability $p = \mathbb{P}(D < \infty)$. Given that p is presumably continuous in the problem data when beyond criticality, the time-in-system would then further degrade gracefully as the system moves into a regime well beyond criticality. This is to be contrasted with first in-first out (FIFO) queues, in which a catastrophic degradation in the distribution of time-in-system occurs as the system moves towards criticality. In this sense, the time-in-system performance of LIFO systems is more robust than that of FIFO systems.

3. Given the above remark, an interesting question is the asymptotics of $p = \mathbb{P}(D < \infty)$ when the load is slightly larger than the critical value given by Theorem 2.3. Some preliminary insight into this question can be obtained by considering the M/M/1 queue with rates λ, μ of arrivals, resp. service. In this setting, the equilibrium time-in-system for FIFO systems is infinite for $\rho \geq 1$ so that $p = 0$ in that parameter region, whereas for our LIFO PRI model, we have the more intriguing relation

$$p \sim 1 - \Big(\frac{2(\rho - 1/2)}{|\log(\rho - 1/2)|}\Big)^{1/2}, \qquad \rho \downarrow 1/2. \tag{5.1}$$

To arrive at this, take $\lambda = 1$, $\mu = 2 - \epsilon$ and write $p(s) = \mathbb{P}(D < \infty \,|\, S = s)$, $p = 1 - \delta$. Fixed point considerations such as those leading to (2.1) then give

$$\begin{aligned} p(s) &= \overline{F}(s) + F(s)pp(s) = \frac{\overline{F}(s)}{1 - pF(s)} \\ &= \frac{\overline{F}(s)}{\overline{F}(s) + \delta F(s)} = \frac{1}{1 + \delta(\mathrm{e}^s - 1)}\,. \end{aligned}$$

It follows that

$$\begin{aligned} 1 - \delta = \mathbb{E}p(S) &= \int_0^\infty \frac{1}{1 + \delta(\mathrm{e}^s - 1)}(2 - \epsilon)\mathrm{e}^{-(2-\epsilon)s}\,\mathrm{d}s \\ &\sim \int_0^\infty \frac{1}{1 + \delta(\mathrm{e}^s - 1)}(2 - \epsilon + \epsilon s)\mathrm{e}^{-2s}\,\mathrm{d}s \\ &= \int_0^\infty \frac{1}{1 + \delta(\mathrm{e}^s - 1)} 2\mathrm{e}^{-2s}\,\mathrm{d}s + \epsilon \int_0^\infty (s - 1)\mathrm{e}^{-2s}\,\mathrm{d}s \\ &= \frac{1}{(1 - \delta)^3}\big[1 - 4\delta + 3\delta^2 - 2\delta^2 \log\delta\big] + (1/4 - 1/2)\epsilon \\ &\sim 1 - \delta + 2\delta^2|\log\delta| + (1/4 - 1/2)\epsilon \end{aligned}$$

where we used the `FullSimplify[Integrate[·]]` command in Mathematica to obtain the value of the integral involving δ. This gives

$$4\delta^2|\log\delta| \sim \epsilon, \quad \delta \sim \Big(\frac{\epsilon}{4|\log\delta|}\Big)^{1/2} \sim \frac{\epsilon^{1/2}}{\big(2|\log\epsilon|\big)^{1/2}}$$

which together with $\rho = 1/(2 - \epsilon) \sim 1/2 - \epsilon/4$ gives (5.1).

4. In the body of the paper, we have interpreted the r.v. D as the time-in-system of a customer. However, in the PRI setting an alternative possibility is to think of D as the duration of the busy period B initiated by the customer. In relation to item 3 above, note that $\mathbb{P}(B < \infty) = 1/\rho$ for standard FIFO M/M/1 queues with $\rho > 1$. Here $1/\rho \sim 1 - (\rho - 1)$ as $\rho \downarrow 0$.
5. The family tree of a customer in the setting of MAP arrivals can be seen as a multitype Galton-Watson tree provided the type of a customer is defined as the Markov state at his arrival. Stability is equivalent to this tree being finite which in turns is well-known to be the same as the offspring mean matrix $\boldsymbol{M}$ having spectral radius ≤ 1. The difficulty in directly applying this observation to the stability of the LIFO MAP/G/1 queue is that $\boldsymbol{M}$ is not directly available. In fact, it can be seen that $\boldsymbol{M} = \boldsymbol{QP}$ in the notation of Section 3 so that the computation of $\boldsymbol{M}$ is essentially equivalent to the problem of computing $\boldsymbol{P}$ that presented the main difficulty in Section 3.
6. A current trend in queueing theory is to generalize the assumption of Poisson arrivals to Lévy input, cf. [10] and references there. In this setting, customers with service requirement s arrive at rate $\nu(\mathrm{d}s)$ where ν is the Lévy measure. Of course, the arriving work to the M/G/1 preemptive repeat models considered in this paper are compound Poisson processes that are special cases of a Lévy process. An extension to full Lévy generality seems problematic in our preemptive repeat setting. For example, suppose that we consider a Lévy input with infinite activity, specifically one in which $\nu(s, \infty) = s^{-\alpha}$ with $\alpha \leq 1$ (needed for non-negativity) with $\nu(\mathrm{d}s) = 0$ for $s > 1$ (for simplicity). Note that such behaviour of ν close to 0 is in fact typical of well-behaved Lévy processes.

   Let D_z, S_z, λ_z etc. refer to the Lévy measure truncated at z, i.e. $\nu_z(\mathrm{d}s) = \nu(\mathrm{d}s)\mathbb{I}(s > z)$. This corresponds to Poisson arrivals at rate $\lambda_z = z^{-\alpha}$ and service time distribution $G_z(\mathrm{d}s) = \nu_z(\mathrm{d}s)/\lambda_z$. Here

$$\mathbb{E}\mathrm{e}^{\lambda_z S_z} = \int_z^1 \mathrm{e}^{s/z^\alpha} z^\alpha \nu(\mathrm{d}s) = \int_{z^{1-\alpha}}^{z^{-\alpha}} \mathrm{e}^y \frac{z^\alpha}{y^{\alpha+1} z^{\alpha(\alpha+1)}} z^\alpha \,\mathrm{d}y$$

   is of order at least $\mathrm{e}^{z^{-\alpha}} z^\alpha$ which goes to ∞ as $z \downarrow 0$. Hence Condition (2.4) fails for small z and since $D = D_0 \geq D_z$, D cannot be finite w.p. 1. So, such preemptive Lévy queues fail to be stable.

Acknowledgements

We are grateful to Onno Boxma, Mor Harchol-Balter and V. Ramaswami for valuable remarks and references. Special thanks go to Jonathan Goodman for providing the differential equation argument in Remark 2.10 that provided the initial scientific impetus for this project. Patrick Laub supplied helpful answers to our programming queries.

References

[1] J. Abate & W. Whitt (1997) Limits and approximations for the M/G/1 LIFO waiting-time distribution. *OR Letters* **20**, 199–206.

[2] Abhisek, M. Boon, M. Mandjes & R. Nunez Queija (2016) Congestion analysis of congested intersections. Working paper, 8 pp.

[3] S. Asmussen (2003) *Applied Probability and Queues* (2nd ed.). Springer-Verlag.

[4] S. Asmussen, P.M. Fiorini, L. Lipsky, T. Rolski & R. Sheahan (2008) Asymptotic behaviour of total times for jobs that must start over if a failure occurs. *Math. Oper. Res.* **33**, 932–944.

[5] S. Asmussen, L. Lipsky & S. Thompson (2016) Markov renewal methods in restart problems in complex systems. In *The Fascination of Probability, Statistics and their Applications, Essays in Honour of Ole E. Barndorff-Nielsen* (M. Podolskij et al., eds.), pp. 501–527. Springer-Verlag,

[6] B. Avi-Itzhak (1963) Preemptive repeat priority queues as a special case of the multipurpose server problem **I-II**. *Oper. Res.* **11**, 597–609, 610–619.

[7] D.A. Bini, G. Latouche & B. Meini (2003) Solving nonlinear matrix equations arising in tree-Like stochastic processes. *Linear Algebra and its Applications.* **366**, 39–64

[8] P.F. Chimento, Jr. & K.S. Trivedi (1993) The completion time of programs on processors subject to failure and repair. *IEEE Trans. on Computers* **42**(1).

[9] R.W. Conway, W.L. Maxwell & L.W. Miller (1968) *Theory of Scheduling.* Addison-Wesley.

[10] K. Debicki & M. Mandjes (2015) *Queues and Lévy Fluctuation Theory.* Springer-Verlag.

[11] T. Field (2206) An analysis of the preemptive repeat queueing discipline. Unpublished notes, 6 pp.

[12] D.P. Gaver, Jr. (1962)A waiting line with interrupted service, including priorities. *J. Roy. Statist. Soc.* **B**, 73–90.

[13] M. Harchol-Balter (2013) *Performance, Modeling and Design of Computer Systems.* Cambridge University Press.

[14] T.E. Harris (1963) *The Theory of Branching Processes.* Springer-Verlag.

[15] Q.-M He & A.S. Alfa (1998) The MMAP[K]/PH[K]/1 queues with a last-come-first-served preemptive service discipline *Queueing Systems* **29**, 269–291.

[16] N.K. Jaiswal (1968) *Priority Queues.* Elsevier.

[17] P. Jelenković & J. Tan (2013) Characterizing heavy-tailed distributions induced by retransmissions. *Adv. Appl. Probab.* **45**, 106–138.

[18] P. Jelenković & E. Skiani (2015) Distribution of the number of retransmissions of bounded documents. *Adv. Appl. Probab.* **47**, 425–44

[19] L. Kleinrock (1976) *Queueing Systems* vol. **II**: *Computer Applications.* Wiley.

[20] V. Kulkarni, V. Nicola & K. Trivedi (1986) On modeling the performance and reliability of multimode systems. *The Journal of Systems and Software* **6**, 175–183.

[21] V. Kulkarni, V. Nicola & K. Trivedi (1987) The completion time of a job on a multi-mode system. *Adv. Appl. Probab.* **19**, 932–954.

[22] G. Latouche & V. Ramaswami (1999) *Introduction to Matrix Analytic Methods in Stochastic Modelling.* SIAM.

[23] V. Limic (2001) A LIFO queue in heavy traffic. *Ann. Appl. Probab.* **11**, 301–331 .

[24] J. Nair, M. Andreasson, L. Andrew, S. Low, and J. Doyle (2010) On channel failures, file fragmentation policies, and heavy-tailed completion times. *Proceedings of IEEE INFOCOM*, 2010.

[25] M.F. Neuts (1989) *Structured Markov Chains of the $M/G/1$ Type and Their Applications.* Marcel Dekker.

[26] K. Sigman (1996) Queues under preemptive LIFO and ladder height distributions for risk processes: a duality. *Stochastic Models* **12**, 725–735.

[27] B. Van Houdt & C. Blondia (2001) Stability and performance of stack algorithms for random access communication modeled as a tree structured QBD Markov chain *Stochastic Models* **17**, 247–270.

[28] R.W. Wolff (1989) *Stochastic Modeling and the Theory of Queues.* Prentice–Hall.

[29] J. Walraevens, D. Flems & H. Bruneel (2006) The discrete-time preemptive repeat identical priority queue. *QUESTA* **53**, 231–243.

Appendix: Computational Experience

The computations of Section 3.1 were done in MATLAB, using the iterative scheme (3.3) to compute $\boldsymbol{P}$. In each step, $\Psi(\boldsymbol{P}_k)$ was evaluated via (3.1) using matrix inversion, rather than the power-series b (3.2). Matrix exponentials and the infinite integral were evaluated via MATLAB's routines

`expm`, resp. `integral(·,0,Inf,'ArrayValued',true)`.

The deviation of $\boldsymbol{P}$ from being stochastic was measured via two criteria, the deficit $\delta = 1 - \boldsymbol{e}^\top \boldsymbol{P}\boldsymbol{e}/q$ of the average rom sum from 1 and calculated values of $\mathbb{E}D$. Ideally, one expects a change point at the critical point for stability, where $\log_{10}\delta$ should jump from $-\infty$ to finite negative values and $\mathbb{E}D$ should go to ∞.

An illustration of the procedure is given in Fig. 1 for the M/M/1 case where the critical value is *known* as $\rho = 0.50$ (cf. Remark 2.4). Three numbers $K = 100, 500, 5000$ of iterations were used. It shows that 1) a sufficiently large value of K is crucial for getting the desired sharp distinction between stability and non-stability. One also notes that 2) $\log_{10}\delta$ is not $-\infty$ in the numerics, but the wiggles for small ρ indicates the numerical precision on δ is about 10^{-15}, 3) The algorithm does not produce ∞ above the critical point, but some other number. This is no contradiction since the expression $\boldsymbol{m} = (\boldsymbol{I} - \boldsymbol{A}_{21})^{-1}\boldsymbol{A}_{2a}$ is only valid assuming stability. The entries beyond that do not give the $\mathbb{E}_i[D;\, D < \infty]$ but rather the ratio of two integrals without any direct interpretation in terms of the queueing problem.

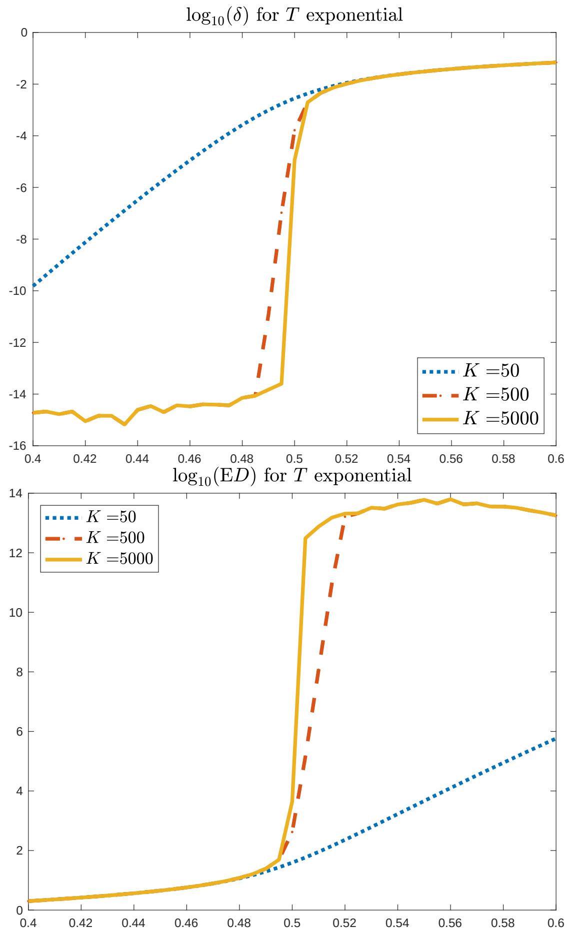
log10(δ) for T exponential
0
-2
-4
-6
-8
-10
-12
-14
-16
0.4
0.42
0.44
0.46
0.48
0.5
0.52
0.54
0.56
0.58
0.6
K =50
K =500
K =5000
log10(ED) for T exponential
14
12
10
8
6
4
2
0
0.4
0.42
0.44
0.46
0.48
0.5
0.52
0.54
0.56
0.58
0.6
K =50
K =500
K =5000

Figure 1: M/M/1

The numbers given in the E_p/M/1 and H_2/M/1 examples were visually assessed from plot similar to Fig. 1. A general feature was that for $p \geq 2$ states the change at the critical value was somewhat less sharp than for $p = 1$ (as in M/M/1) with the same number K of iterations, so that sometimes up to $K = 20{,}000$ iterations were required around the critical value; note that the interpretation of $\boldsymbol{P}_k$ in terms of the tree depth indicates that particularly many iterations are needed here. Obviously, the iteration scheme is time-consuming since each step involves computing an infinite integral of a function involving a matrix inverse and matrix-exponentials. Finite machine precisions sets further limit. Nevertheless, the results look reliable to us up to the given number of digits. One of the most demanding examples is presented in Fig. 2.

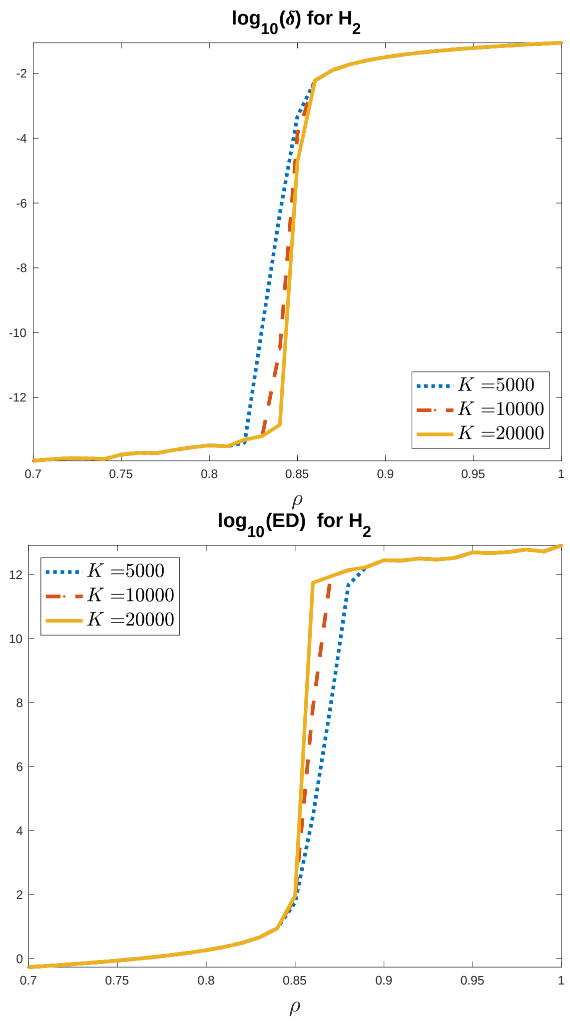

Figure 2: H_2 arrivals, $\theta = 1/8$, $\eta = 14.6$