Scheduling Analysis with Martingales

Felix Poloczek[a,b], Florin Ciucu[a]

[a]*University of Warwick, Departement of Computer Science, Coventry, CV4 7AL, UK*
[b]*Technische Universität Berlin, Research Group INET, MAR 4-4, Marchstr. 23, 10587 Berlin, Germany*

Abstract

This paper proposes a new characterization of queueing systems by bounding a suitable exponential transform with a martingale. The constructed martingale is quite versatile in the sense that it captures queueing systems with Markovian and autoregressive arrivals in a unified manner; the second class is particularly relevant due to Wold's decomposition of stationary processes. Moreover, using the framework of stochastic network calculus, the martingales allow for a simple handling of typical queueing operations: 1) flows' *multiplexing* translates into multiplying the corresponding martingales, and 2) *scheduling* translates into time-shifting the martingales. The emerging calculus is applied to estimate the per-flow delay for FIFO, SP, and EDF scheduling. Unlike state-of-the-art results, our bounds capture a fundamental exponential leading constant in the number of multiplexed flows, and additionally are numerically tight.

The theory of effective bandwidth emerged in the 1990s as a *unified framework* to analyze the queueing behavior of broad classes of arrivals (e.g., deterministically regulated, Markovian, long-range dependent); for an excellent overview see Kelly [1]. The effective bandwidth is associated to an arrival flow and is essentially a number between the flow's average and peak rates, depending on some predefined Quality-of-Service (QoS) constraint (e.g., margins on the buffer overflow probabilities). Effective bandwidths posses a fundamental additive property which makes them very attractive from an engineering point of view, e.g., admission control under some QoS requirements can be performed by allowing flows as long as the sum of their effective bandwidths does not exceed the available capacity.

This attractive additive property, however, was raised up as a fundamental drawback of effective bandwidths by Shroff and Schwartz [2]. Indeed, since only Poisson flows are closed under superposition, the effective bandwidth framework is prone to very pessimistic estimates. An alternative convincing technical explanation, on the pitfalls of effective bandwidths, was provided by Choudhury *et al.* [3]. Considering in particular the delay distribution W of some multiplexed flows, the corresponding *effective bandwidth approximation* states *asymptotically* that

$$\mathbb{P}\left(W > d\right) \sim \alpha e^{-\theta d} ,\tag{1}$$

where α is the *asymptotic constant*, θ is the *asymptotic decay rate*, and $f(x) \sim g(x)$ means that $f(x)/g(x) \rightarrow 1$ as $x \rightarrow \infty$. While the asymptotic decay rate θ is exact, the asymptotic constant α can be a very loose estimate. Indeed, for burstier than Poisson flows, it was convincingly shown by numerical results that, *non-asymptotically* (in d),

$$\mathbb{P}\left(W > d\right) \approx \kappa^{\#flows} e^{-\theta d} ,\tag{2}$$

where $0 < \kappa < 1$. Sharp bounds capturing this behavior were formally obtained by Duffield [4] in the case of Markov-Modulated On-Off processes satisfying a certain burstiness condition.

From a technical point of view, the inaccuracy (in non-asymptotic regimes) of the effective bandwidth approximation from Eq. (1) stems from applying Boole's inequality, i.e.,

$$\mathbb{P}(\sup_n X_n \geq \sigma) \leq \sum_n \mathbb{P}(X_n \geq \sigma) .\tag{3}$$

which is typically used in the large deviations framework to bound the supremum of a stochastic process. It is known that this inequality is very loose, especially in non-Poisson scenarios (see Talagrand [5]). To improve the accuracy a different approach was undertaken in [4] by extending the classical GI/GI/1 bounds of Kingman [6]. Concretely, the key idea is to avoid Boole's inequality and apply instead Doob's inequality

$$\mathbb{P}(\sup_n X_n \geq \sigma) \leq \mathbb{E}[X_0]\sigma^{-1} \ , \tag{4}$$

for a suitably constructed *martingale* X_n in the case of Markov-Modulated On-Off (MMOO) processes.

Let us emphasize at this point that Eq. (2) holds at the *aggregate level* in a flows' multiplexing scenario. An immediate question of practical interest, which constitutes the scope of this paper, concerns whether this fundamental result holds at the *per-flow level* in a scheduling scenario as well. Unfortunately, classical state-of-the art results in scheduling systems with Markovian arrivals share the pitfall of Boole's inequality; see Courcoubetis and Weber [7] for FIFO, Berger and Whitt [8] and Wischik [9] for Static Priority (SP), and Sivaraman and Chiussi [10] for Earliest-Deadline-First (EDF), and Bertsimas *et al.* [11] for Weighted-Fair-Queueing (WFQ). More recently, Ciucu *et al.* [12] demonstrated that Eq. (2) does hold at the per-flow level for FIFO, SP, EDF, and WFQ scheduling, in the specific case of Markov fluid processes. The central idea consists in invoking Doob's inequality on a suitable martingale construction by Ethier and Kurz ([13], p. 175), after appropriately decoupling scheduling using the framework of the stochastic network calculus.

In this paper we develop a *unified framework* in which Eq. (2) holds in great generality, at the *per-flow level*, for FIFO, SP, and EDF scheduling. To this end, we propose a novel representation of a queueing system through a *martingale-envelope*. Moreover, we integrate it in the framework of the stochastic network calculus (SNC), which is a fairly general and unified framework to compute bounds in queueing networks for a broad class of arrivals and scheduling algorithms (see Ciucu and Schmitt [14]). Unlike existing SNC envelope models, which set bounds on the arrival processes alone, the proposed martingale-envelope sets bounds on a suitable exponential transformation, involving both the arrival process and the allocated capacity. The crucial insight enabling the developed unified framework is that typical queueing operations simply translate into martingale operations:

1. *Multiplexing* of flows translates into multiplying the corresponding martingales.
2. *Scheduling* translates into time-shifting the martingales, corresponding to the scheduled flows, at a specific shifting time parameter depending on the scheduling algorithm itself.

The second operation in particular highlights the instrumental role of the emerged SNC martingale framework to deal with the difficult problem of scheduling in a unified manner, roughly by decoupling scheduling through a *shifting parameter*; this shifting parameter can be tuned depending on scheduling, i.e., FIFO, SP, and EDF.

We apply our unified framework to the class of *discrete-time* Markovian arrivals, and demonstrate the first-time manifestation of Eq. (2) at the per-flow level, for FIFO, SP, and EDF; our results can be regarded as per-flow level extensions of the aggregate level results by Duffield [4]. Remarkably, unlike Markov fluid arrivals for which Eq. (2) always holds at the per-flow level (see [12]), the discrete-time counterpart bounds only hold in certain burstiness scenarios; we will explain this idiosyncrasy by an underlying embeddability property of Markov chains.

We will also consider p-order autoregressive processes which can approximate the *whole* class of stationary processes (this property is typically referred to as Wold's decomposition, [15], p. 187). Although the autoregressive processes are also Markovian, their particular representation allows for a closed-form derivation of the performance bounds (the more general Markovian processes are subject to bounds in terms of implicit eigenvalues/vectors equations). More remarkably, unlike the results from [16], p. 340, which yield bogus (infinite) bounds when fitted for unbounded increment distributions, our results provide numerically accurate bounds.

For the rest of the paper we first develop the theory of martingale-envelopes in Section 1. Then, in Section 2, we apply the emerging SNC framework to several classes of processes (independent increments, general Markovian arrivals, and p-order autoregressive processes). We finally discuss several interesting technical issues in Section 3 and then conclude the paper.

1. A Calculus with Martingale-Envelopes

(a) Scenario 1: single flow   (b) Scenario 2: multiplexed flows   (c) Scheduling abstraction

Figure 1: Two scenarios: (a) consists of a single flow A, whereas (b) has an additional cross-flow A', in (c) the cross-flow from (b) is encoded in the dynamic service process S.

We consider two queueing scenarios as depicted in Figure 1, in a *discrete-time* model. In the first scenario (Figure 1(a)) a single flow A arrives at a server with capacity $c > 0$, whereas in the second (Figure 1(b)), two flows A (through-flow) and A' (cross-flow) compete for a shared server with total capacity $C = c + c'$.

The cumulative arrivals are given by stochastic processes

$$A(m,n) = \sum_{k=m+1}^{n} a_k \ , \quad A'(m,n) = \sum_{k=m+1}^{n} a'_k \ , \tag{5}$$

where $(a_n)_n$ and $(a'_n)_n$ are the instantaneous arrival processes which are throughout assumed to be stationary. As a consequence of Kolmogorov's extension theorem, both processes $(a_n)_{n \in \mathbb{N}}$ and $(a'_n)_{n \in \mathbb{N}}$ can be extended to stationary processes $(a_n)_{n \in \mathbb{Z}}$ and $(a'_n)_{n \in \mathbb{Z}}$ having the same finite dimensional distributions.

In the second scenario, there exists a variety of *scheduling policies* determining the priority of the data from flow A and A', respectively. In this paper we will consider *static priority* (SP), *first in, first out* (FIFO), and *earliest deadline first* (EDF).

The network calculus approach to address scheduling queueing systems is to transform the system from Figure 1(b) into the system from Figure 1(c). The transformation occurs by encoding information about the capacity, the cross-flow, and the scheduling into a single *service process* $S(m,n)$ which satisfies

$$D(n) \geq (A * S)(n) := \inf_{0 \leq k \leq n} \{A(0,k) + S(k,n)\} \ , \tag{6}$$

for *any* arrival flow $A(n)$. In some sense, the service process $S(m,n)$ is intimately related to the impulse-response of a linear and time invariant (LTI) system (for a discussion of this analogy see, e.g., [17, 18, 14]).

The performance metrics we are interested in are 1) the stationary *queue size* Q, i.e., the amount of data in the system at time n, and 2) the *virtual delay*

$$W(n) := \inf\{k \in \mathbb{N} \mid A(n-k) \leq D(n)\} \ ,$$

i.e., the time a data unit would have stayed in the system had it departed at time n. By the stationarity assumption, Q has the following representation (see [16])

$$Q :=_{\mathcal{D}} \sup_{n \in \mathbb{N}} \{A^r(0,n) - Cn\} \ , \tag{7}$$

where A^r stands for the *reversed* process, i.e.,

$$A^r(m,n) := \sum_{k=m+1}^{n} a_{-k} \ ,$$

(where by convention $A^r(0,0) := 0$) and $=_{\mathcal{D}}$ denotes equality in distribution.

We next introduce our characterization of a queueing system by a certain supermartingale. First we introduce two helpful technical notations:

Notation 1. *Denote by $\overrightarrow{a_n}$ the p-dimensional vector*

$$\overrightarrow{a_n} := (a_n, a_n + a_{n-1}, \ldots, a_n + \cdots + a_{n-p+1}) = \left(\sum_{k=1}^{i} a_{n-k+1}\right)_{1 \leq i \leq p} .$$

Notation 2. *For functions $h_1, \ldots, h_p$ let Πh denote the product*

$$\Pi h(x_1, \ldots, x_p) := \prod_{i=1}^{p} h_i(x_i) .$$

For brevity, we omit the parameter p in Notation 2, because its value is clear from the context. We remark that we will consider $p = 1$ for the class of Markovian arrivals (see Section 2.2), and any values of p for the class of p-order auto-regressive processes (see Section 2.3).

Definition 3 (MARTINGALE-ENVELOPE). *For $p \geq 0$ and monotonically increasing functions $h_1, \ldots, h_p : \mathbb{R}_+ \to \mathbb{R}_+$, and $\theta > 0$, we say the flow A admits a $(\Pi h, \theta, c)$-martingale-envelope if for every $m \geq 0$ the process*

$$\Pi h\left(\overrightarrow{a_n}\right) e^{\theta(A^r(m,n) - (n-m)c)} \leq M_m(n) \tag{8}$$

is almost surely bounded by a supermartingale M_m.

An intuition for this definition is the following: In order to keep a queueing system in a stable regime, by Loynes' condition, the average arrival rate has to be strictly less than the service rate. If one ignores the positivity constraint on the buffer, its expected increment (drift) is negative and thus the buffer content 'resembles' a supermartingale. The conceptual reason for the exponential transform is that its shape directly determines the decay rate of queueing metrics (which for Markovian arrivals are exponential). From a technical point of view, the (convex) exponential transform assigns more weight to larger arrivals, reducing the negative drift and consequently the gap between the constructed supermartingale and a martingale. Moreover, since Doob's inequality does not differentiate between a supermartingale and a martingale, one looks to minimize the previous gap by maximizing the decay factor θ, which eventually determines the decay rate of the queueing metrics. Finally, the function h compensates for potential correlations among the increments; in particular, for i.i.d. increments, h is a constant.

The monotonicity of the h_i is a technical condition needed for the following important Lemma:[1]

Lemma 4. *For $\sigma > 0$, let*

$$N := \inf\{n \geq 0 \mid A^r(0, n) - cn \geq \sigma\} \tag{9}$$

denote the first point in time where the supremum in Eq. (7) is attained. Then for any $k \geq 1$,

$$\sum_{i=0}^{k-1} a_{N-i} \geq kc .$$

For the special case when $k = 1$, the inequality in Lemma 4 can be slightly strengthened to $A_N \geq \tau$, where τ is defined by

$$\tau := \inf\{x > c \mid \mathbb{P}(a_n \in [x, \infty)) > 0\} ,$$

i.e., the smallest possible instantaneous arrival such that the buffer content increases. Note that this is only of importance if discrete distributions are considered. For continuous distributions τ is simply equal to c, and the statement is contained in Lemma 4.

The next theorems and corollaries are the central results, describing how martingale-envelopes can be used to derive bounds on the performance metrics Q and W. We start with the first scenario from Figure 1(a), i.e., considering the case of a single flow:

---

[1] For the sake of readability, all the proofs are given in the Appendix.

Theorem 5 (SINGLE FLOW BOUND)**.** *If the flow A admits a (h, θ, c)-martingale-envelope, then we have the following upper bound on the backlog and the virtual delay respectively:*

$$\mathbb{P}(Q \geq \sigma) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]}{\Pi h(c, 2c, \ldots, pc)} e^{-\theta\sigma} \ , \quad \mathbb{P}(W(n) \geq k) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]}{\Pi h(c, 2c, \ldots, pc)} e^{-\theta ck} \ .$$

Consider now the second scenario from Figure 1(b): two single flows A and A' with allocated capacities c and c', respectively, are multiplexed into one queueing system with a shared total capacity of $C = c + c'$. The resulting system can be analyzed in two different ways: Firstly, for the aggregate system both metrics Q and W can be estimated (*aggregate analysis*), and secondly, the virtual delay W for a single flow in the multiplexed system can be analyzed for several scheduling policies (*per-flow analysis*).

For both tasks, a technical definition is required:

Definition 6. *For two monotonically increasing functions $h, h' : \mathbb{R}_+ \to \mathbb{R}_+$, define the $(\min, \times)$-convolution by*

$$(h \otimes h')(t) := \inf_{0 \leq s \leq t} h(s)h'(t - s) \ ,$$

for all $t \in \mathbb{R}_+$. As above, if h and h' are families of functions, $\Pi(h \otimes h')$ is defined componentwise:

$$\Pi h \otimes h' = \prod_i h_i \otimes h'_i \ .$$

It is easy to check that $h_1 \otimes h_2$ is monotonically increasing as well, and that, by definition, for all a, b:

$$h \otimes h'(a + b) \leq h(a)h'(b) \ . \tag{10}$$

1.1. Aggregate Analysis

The next theorem addresses the *aggregate analysis* for the queueing system with aggregate arrivals $A + A'$:

Theorem 7 (AGGREGATE ENVELOPE)**.** *Assume two independent arrivals A and A' described by martingale-envelopes with parameters $(\Pi h, \theta, c)$ and $(\Pi h', \theta, c')$, respectively. Then the aggregate flow $A + A'$ admits a $(\Pi h \otimes h', \theta, C)$-martingale-envelope, where $C := c + c'$.*

The advantage of this theorem is that an aggregate flow can be handled in the same way as a single flow, e.g., for the constructed martingale-envelope, Theorem 5 can be evoked to derive the bounds on the backlog Q and the virtual delay W.

Note that in Theorem 7 the flows are required to be *homogeneous* in the sense that they admit the same θ in their respective envelopes. If this is not the case, the following transform of martingale-envelopes can be used:

Lemma 8. *If A admits a $(\Pi h, \theta, c)$-martingale-envelope and $\theta' < \theta$, then A admits a $\left(\Pi h^{\frac{\theta'}{\theta}}, \theta', c\right)$-martingale-envelope as well.*

1.2. Per-Flow Analysis

We now turn to the *per-flow analysis* of flow A in the multiplexed queueing system equipped with a scheduling policy (Figure 1(b)). The key element is the following technical lemma:

Lemma 9. *Assume the same situation as in Theorem 7. Then for every $l \geq 0$ and $\sigma > 0$ the following bound on the sample path holds:*

$$\mathbb{P}\left(\sup_{n \geq l}\{A^r(l, n) + A'^r(0, n) - Cn\} \geq \sigma\right) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta(\sigma + cl)} \ .$$

The crucial parameter in Lemma 9 is the parameter l, indicating how many points in time the process A is delayed. This parameter can be adjusted according to the scheduling policy under consideration, or more precisely to the expression of the service process S depicted in Figure 1(c). We will next apply Lemma 9 and properly tune the parameter l for SP, FIFO, and EDF scheduling.

Let us first describe a common step. Let D denote the departure process of flow A. For every policy, for which a service process S was constructed, the bounding procedure starts with a computation similar to the one of the virtual delay in Theorem 5:

$$\mathbb{P}(W(n) \geq k) = \mathbb{P}(A(0, n-k) \geq D(n)) \leq \mathbb{P}\left(A(0, n-k) \geq A * S(n)\right)$$
$$= \mathbb{P}\left(\sup_{0 \leq m \leq n} \{A(m, n-k) - S(m, n)\} \geq 0\right) \leq \mathbb{P}\left(\sup_{n \geq k} \{A^r(k, n) - S^r(0, n)\} \geq 0\right) , \quad (11)$$

where we again used the monotonicity of A and the reversed representation.

Static Priority (SP). This scheduling policy always gives priority to the cross-flow A'. The service process $S(m, n)$ is given by (see [19]):
$$S(m, n) = [C(n-m) - A'(m, n)]_+ , \quad (12)$$
where $[x]_+ = \max\{0, x\}$.

Corollary 10 (SP PER-FLOW BOUND). *Consider the situation as in Theorem 7, with SP as the scheduling policy. Then for the virtual delay $W(n)$ holds:*

$$\mathbb{P}(W(n) \geq k) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta c k} .$$

First In, First Out (FIFO). For FIFO the service process $S(m, n)$ is given by (see [20]):
$$S(m, n) = [C(n-m) - A'(m, n-x)]_+ 1_{\{n-m>x\}} , \quad (13)$$
where $x \geq 0$ is a parameter freely chosen, but fixed.

Corollary 11 (FIFO PER-FLOW BOUND). *Consider the situation as in Theorem 7, with FIFO as the scheduling policy. Then for the virtual delay $W(n)$ holds:*

$$\mathbb{P}(W(n) \geq k) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta C k} .$$

Note the difference in the decay rate: Whereas for SP it is the per-flow capacity c, for FIFO we have the total capacity $C = c + c'$.

Earliest Deadline First (EDF). Now consider the case of EDF scheduling. Let d and d' denote the relative deadlines for the data units of flows A and A', respectively. The service process $S(m, n)$ is given by (see [21]):
$$S(m, n) = [C(n-m) - A'(m, n-x+\min\{x,y\})]_+ 1_{\{n-m>x\}} , \quad (14)$$
where $x \geq 0$ is again a free parameter, and $y := d - d'$ denotes the difference between the respective deadlines. It is convenient to distinguish between the cases $y \geq 0$ and $y < 0$.

Let us first consider the case $y \geq 0$:

Corollary 12 (EDF PER-FLOW BOUND, $y \geq 0$). *Assume EDF is used as scheduling policy, $y \geq 0$, and consider the situation as in Theorem 7. Then for the virtual delay $W(n)$ holds:*

$$\mathbb{P}(W(n) \geq k) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta(Ck - c'\min\{k,y\})} .$$

Consider now the case $y = d - d' < 0$. This is more difficult as now $\min\{k, y\} = y < 0$, so that for

$$n_0 \in B := \{n \geq k \mid n < k - y\} \ ,$$

the argument $n_0 - k + \min\{k, y\}$ is negative as well. By definition (again from [21]), for those $n_0 \in B$:

$$A'^r(n_0 - k + \min\{k, y\}) = 0 \ . \tag{15}$$

Corollary 13 (EDF Per-Flow Bound, $y < 0$). *Assuming EDF scheduling with $y < 0$, for the virtual delay $W(n)$ holds:*

$$\mathbb{P}(W(n) \geq k) \leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta(Ck + c'y)} + \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]}{\Pi h(c, 2c, \ldots, pc)} e^{-\tilde{\theta}Ck} \ ,$$

where $\tilde{\theta}$ is the parameter to which the flow A admits a $(\Pi h, \tilde{\theta}, C)$-martingale-envelope. Note that as $C > c$, such a $\tilde{\theta}$ exists and is greater than θ.

2. Applications

In this section we demonstrate the versatility of the proposed martingale-envelope calculus to address several broad classes of arrival processes: with independent increments, with Markovian increments, and p-order autoregressive.

2.1. Processes with Independent Increments

One of the simplest traffic model is given by a process with independent increments. Although not realistic, it is included here because it provides a good intuition on how the martingale-envelope calculus works. Let $a_1, a_2 \ldots$ denote nonnegative i.i.d. random variables. The arrival process is thus $A(m, n) = \sum_{k=m+1}^{n} a_k$. Let the capacity $c > 0$ satisfy the two stability conditions

$$\mathbb{E}[a_1] < c < \sup a_1 \ , \tag{16}$$

to avoid the trivial scenarios of no queueing at all and infinite queue size, respectively.

Lemma 14. *In the situation above there is a $\theta_0 > 0$ such that A admits a $(1, \theta_0, c)$-martingale-envelope.*

Combining the martingale-envelope from Lemma 14 with the general theory from Section 1 the following bounds hold:

Corollary 15. *Consider an i.i.d. arrival flow $(a_n)_n$, and a capacity c such that the condition from Eq. (16) holds. Then for this single flow:*

$$\mathbb{P}(Q \geq \sigma) \leq e^{-\theta_0 c} \ , \quad \text{and} \quad \mathbb{P}(W(n) \geq k) \leq e^{-\theta_0 ck} \ .$$

With an additional i.i.d. cross-flow $(a'_n)_n$ and capacity c' (satisfying the corresponding stability conditions), for flow A holds in the multiplexed queueing system under scheduling:

FIFO: $\quad \mathbb{P}(W(n) \geq k) \leq e^{-\theta Ck}$ $\qquad$ SP: $\quad \mathbb{P}(W(n) \geq k) \leq e^{-\theta ck}$

EDF1: $\quad \mathbb{P}(W(n) \geq k) \leq e^{-\theta(Ck - c'\min\{k,y\})}$ $\qquad$ EDF2: $\quad \mathbb{P}(W(n) \geq k) \leq e^{-\theta(Ck + c'y)} + e^{\tilde{\theta}c'k} \ ,$

where θ'_0 is the parameter in the martingale-envelope for A', $\theta = \min\{\theta_0, \theta'_0\}$, $y = d - d'$, $C = c + c'$, and $\tilde{\theta}$ is the parameter of flow A in the system with total capacity C.
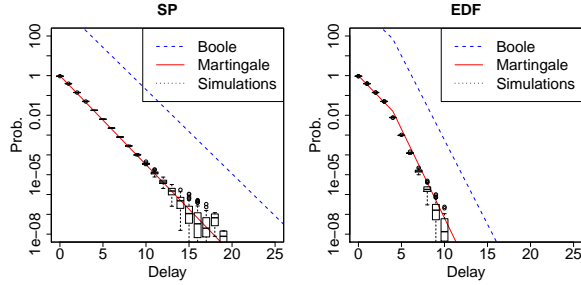
Figure 2: CCDF of the packet delay of $10 + 10$ exponentially distributed subflows with $\lambda = 1$, utilization $\rho = 0.95$, and, for EDF, $y = d - d' = 4$.

EDF1 and EDF2 correspond to the cases $y \geq 0$ and $y < 0$, respectively (see Corollaries 12 – 13).

Note that the aggregate analysis of the whole system (as in Subsection 1.1) is contained in the first part of Lemma 15, as the resulting aggregate flow $(a_n + a'_n)_n$ is still i.i.d.

In Figure 2 simulations of the MMOO and the corresponding bounds for SP and EDF are displayed[2]. The Martingale bounds (from Corollary 15) almost match the simulations, whereas the bounds computed with Boole's inequality are off by several orders of magnitude. As a side remark, all comparisons between bounds on virtual delays and packet delays account for the underlying Palm change of measure; moreover, we restrict to SP and the first EDF scenario; FIFO is a particular case of EDF.

2.2. Processes with Markovian Increments

The previous independence assumption on the increments is now replaced by a *Markovian correlation structure*, i.e., the process a_n is a Markov chain with state space $S = \{s_i \mid 1 \leq i \leq m\}$. To ensure stationarity, we assume a_n to be in steady state.

Let π denote the stationary distribution and T the transition matrix of the *reversed process*, i.e.,

$$\pi(i) = \mathbb{P}(a_n = s_i) \quad \text{and} \quad T(i,j) = \mathbb{P}(a_{n-1} = s_j \mid a_n = s_i) .$$

In many cases the Markov chain is *reversible* and the matrix T coincides with the transition matrix of a_n itself. Now, for any $\theta \geq 0$, let T_θ denote the *exponentially transformed* transition matrix, i.e.,

$$T_\theta(i,j) = T(i,j)e^{\theta s_j} .$$

Clearly, $T = T_0$. Further, let $\lambda(\theta)$ denote the spectral radius of T_θ and v a corresponding eigenvector. The function $\lambda(\theta)$ plays a similar role as the moment generating function φ_1 in Section 2.1. It can be shown (see [22]) that v can be chosen to be positive and that

$$\lambda(\theta) = \lim_{n \to \infty} \mathbb{E}[e^{\theta A(n)}]^{\frac{1}{n}} ,$$

so especially $\lambda(\theta) \geq 1$. Further, if the usual stability condition $\mathbb{E}[a_n] < c$ holds,

$$\frac{d}{d\theta}\lambda(\theta)\Big|_{\theta=0} = \lim_{n \to \infty} \frac{1}{n}\mathbb{E}[A_n] = \mathbb{E}[a_n] < c = \frac{d}{d\theta}e^{\theta c}\Big|_{\theta=0} ,$$

(a more rigorous proof can be found in [4]). This means that by a similar argument as in the proof of Lemma 14, θ can be chosen such that

$$\lambda(\theta) = e^{\theta c} . \tag{17}$$

The following martingale construction can be found in [4]:

---

[2]For this figure (and the figures below), 100 independent simulations were run, each consisting of 10^9 packets. To ensure a stationary regime, the first 10^8 packets in each run were discarded. The resulting (empirical) CCDFs are presented as box-plots.

Lemma 16. *In the situation above (i.e., such that Eq. (17) holds), if the function h defined by $h(s_i) = v(i)$ is monotonically increasing, then the flow A admits a (h, θ, c)-martingale-envelope.*

The martingale-envelope constructed in Lemma 16 can be easily extended to Markov chains with continuous state space (see again [4]).

As an application of Lemma 16 consider the arrival model as a *Markov Modulated On-Off Process* (MMOO), i.e., a Markov chain a_n jumping between the two states *On* and *Off* with probabilities α and β, respectively. While in state *On* it transmits R data units per time unit, and while in state *Off* it does not transmit any data. The stationary distribution is given by:

$$\pi_0 := \mathbb{P}(a_n = 0) = \frac{\beta}{\alpha + \beta}, \quad \pi_1 := \mathbb{P}(a_n = R) = \frac{\alpha}{\alpha + \beta} \ ,$$

and that the process is reversible, i.e., $A = A^r$. Further, in [23] it was shown that the eigenfunction h (as defined in Lemma 16) is monotonically increasing if and only if:

$$\text{Cov}[a_n, a_{n+1}] > 0 \Leftrightarrow \alpha < 1 - \beta \ . \tag{18}$$

As an immediate consequence of Theorem 5 we now have:

Corollary 17. *For the MMOO arrival model above and a (per-flow) capacity c satisfying $c > R\pi_1 = \mathbb{E}[a_n]$, it holds for the backlog Q and the virtual delay $W(n)$:*

$$\mathbb{P}(Q \geq \sigma) \leq \kappa e^{-\theta\sigma} \ , \quad and \quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta kC} \ ,$$

where $\kappa := \frac{\alpha + \beta h(0)/h(R)}{\alpha + \beta}$. Moreover, $\kappa < 1$.

We now consider the case of N such queueing systems (A_i, c) being multiplexed. Instead of writing down the transition matrix for the resulting process, we simply can apply Theorem 7 and Lemmas 10 – 13 to obtain bounds on the aggregate and per-flow analysis, respectively:

Corollary 18. *Let*

$$\kappa = \frac{(\pi_0 v_0 + \pi_1 v_1)^N}{v_0^{N - \lceil CR^{-1} \rceil} v_1^{\lceil CR^{-1} \rceil}} \ .$$

Then in the multiplexed queueing system with total capacity $C = Nc$, it holds for the aggregate flow:

$$\mathbb{P}(Q \geq \sigma) \leq \kappa e^{-\theta\sigma} \ , \quad and \quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta Ck} \ ,$$

and for a single flow comprising $N_1 < N$ subflows under scheduling:

FIFO: $\quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta Ck} \qquad\qquad$ SP: $\quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta N_1 ck}$

EDF1: $\quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta(Ck - (N - N_1)c\min\{k,y\})} \qquad$ EDF2: $\quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta(Ck + (N - N_1)cy)}$
$$+ \tilde{\kappa} e^{-\tilde{\theta} Nck} \ ,$$

where $y := d - d'$, and EDF1 and EDF2 correspond to $y \geq 0$ and $y < 0$, respectively. For EDF2, $\tilde{\kappa}$ and $\tilde{\theta}$ denote the corresponding parameters in the queueing system which has the total capacity $C = Nc$ but only the N_1 subflows as arrivals.

It can be shown that the leading constant is exponential in N (see [23]) and thus the fundamental property from Eq. (2) is captured. As a side remark, the corresponding leading constant from [16], p. 340, is greater than one.

We point out that while the bounds in Corollary 18 for the *aggregate flow* have already been obtained in [23], the *per-flow* bounds (i.e., for SP, FIFO, and EDF) represent the contribution of this paper.

In Figure 3 simulations of the MMOO and the corresponding bounds for SP and EDF are displayed for different link utilizations. As in the case of independent increments, the Martingale bounds (from Corollary 18) are tight even at high utilizations (i.e., $\rho = 0.95$), whereas the bounds calculated with Boole's inequality (see Eq. (3)) are off by several orders of magnitude.

As a last remark, the martingale-envelope constructed in this section can be extended to p-order Markov chains by considering the cartesian product S^p as the new state space.
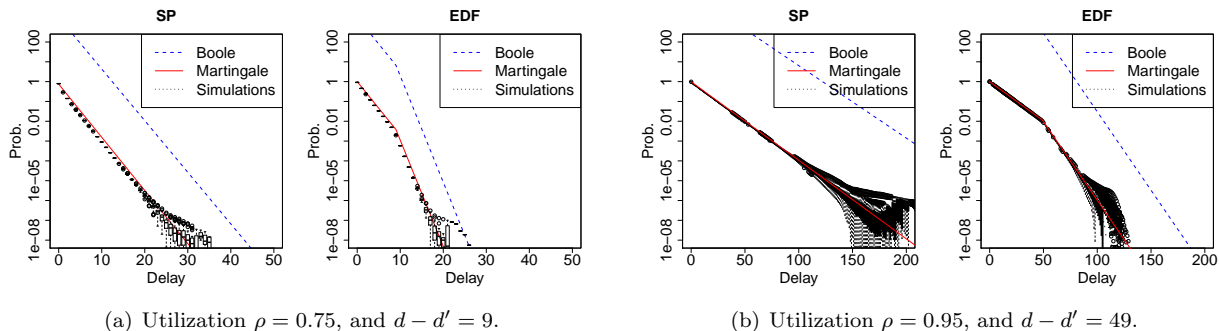
(a) Utilization $\rho = 0.75$, and $d - d' = 9$.

(b) Utilization $\rho = 0.95$, and $d - d' = 49$.

Figure 3: CCDF of the packet delay with $N_1 = \frac{1}{2}N = 10$, $\alpha = 0.1$, $\beta = 0.5$, and $R = 1$.

2.3. Autoregressive Arrival Models

As a third example we consider autoregressive processes. Roughly, a p-order autoregressive process $(AR(p))$ evolves by rescaling the p previous values of the process and adding *Gaussian white noise*, i.e., uncorrelated Gaussian random variables. Although implicitly contained in the theory of (p-order) Markov processes, the different representation has the advantage of providing closed form solutions to the performance metrics.

We start with the formal definition of $AR(p)$. For simplicity, we assume throughout that the white noise is not only uncorrelated but independent.

Definition 19. *Let $p \geq 1$, $Z_0, Z_1, Z_2, \cdots \sim \mathcal{N}_{0,1}$ i.i.d., $\varphi_1, \ldots, \varphi_p \in [0, 1)$, $\varphi = \sum_{k=1}^{p} \varphi_k$, and $\mu, \sigma > 0$. If the relation*

$$a_n = \sum_{k=1}^{p} \varphi_k a_{n-k} + (1 - \varphi)\,\mu + (1 - \varphi)\,\sigma Z_n \tag{19}$$

holds, the process $(a_n)_n$ is called the p-order autoregressive process, $AR(p)$.

It can be shown (see, e.g., [15], p. 85) that if all the (complex) roots of the *characteristic polynomial*

$$\chi(z) = 1 - \sum_{k=1}^{p} \varphi_k z^k$$

lie outside the unit interval, i.e., $\chi(z) = 0 \Rightarrow |z| > 1$, then the process $AR(p)$ is stationary. We assume throughout that this is fulfilled. As above, we apply Kolmogorov's theorem to obtain an extended process $(a_n)_{n \in \mathbb{Z}}$ which is still stationary and satisfying Eq. (19). Moreover, as $AR(p)$ is clearly a Gaussian process itself, it is also reversible (see [24]), i.e., $A^r = A$.

Note that although $\mathbb{E}[a_n] = \mu$ for all $n \in \mathbb{Z}$, by the correlation of $AR(p)$ the variance $\mathbb{V}[a_n]$ is not equal to σ, but must be derived using the Yule-Walker-Equations (see again [15], p. 239).

As in the previous examples we interpret a_n as the instantaneous arrival at time n, i.e.,

$$A_n := \sum_{k=1}^{n} a_k$$

represents the cumulative arrivals up to time n. In the next theorem a martingale-envelope for $AR(p)$ is constructed. As usual, $c > \mu$ denotes the flow's allocated *capacity*.

Theorem 20. *Let $\theta = 2\frac{c - \mu}{\sigma^2}$ and for $1 \leq k \leq p$*

$$h_i(t) := e^{\frac{\theta}{1 - \varphi} \varphi_k t} \ .$$

Then the flow A admits a $(\Pi h, \theta, c)$-martingale-envelope.

10

Note that for $p = 0$ we recover the case of independent increments as in Subsection 2.1.

Clearly, the product function Πh is monotonically increasing in its parameters $\left(\sum_{i=1}^{k} t_i\right)_{0 \leq k \leq p}$. So Lemma 4 can be applied to obtain

$$\Pi h(\overrightarrow{a_N}) \geq \Pi h(c, \dots, pc) = e^{\frac{\theta}{1-\varphi} c \sum_{k=1}^{p} k \varphi_k} , \tag{20}$$

where N denotes the stopping time from Eq. (9).

Let now

$$Y := \sum_{k=1}^{p} \varphi_k \sum_{i=1}^{k} a_{n-i+1}$$

(note that in distribution this is independent of n). Y is normally distributed with $\mathbb{E}[Y] = \mu \sum_{k=1}^{p} k\varphi_k$. Let $\nu^2 := \mathbb{V}[Y]$ denote its variance, which again can be calculated using the Yule-Walker-Equations.

Considering the single flow scenario from Figure 1(a) and Theorem 5, the following bounds hold:

Corollary 21. *For the autoregressive arrival model $AR(p)$ with a capacity c satisfying $c > \mu$, let*

$$\kappa = e^{\frac{\theta(\mu-c)}{1-\varphi}\left(\sum_{k=1}^{p} k\varphi_k - \frac{\nu^2}{(1-\varphi)\sigma^2}\right)} , \quad and \quad \theta = 2\frac{c-\mu}{\sigma^2} .$$

Then for the backlog Q and virtual delay $W(n)$ hold

$$\mathbb{P}(Q \geq \sigma) \leq \kappa e^{-\theta\sigma} , \quad and \quad \mathbb{P}(W(n) \geq k) \leq \kappa e^{-\theta ck} .$$

Let us consider the case of $p = 1$, i.e.,:

$$a_n = \varphi a_{n-1} + (1-\varphi)\mu + (1-\varphi)\sigma Z_n .$$

This special case allows an explicit calculation of the variance ν^2:

$$\nu^2 = \mathbb{V}[\varphi a_n] = \mathbb{V}[\varphi a_{n+1}] = \varphi^2 \mathbb{V}[\varphi a_n + \sigma(1-\varphi)Z_{n+1}] = \varphi^2\left(\nu^2 + \sigma^2(1-\varphi)^2\right) ,$$

and thus $\nu^2 = \sigma^2 \frac{(1-\varphi)\varphi^2}{1+\varphi}$. The leading constant κ from Corollary 21 reduces to

$$\kappa = \frac{\mathbb{E}[h(a_n)]}{h(c)} = e^{\frac{\theta(\mu-c)}{1-\varphi}\left(\varphi - \frac{\nu^2}{(1-\varphi)\sigma^2}\right)} = e^{\frac{\theta(\mu-c)}{1-\varphi}\left(\varphi - \frac{\varphi^2}{1+\varphi}\right)} = e^{\frac{\theta\varphi(\mu-c)}{1-\varphi^2}} . \tag{21}$$

Note that in this case $\kappa \in (0, 1]$. Therefore, with regards to the queue size Q, the following bound holds:

$$\mathbb{P}(Q > \sigma) \leq e^{\frac{\theta\varphi(\mu-c)}{1-\varphi^2}} e^{-\theta\sigma} .$$

This bound improves the known results drastically: e.g., in [16], p. 340, an additional factor occurs, which depends on an upper bound on the increment process. As the *Gaussian white noise* is unbounded, the corresponding bound from [16] is meaningless.

Now consider the second scenario as in Figure 1(b): We assume that two homogeneous and independent autoregressive arrival flows are multiplexed.

Corollary 22. *With the definitions as in Corollary 21 for the multiplexed queueing system with aggregate capacity $2c$ holds:*

$$\mathbb{P}(Q \geq \sigma) \leq \kappa^2 e^{-\theta\sigma} , \quad and \quad \mathbb{P}(W(n) \geq k) \leq \kappa^2 e^{-\theta 2ck} ,$$

and for a single flow under scheduling:

FIFO: $\mathbb{P}(W(n) \geq k) \leq \kappa^2 e^{-\theta 2ck}$ SP: $\mathbb{P}(W(n) \geq k) \leq \kappa^2 e^{-\theta ck}$

EDF1: $\mathbb{P}(W(n) \geq k) \leq \kappa^2 e^{-\theta(2ck - c\min\{k,y\})}$ EDF2: $\mathbb{P}(W(n) \geq k) \leq \kappa^2 e^{-\theta(2ck+cy)} + \tilde{\kappa} e^{-\tilde{\theta}2ck} .$

Again, $y := d - d'$, and EDF1 and EDF2 correspond to $y \geq 0$ and $y < 0$, respectively; $\tilde{\kappa}$ and $\tilde{\theta}$ denote the constants κ and θ with c exchanged by $2c$.
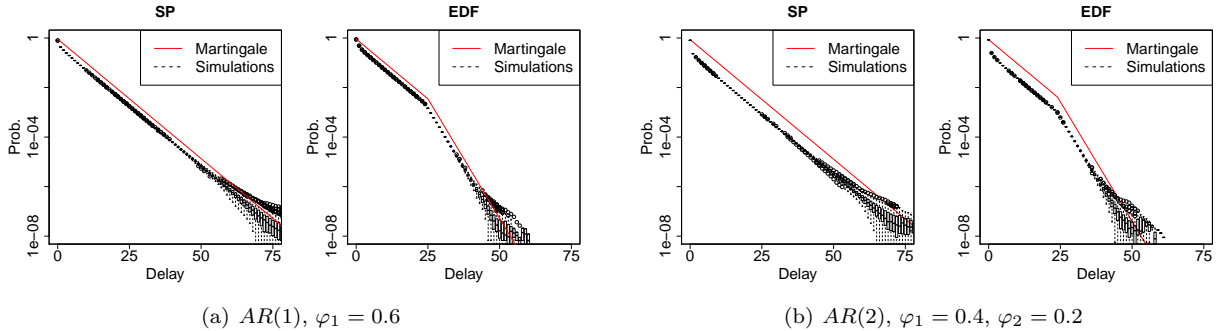
(a) $AR(1)$, $\varphi_1 = 0.6$ (b) $AR(2)$, $\varphi_1 = 0.4$, $\varphi_2 = 0.2$

Figure 4: CCDF of the packet delay for $AR(1)$ ((a)) and $AR(2)$ ((b)), with parameters $\mu = 0.5$, $\sigma = 1$, utilization $\rho = 0.75$, and, for EDF, $y = d - d' = 24$.
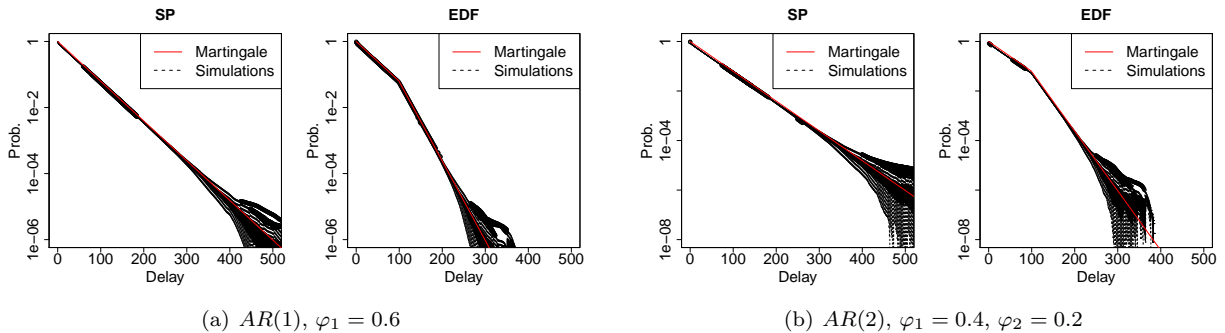

(a) $AR(1)$, $\varphi_1 = 0.6$ (b) $AR(2)$, $\varphi_1 = 0.4$, $\varphi_2 = 0.2$

Figure 5: CCDF of the packet delay for $AR(1)$ ((a)) and $AR(2)$ ((b)), with parameters $\mu = 0.5$, $\sigma = 1$, utilization $\rho = 0.95$, and, for EDF, $y = d - d' = 99$.

Note that as the sum of independent autoregressive processes is still autoregressive, the aggregate bounds in the first part of Corollary 22 could also be obtained by applying Corollary 21 to the *single* flow $A_n + A'_n$. As the corresponding κ is independent of the number of flows, applying Eq. (A.2) iteratively leads to bounds retaining the fundamental exponential decay property from Eq. (2).

In Figures 4 and 5, simulations of the $AR(p)$ and the corresponding bounds for SP and EDF are displayed for different link utilizations. Unlike in the two previous examples, Boole's inequality could not be evoked to obtain bounds, since the sum on the right hand side in Eq. (3) seems not to converge. A follow-up discussion on tightness will be given in the next section.

3. Discussion

In this section we discuss on the tightness issues of the martingale-based method and provide some insight into a contriving idiosyncrasy between discrete vs. continuous-time results. We divide the discussion according to the flows' burstiness level.

3.1. Bursty Flows

Although the bounds illustrated in Figures 2-5 are seemingly accurate, the bounds degrade with the level of correlations within the arrivals. This trend can be particularly noticed for 1-order vs. 2-order autoregressive processes (see Figure 4(a) vs. 4(b)); the same could be observed by reducing the scale of the x-axis in Figures 5(a) and 5(b). One explanation is that on a logarithmic y-axis the simulations throughout are seemingly convex, i.e., the probabilities in an initial phase decay faster than asymptotically (this behavior has been in fact convincingly shown to hold for bursty flows in [3]). In contrast, as the martingale-envelope
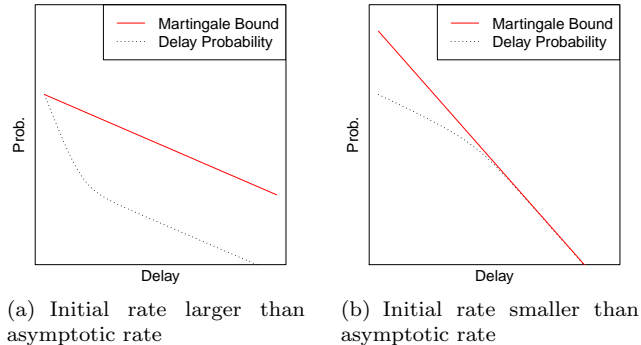
12

(a) Initial rate larger than asymptotic rate

(b) Initial rate smaller than asymptotic rate

Figure 6: Possible CCDF of the delay. Depending on the flows' burstiness the martingale (exponential) bounds are inevitably loose for small or large delays.

is based on an exponential transform, it *can only* render bounds of the form of the (generalized) exponential distribution (i.e., $Prob \leq \kappa e^{-\theta x}$), whence the straight lines in the plots. In other words, the longer the "initial phase" of the true distribution is, or more generally the level of long-range correlations, the larger the gap is between the distribution and the obtained bounds. We point out that the bounds match in fact simulations at the starting point (although difficult to visually perceive from the printed plots) of the true distribution. To more conveniently contrast the true distribution and the *best possible* exponential bounds, see Figure 6(a).

A possible approach to reduce this inherent gap would be to use hyperexponential rather than exponential transforms, i.e., functions of the form $p_1 e^{\theta_1 x} + p_2 e^{\theta_2 x}$, where the parameters p_1, λ_1 and p_2, λ_2 are scaled accordingly to the initial and the tail periods, respectively.

3.2. Less Bursty Flows

We now address the case of Markovian arrivals which are 'less bursty (i.e., *better*) than independent increments'. This characterization is the analogous of 'less bursty than Poisson' (see [3] or [25]) in discrete time.

Let us recall the martingale-envelope from Subsection 2.2, but now assume that the "burstiness"-condition from Eq. (18) is *not* fulfilled. This condition was used to prove that the leading constant κ, in the bounds from Corollary 17, is strictly less than 1. For multiplexed flows, the condition finally implied that the bounds are *exponentially decaying* in the number of flows (see Corollary 18).

The condition fails when $\alpha \geq 1 - \beta$, or, equivalently when $h(0) \geq h(R)$. Let the stopping time N be defined as in Eq. (9). For this specific Markov chain, we have again $h(a_N) = h(R)$, as there are only two states. Hence, proceeding as in the proof of Theorem 5:

$$\mathbb{P}(Q \geq \sigma) \leq \frac{\alpha + \beta \frac{h(0)}{h(R)}}{\alpha + \beta} e^{-\theta \sigma} \ ,$$

but now with a leading constant greater or equal to 1. This indicates that the method developed in this paper does not lead to sharp bounds, because of the now *exponentially increasing* leading constant in the number of flows, for better than Poisson flows.

For such flows, the shape of the performance metrics was observed to be concave, on a logarithmic scale of the y-axis. For a convenient illustration see Figure 6(b), and for a concrete illustration of our bounds see Figure 7 (with a deliberately high utilization). In contrast to the case of bursty flows, whereby the bounds are only tight at the starting point due to the convex shape of the distribution, the bounds for less bursty flows are loose in the initial phase even for a small number of flows. Moreover, due to exponential increase of the leading constant, the bounds eventually blow up in the number of flows, albeit they exactly capture the decay rate θ.
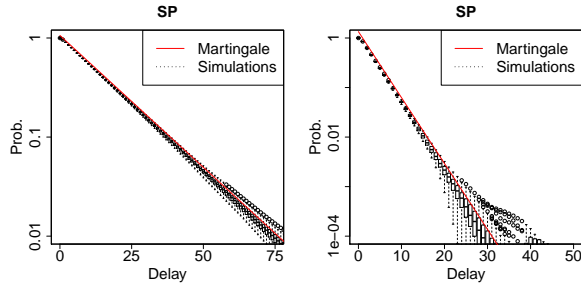
Figure 7: CCDF of the packet delay of with probabilities $\alpha = 0.7$, $\beta = 0.9$, utilization $\rho = 0.999$ and $10 + 10$ (left) and $50 + 50$ (right) through- and cross-flows, respectively.

Let us finally make a connection to the parallel recent results from [12]. Therein, the authors studied a similar queueing system, with the seemingly unimportant difference that they take place in *continuous* rather than *discrete* time. A key finding was that *any* arrival flow driven by a (continuous-time) Markov process with two states admits performance metrics which are decaying exponentially in the number of flows. Let us next explain the roots of this contriving idiosyncracy between continuous and discrete-time models.

The elementary explanation is that while for a continuous-time Markov process $(X_t)_{t \geq 0}$ the sub-process $(X_n)_{n \in \mathbb{N}}$ is a discrete-time Markov chain, the converse does *not* hold in general since not every Markov chain is *embeddable* into a continuous process. To give some background, a discrete Markov chain with a transition matrix P is said to be embeddable if there is a continuous-time Markov process with generator Q such that

$$P = e^Q \ .$$

The problem of embeddability was first addressed in [26], whereas in [6] it was shown that in the very special case of two states, a chain is embeddable if and only if $\det(P) > 0$. Remarkably, this condition is equivalent to the condition $\alpha < 1 - \beta$ from Eq. (18) (recall that this condition was previously established in [4]). Unfortunately, for the general case (i.e., chains with more than two states), no explicit condition is known to the best of our knowledge.

4. Conclusion

In this paper we have proposed a novel characterization of queueing systems by martingale-envelopes and developed a related *unified calculus* dealing with flows' multiplexing and scheduling. The crucial result of this calculus is that the scheduling operation translates into a time shifting operation of the underlying martingale-envelopes, enabling thus the derivation of tight per-flow performance bounds by leveraging a variant of Doob's inequality. We applied this calculus to Markovian and p-order autoregressive arrival flows and derived bounds on the per-flow delay distributions for several scheduling policies (FIFO, SP, and EDF). In certain burstiness scenarios, the obtained per-flow bounds capture for the first-time a fundamental exponential decay factor in the number of flows. Moreover, the bounds almost match simulations, improving over classic results (e.g., FIFO: [7, 16], SP: [8, 9], EDF: [10]) by arbitrary orders of magnitude, especially at extreme utilizations.

However, as the martingale-envelope is based on an exponential transform, the bounds' accuracy can degrade in situations where the true distribution diverges from the exponential. The development of more contrived martingale-envelopes encapturing this behavior, together with the analysis of other transformations of queueing systems (e.g., of a queueing network), remain open issues for further research in this area.

14

References

[1] F. P. Kelly, Notes on effective bandwidths, in: Stochastic Networks: Theory and Applications. (Editors: F.P. Kelly, S. Zachary and I.B. Ziedins) Royal Statistical Society Lecture Notes Series, 4, Oxford University Press, 1996, pp. 141–168.

[2] N. B. Shroff, M. Schwartz, Improved loss calculations at an ATM multiplexer, IEEE/ACM Transactions on Networking 6 (4) (1998) 411–421.

[3] G. Choudhury, D. Lucantoni, W. Whitt, Squeezing the most out of ATM, IEEE Transactions on Communications 44 (2) (1996) 203–217.

[4] N. G. Duffield, Exponential bounds for queues with Markovian arrivals, Queueing Systems 17 (3-4) (1994) 413–430.

[5] M. Talagrand, Majorizing measures: The generic chaining, Annals of Probability 24 (3) (1996) 1049–1103. URL http://www.math.ohio-state.edu/~talagran/preprints/majmeas.dvi

[6] J. F. C. Kingman, The imbedding problem for finite markov chains, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 1 (1) (1962) 14–24.

[7] C. Courcoubetis, R. Weber, Effective bandwidths for stationary sources, Probability in Engineering and Informational Sciences 9 (2) (1995) 285–294.

[8] A. W. Berger, W. Whitt, Effective bandwidths with priorities, IEEE/ACM Transactions on Networking 6 (4) (1998) 447–460.

[9] D. Wischik, Sample path large deviations for queues with many inputs, Annals of Applied Probability 11 (2) (2001) 379–404.

[10] V. Sivaraman, F. M. Chiussi, Statistical analysis of delay bound violations at an earliest deadline first scheduler, Performance Evaluation 36 (1) (1999) 457–470.

[11] D. Bertsimas, I. Paschalidis, J. Tsitsiklis, Large deviations analysis of the generalized processor sharing policy, Queueing Systems 32 (4) (1999) 319–349. doi:10.1023/A:1019151423773.

[12] F. Ciucu, F. Poloczek, J. Schmitt, Sharp bounds in stochastic network calculus, CoRR abs/1303.4114.

[13] S. N. Ethier, T. G. Kurtz, Markov processes – characterization and convergence, John Wiley & Sons Inc., 1986.

[14] F. Ciucu, J. Schmitt, Perspectives on network calculus: no free lunch, but still good value, in: ACM Sigcomm, 2012, pp. 311–322.

[15] P. J. Brockwell, R. A. Davis, Time Series: Theory and Methods, Springer-Verlag, 1991.

[16] C.-S. Chang, Performance Guarantees in Communication Networks, Springer Verlag, 2000.

[17] R. L. Cruz, C. Okino, Service guarantees for window flow control, in: 34th Allerton Conference on Communications, Control and Computating, 1996.

[18] J.-Y. Le Boudec, P. Thiran, Network Calculus, Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.

[19] M. Fidler, An end-to-end probabilistic network calculus with moment generating functions, in: IEEE International Workshop on Quality of Service (IWQoS), 2006, pp. 261–270.

[20] R. L. Cruz, SCED+: Efficient management of quality of service guarantees, in: IEEE Infocom, 1998, pp. 625–634.

[21] J. Liebeherr, Y. Ghiassi-Farrokhfal, A. Burchard, On the impact of link scheduling on end-to-end delays in large networks, IEEE Journal on Selected Areas in Communications 29 (5) (2011) 1009–1020. doi:10.1109/JSAC.2011.110511.

[22] I. Iscoe, P. Ney, E. Nummelin, Large deviations of uniformly recurrent markov additive processes, Advances in Applied Mathematics 6 (4) (1985) 373–412. doi:10.1016/0196-8858(85)90017-X. URL http://dx.doi.org/10.1016/0196-8858(85)90017-X

[23] E. Buffet, N. G. Duffield, Exponential upper bounds via martingales for multiplexers with Markovian arrivals, Journal of Applied Probability 31 (4) (1994) 1049–1060.

[24] G. Weiss, Time-reversibility of linear stochastic processes, Journal of Applied Probability 12 (4) (1975) 831–836.

[25] T. Bonald, A. Proutiere, J. Roberts, Statistical performance guarantees for streaming flows using expedited forwarding, in: IEEE Infocom, 2001, pp. 1104–1112.

[26] G. Elfving, Zur Theorie der Markoffschen Ketten, Acta Societatis Scientiarum Fennicae 2 (8).

[27] P. Billingsley, Probability and Measure (3^{rd} Edition), Wiley, 1995.

[28] A. Cherny, Some particular problems of martingale theory, in: From Stochastic Calculus to Mathematical Finance, Springer, 2006, pp. 109–124.

AppendixA. Proofs

Proof of Lemma 4. Assume that $\sum_{i=0}^{k-1} a_{N-i} \leq kc$ for some $k \geq 1$. Then

$$A_{N-k} - (N-k)c = (A_N - Nc) - \left(\sum_{i=0}^{k-1} a_{N-i} - kc \right)$$

$$\geq \sigma ,$$

contradicting the minimal property of N. □

Proof of Theorem 5. Let the stopping time N as in Eq. (9). By definition, $\mathbb{P}(Q \geq \sigma) = \mathbb{P}(N < \infty)$. Applying the optional sampling theorem (see [27]) to the supermartingale from Eq. (8) (with $m = 0$) yields for every $k \in \mathbb{N}$:

$$\mathbb{E}[\Pi h(\overrightarrow{a_n})] = \mathbb{E}[M_0(0)] \geq \mathbb{E}[M_0(N \wedge k)] \geq \mathbb{E}[M_0(N \wedge k)1_{N<k}]$$
$$= \mathbb{E}[\Pi h(\overrightarrow{a_N})e^{\theta(A^r(0,N)-cN)}1_{N<k}]$$
$$\geq \Pi h(c, 2c, \ldots, pc)e^{\theta\sigma}\mathbb{P}(N < k) ,$$

where we used the monotonicity of Πh and Lemma 4. Now simply let $k \to \infty$.

For the virtual delay, first note that:

$$\mathbb{P}(W(n) \geq k) = \mathbb{P}(A(0, n-k) \geq D(n))$$
$$\leq \mathbb{P}(A(0, n-k) \geq (A * S)(n))$$
$$= \mathbb{P}(A(0, n-k) \geq \inf_{0 \leq l \leq n} \{A(0,l) - C(n-l)\})$$
$$\leq \mathbb{P}(\sup_{0 \leq l \leq n-k} \{A(l, n-k) - C(n-l)\} \geq 0)$$
$$= \mathbb{P}(\sup_{l \geq 0}\{A^r(0,l) - Cl\} \geq Ck) ,$$

where we used the service process from Eq. (6) in the first and the monotonicity of A in the second inequality; in the final line, by stationarity, the time was shifted to $n - k$.

As a final step, let N as in Eq. (9) with $\sigma = Ck$, and proceed as for the backlog. $\square$

Proof of Theorem 7. Let a_n and a'_n denote the respective increment processes. Clearly, Eq. (10) implies for all n:

$$(\Pi h \otimes h')(\overrightarrow{a_n} + \overrightarrow{a_n}')e^{\theta(A(m,n)+A'(m,n)-C(n-m))}$$
$$\leq \Pi h(\overrightarrow{a_n})e^{\theta(A(m,n)-c(n-m))}\Pi h'(\overrightarrow{a_n}')e^{\theta(A'(m,n)-c'(n-m))} ,$$

i.e., the product of the two martingale envelopes. By the independence assumption, this product is a martingale as well (see [28]), and the proof is done. $\square$

Proof of Lemma 8. The function $\Phi(t) = t^{\frac{\theta'}{\theta}}$ is concave and monotonically increasing in $t \geq 0$. Now

$$\mathbb{E}\left[\Pi h^{\frac{\theta'}{\theta}}(\overrightarrow{a_{n+1}})e^{\theta'(A^r(m,n+1)-c(n-m-1))} \mid \mathcal{F}_n\right]$$
$$= \mathbb{E}\left[\Phi\left(\Pi h(\overrightarrow{a_{n+1}})e^{\theta(A^r(m,n+1)-c(n-m-1))}\right) \mid \mathcal{F}_n\right]$$
$$\leq \Phi\left(\mathbb{E}\left[\Pi h(\overrightarrow{a_{n+1}})e^{\theta(A^r(m,n+1)-c(n-m-1))} \mid \mathcal{F}_n\right]\right)$$
$$\leq \Phi\left(\Pi h(\overrightarrow{a_n})e^{\theta(A^r(m,n)-c(n-m))}\right)$$
$$= \Pi h^{\frac{\theta'}{\theta}}(\overrightarrow{a_n})e^{\theta'(A^r(m,n)-c(n-m))} ,$$

using Jensen's inequality for conditional expectations (see [27]) in the third line. $\square$

Proof of Lemma 9. We proceed similarly as in the proof of Theorem 7. Consider the two supermartingales

$$M(n) = \Pi h(\overrightarrow{a_n})e^{\theta(A^r(l,n)-c(n-l))} , \qquad\qquad n \geq l , \text{ and}$$
$$M'(n) = \Pi h(\overrightarrow{a_n}')e^{\theta\left(A'^r(0,n)-c'n\right)} , \qquad\qquad n \geq 0$$

from the definition of the envelopes. By the independence assumption, the process

$$\tilde{M}(n) = M(n)M'(n)$$

16

is a supermartingale in the *time-shifted* domain $\{l, l+1, l+2, \dots\}$. Let N denote a stopping time similar to the one from Eq. (9):

$$N := \inf\{n \geq l \mid A^r(l, n) + A^r(0, n) - Cn \geq cl + \sigma\} .\tag{A.1}$$

Again, the desired probability is equal to $\mathbb{P}(N < \infty)$. By applying the optional stopping theorem, one has for $k \geq l$:

$$
\begin{aligned}
\mathbb{E}[\tilde{M}(l)] &\geq \mathbb{E}[\tilde{M}(N \wedge k)]\\
&\geq \mathbb{E}[\Pi h(\overrightarrow{a_n})\Pi h'(\overrightarrow{a_n}')e^{\theta\left(A^r(l,n)+A'^r(0,n)-Cn+cl\right)}1_{N<k}]\\
&\geq \Pi h(c, 2c, \dots, pc)\Pi h'(c, 2c, \dots, pc)e^{\theta(\sigma+cl)}\mathbb{P}(N<k)\\
&\geq \Pi h \otimes h'(c, 2c, \dots, pc)e^{\theta(\sigma+cl)}\mathbb{P}(N<k) .
\end{aligned}
$$

Now, by independence and the supermartingale property of M':

$$
\begin{aligned}
\mathbb{E}[\tilde{M}(l)] &= \mathbb{E}[\Pi h(\overrightarrow{a_n})M'(l)] = \mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[M'(l)]\\
&\leq \mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[M'(0)] = \mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a_n'})] .
\end{aligned}
$$

As above, we finally let $N \to \infty$ to complete the proof. $\qquad\square$

Proof of Corollary 10. In continuation of Eq. (11) with the service process as in Eq. (12):

$$
\begin{aligned}
\mathbb{P}(W(n) \geq k) &\leq \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) - S^r(0, n)\} \geq 0)\\
&= \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) - [Cn - A'^r(0, n)]_+\} \geq 0)\\
&\leq \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) + A'^r(0, n) - Cn \geq 0\}) .
\end{aligned}
$$

Now simply plug in the parameters $\sigma = 0$ and $l = k$ into Lemma 9. $\qquad\square$

Proof of Corollary 11. For the free parameter choose $x = k$. Then Eq. (11), with the service process from Eq. (13), continues to:

$$
\begin{aligned}
\mathbb{P}(W(n) \geq k) &\leq \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) - S^r(0, n)\} \geq 0)\\
&= \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) - [Cn - A'^r(0, n-k)]_+ 1_{\{n>k\}}\} \geq 0)\\
&\leq \mathbb{P}(\sup_{n \geq 0}\{A^r(0, n) + A'^r(0, n) - C(n+k)\} \geq 0) .
\end{aligned}
$$

Now apply Lemma 9 with $l = 0$ and $\sigma = Ck$. $\qquad\square$

Proof of Corollary 12. Again, let $x := k$. Eq. (11) with the service process from Eq. (14) gives:

$$
\begin{aligned}
&\mathbb{P}(W(n) \geq k)\\
&\leq \mathbb{P}(\sup_{n \geq k}\{A^r(k, n) + A'^r(0, n-k+\min\{k, y\}) - Cn\} \geq 0)\\
&\leq \mathbb{P}(\sup_{\tilde{n} \geq \min\{k,y\}}\{A^r(k, \tilde{n}+k-\min\{k, y\}) + A'^r(0, \tilde{n})\\
&\qquad\qquad\qquad - C(\tilde{n}+k-\min\{k, y\}\} \geq 0)\\
&\leq \mathbb{P}(\sup_{\tilde{n} \geq \min\{k,y\}}\{A^r(\min\{k, y\}, \tilde{n}) + A'^r(0, \tilde{n}) - C\tilde{n}\}\\
&\qquad\qquad\qquad \geq C(k - \min\{k, y\})) ,
\end{aligned}
$$

17

where we used the substitution

$$\tilde{n} = n - k + \min\{k, y\}$$

in the third, and the stationarity of A^r in the fourth line. Now apply Lemma 9 with $l = \min\{k, y\}$, and $\sigma = C(k - \min\{k, y\})$; hereby note that $l \geq 0$ and $\sigma - cl = Ck - c'\min\{k, y\}$. $\qquad\square$

Proof of Corollary 13. By splitting up the probability in Eq. (11) using the union bound

$$\mathbb{P}(W(n) \geq k) \leq \mathbb{P}(\sup_{n \geq k: n \notin B} \{A^r(k, n) - S^r(0, n)\} \geq 0)$$
$$+ \mathbb{P}(\sup_{n \geq k: n \in B} \{A^r(k, n) - S^r(0, n)\} \geq 0) ,$$

one has for the first probability:

$$\mathbb{P}(\sup_{n \geq k: n \notin B} \{A^r(k, n) - S^r(0, n)\} \geq 0)$$
$$\leq \mathbb{P}(\sup_{n \geq k-y} \{A^r(k, n) + A'^r(0, n - k + y) - Cn\} \geq 0)$$
$$\leq \mathbb{P}(\sup_{n \geq -y} \{A^r(0, \tilde{n}) + A'^r(-y, \tilde{n}) - C\tilde{n}\} \geq Ck)$$
$$\leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]\mathbb{E}[\Pi h(\overrightarrow{a'_n})]}{\Pi h \otimes h'(c, 2c, \ldots, pc)} e^{-\theta(Ck + c'y)} .$$

In the third line, stationarity and the substitution $\tilde{n} = n + k$ was used, and in the fourth line Lemma 9 was applied with $\sigma = Ck$, $l = -y$, and the roles of A and A' were interchanged.

For the second probability with Eq. (15):

$$\mathbb{P}(\sup_{n \geq k: n \in B} \{A^r(k, n) - S^r(0, n)\} \geq 0)$$
$$\leq \mathbb{P}(\sup_{k \leq n < k-y} \{A^r(k, n) - Cn\} \geq 0)$$
$$= \mathbb{P}(\sup_{0 \leq \tilde{n} < -y} \{A^r(0, \tilde{n}) - C(\tilde{n} + k)\} \geq 0)$$
$$\leq \mathbb{P}(\sup_{\tilde{n} \geq 0} \{A^r(0, \tilde{n}) - C\tilde{n}\} \geq Ck)$$
$$\leq \frac{\mathbb{E}[\Pi h(\overrightarrow{a_n})]}{\Pi h(c, 2c, \ldots, pc)} e^{-\tilde{\theta}Ck} ,$$

with the usual substitution $\tilde{n} = n - k$ and the stationarity assumption in the fourth line. In the last line Eq. (7) and Theorem 5 with $\sigma = Ck$ were used. $\qquad\square$

Proof of Lemma 14. Consider the two continuous functions

$$\varphi_1(\theta) := \mathbb{E}[e^{\theta a_1}] \text{ and } \varphi_2(\theta) := e^{\theta c} .$$

Due to the first stability condition from Eq. (16) we know that

$$\frac{d}{d\theta}\varphi_1(\theta)\bigg|_{\theta=0} = \mathbb{E}[a_1] < c = \frac{d}{d\theta}\varphi_2(\theta)\bigg|_{\theta=0} ,$$

i.e., (since $\varphi_1(0) = \varphi_2(0) = 1$) there is $\varepsilon > 0$ such that $\varphi_1 < \varphi_2$ on $[0, \varepsilon]$. Due to the second stability condition, φ_1 will eventually become larger than φ_2, and so by continuity there exists $\theta_0 > 0$ such that $\varphi_1(\theta_0) = \varphi_2(\theta_0)$. Now,

$$\mathbb{E}[e^{\theta_0(A(n+1)-(n+1)c)} \mid a_1, \ldots, a_n] = e^{\theta_0(A(n)-nc)}\mathbb{E}[e^{\theta_0 a_1}]e^{-\theta_0 c}$$
$$= e^{\theta_0(A(n)-nc)} ,$$

and so $e^{\theta_0(A_n - nc)}$ is a martingale. $\qquad\square$

Proof of Corollary 15. Use the martingale-envelope from Lemma 14. For the first part, apply Theorem 5. For the second apply Corollaries 10 – 13 using the transform from Lemma 8. □

Proof of Lemma 16. Given $M(n)$ as in Eq. (8), we have to show that it is a martingale. Note that:

$$\mathbb{E}[h(a_{n+1})e^{\theta a_{n+1}} \mid a_n = s_i] = \sum_{j=1}^m h(s_j)e^{\theta s_j}T(i,j)$$
$$= \sum_{j=1}^m T_\theta(i,j)v(j) ,$$

so that

$$\mathbb{E}[h(a_{n+1})e^{\theta(a_{n+1}-c)} \mid \mathcal{F}_n] = e^{-\theta c}T_\theta v$$
$$= e^{-\theta c}\lambda(\theta)h(a_n) = h(a_n) .$$

Now simply multiply both sides with $e^{\theta(A(m,n)-(n-m)c)}$. □

Proof of Corollary 17. With the remark below Lemma 4,

$$\inf\{x > c \mid \mathbb{P}(a_n \in [x,\infty) > 0)\} = R .$$

Now apply Theorem 5 to the martingale-envelope constructed in Lemma 16. □

Proof of Corollary 18. By the remark below Lemma 4 we can use $\tau = R\lceil CR^{-1}\rceil$ as the smallest possible value such that the aggregate input is larger than the capacity. Clearly:

$$h^{\otimes N}(\tau_{\sum_{i=1}^N A_i,C}) = v_0^{N-\lceil CR^{-1}\rceil}v_1^{\lceil CR^{-1}\rceil} .$$

Now simply apply Theorem 7 and Lemmas 10 – 13 to the envelope of Lemma 16. □

Proof of Theorem 20. Note first that

$$\Pi h(\overrightarrow{a_n}) = e^{\frac{\theta}{1-\varphi}\sum_{k=1}^p \varphi_k \sum_{i=1}^k a_{n-i+1}} .$$

For $n \geq 0$, let $\mathcal{F}_n := \sigma\{Z_m \mid m \leq n\}$ and

$$M_n := h(a_n,\ldots,a_{n-p+1})e^{\theta(A_n-nc)} .$$

We show that M_n is a martingale w.r.t. $\mathcal{F}_n$. Note that

$$\mathbb{E}[h(a_n,\ldots,a_{n-p+1})e^{\theta(a_n-c)} \mid \mathcal{F}_{n-1}]$$
$$=\mathbb{E}[e^{\theta\left(\frac{1}{1-\varphi}\left(\sum_{k=1}^p \varphi_k \sum_{i=1}^k a_{n-i+1}\right)+a_n-c\right)} \mid \mathcal{F}_{n-1}]$$
$$=\mathbb{E}[e^{\theta\left(\frac{1}{1-\varphi}\left(\sum_{k=1}^p \varphi_k \sum_{i=2}^k a_{n-i+1}\right)+\frac{\varphi}{1-\varphi}a_n+a_n-c\right)} \mid \mathcal{F}_{n-1}]$$
$$=\mathbb{E}[e^{\theta\left(\frac{1}{1-\varphi}\left(\sum_{k=1}^p \varphi_k \sum_{i=1}^{k-1} a_{n-i}\right)+\sum_{k=1}^p \frac{\varphi_k}{1-\varphi}a_{n-k}+\mu+\sigma Z_n-c\right)} \mid \mathcal{F}_{n-1}]$$
$$=\mathbb{E}[e^{\theta\left(\frac{1}{1-\varphi}\left(\sum_{k=1}^p \varphi_k \sum_{i=1}^k a_{n-i}\right)+\mu+\sigma Z_n-c\right)} \mid \mathcal{F}_{n-1}]$$
$$=h(a_{n-1},\ldots,a_{n-p})\mathbb{E}[e^{\theta(\mu+\sigma Z_n-c)}]$$
$$=h(a_{n-1},\ldots,a_{n-p}) .$$

Multiplying both sides by $e^{\theta(A_{n-1}-(n-1)c)}$ yields

$$\mathbb{E}[M_n \mid \mathcal{F}_{n-1}] = M_{n-1}$$

and the proof is complete. □

19

Proof of Corollary 21. Evaluating the leading constant from Theorem 5 with the given parameters yields:

$$\frac{\mathbb{E}[h(a_n,\ldots,a_{n-p+1})]}{h(c,\ldots,pc)} = \frac{\mathbb{E}[e^{\frac{\theta}{1-\varphi}\left(\mu\sum_{k=1}^{p}k\varphi_k+\nu Z_0\right)}]}{e^{\frac{\theta}{1-\varphi}\left(\sum_{k=1}^{p}\varphi_k\sum_{i=1}^{k}a_{n-i+1}\right)}} = \frac{e^{\frac{\theta}{1-\varphi}\left(\mu\sum_{k=1}^{p}k\varphi_k+\frac{\theta\nu^2}{(1-\varphi)2}\right)}}{e^{\frac{\theta}{1-\varphi}\sum_{k=1}^{p}\varphi_k kc}}$$

$$= e^{\frac{\theta(\mu-c)}{1-\varphi}\left(\sum_{k=1}^{p}k\varphi_k-\frac{\nu^2}{(1-\varphi)\sigma^2}\right)} = \kappa \ .$$

Now simply apply Theorem 5 and Theorem 20. $\qquad\square$

Proof of Corollary 22. By definition of h_i in Theorem 20:

$$h_i \otimes h_i(t) = h_i(t)^2 \ . \tag{A.2}$$

Now simply apply Theorem 7 and Corollaries 10 – 13 to the envelope of Theorem 20. $\qquad\square$