



HAL
open science

Network-Calculus Service Curves of the Interleaved Regulator

Ludovic Thomas, Jean-Yves Le Boudec

► **To cite this version:**

Ludovic Thomas, Jean-Yves Le Boudec. Network-Calculus Service Curves of the Interleaved Regulator. International Teletraffic Congress (ITC 35), Turin, Italy, Oct 2023, Turin (IT), Italy. 10.48550/ARXIV.2305.18036 . hal-04423400

HAL Id: hal-04423400

<https://hal.science/hal-04423400v1>

Submitted on 29 Jan 2024

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

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

Copyright

Network-Calculus Service Curves of the Interleaved Regulator

Ludovic Thomas

École Polytechnique Fédérale de Lausanne, Switzerland
ludovic.thomas@epfl.ch

Jean-Yves Le Boudec

École Polytechnique Fédérale de Lausanne, Switzerland
jean-yves.leboudec@epfl.ch

Abstract—The interleaved regulator (implemented by IEEE TSN Asynchronous Traffic Shaping) is used in time-sensitive networks for reshaping the flows with per-flow contracts. When applied to an aggregate of flows that come from a FIFO system, an interleaved regulator that reshapes the flows with their initial contracts does not increase the worst-case delay of the aggregate. This shaping-for-free property supports the computation of end-to-end latency bounds and the validation of the network’s timing requirements. A common method to establish the properties of a network element is to obtain a network-calculus service-curve model. The existence of such a model for the interleaved regulator remains an open question. If a service-curve model were found for the interleaved regulator, then the analysis of this mechanism would no longer be limited to the situations where the shaping-for-free holds, which would widen its use in time-sensitive networks. In this paper, we investigate if network-calculus service curves can capture the behavior of the interleaved regulator. We find that an interleaved regulator placed outside of the shaping-for-free requirements (after a non-FIFO system) can yield unbounded latencies. Consequently, we prove that no network-calculus service curve exists to explain the interleaved regulator’s behavior. It is still possible to find non-trivial service curves for the interleaved regulator. However, their long-term rate cannot be large enough to provide any guarantee (specifically, we prove that for the regulators that process at least four flows with the same contract, the long-term rate of any service curve is upper bounded by three times the rate of the per-flow contract).

Index Terms—Network Calculus, Service Curve, Interleaved Regulator (IR), Time-Sensitive Networking (TSN), Asynchronous Traffic Shaping (ATS)

I. INTRODUCTION

Time-sensitive networks, as specified by the time-sensitive networking (TSN) task group of the Institute of Electrical and Electronics Engineers (IEEE), support safety-critical applications in the aerospace, automation, and automotive domains [1]. To do so, time-sensitive networks provide a deterministic service with guaranteed bounded latencies.

These guarantees must be validated through a deterministic worst-case timing analysis that can be performed with network calculus. This mathematical framework obtains worst-case performance bounds by modeling the flows with the concept of arrival curves and the network elements with the concept of service curves. Service curves constrain the minimum amount of service that network elements provide to a flow or aggregate of flows.

Time-sensitive networks can also rely on a set of mechanisms that improve the traditional forwarding process of an

output port. The traffic regulators are such hardware elements that support higher scalability and efficiency of time-sensitive networks. Placed after a multiplexing stage, they reshape the flows with per-flow shaping curves (by delaying packets if required) and remove the increase of the flows’ burstiness due to their interference with other flows.

Traffic regulators come in two flavors: per-flow regulators (PFRs) and interleaved regulators (IRs). A PFR processes a unique flow. It stores the packets of the flow in a first in, first out (FIFO) queue and releases the head-of-line packet as soon as doing so does not violate the configured shaping curve for the flow. In contrast, the IR processes an aggregate of flows with a unique FIFO queue. Each flow has its own configured shaping curve, but the IR analyses only the head-of-line packet and releases it as soon as doing so does not violate the shaping curve of the associated flow. The packets in the IR queue are blocked by the head-of-line even if they belong to other flows. This second flavor is implemented within IEEE TSN under the name asynchronous traffic shaping (ATS) [2].

In time-sensitive networks that contain traffic regulators, end-to-end latency bounds are obtained from the knowledge of the shaping curves enforced by the regulators and from the essential “shaping-for-free” property. It states that the traffic regulators do not increase the worst-case latency of the flow (or of the flow aggregate) under certain conditions that depend on the type of the regulator (PFR or IR). Most analyses of traffic regulators rely on this property.

For the PFR, the shaping-for-free property is well understood because a PFR with a concave shaping curve can be modeled with a context-agnostic service curve, *i.e.*, a service curve that only depends on the configuration of the PFR but not on the context in which the PFR is placed. This service curve proves the shaping-for-free property when the PFR is placed in the appropriate context. When the PFR’s context deviates from the shaping-for-free requirements, the context-agnostic service curve still provides performance bounds for the PFR, and slight deviations of the context lead to bounded delay penalties. On the contrary, the only context-agnostic service curve known for the IR is the trivial function $t \mapsto 0$. The non-trivial service-curve models published in the literature [3], [4] are context dependent and always assume that the shaping-for-free property holds. Without a context-agnostic service-curve model, performance bounds cannot be obtained for an IR placed outside the shaping-for-free requirements,

which restrains its use in time-sensitive networks.

In this paper, we investigate if the behavior of the IR can be modeled by a context-agnostic network-calculus service curve. Our contributions are:

- As opposed to the shaping-for-free property when the IR is placed after a FIFO system, we prove that the IR can yield unbounded latencies when placed after a non-FIFO system, even if the latter is FIFO-per-flow and lossless.
- We prove that the shaping-for-free property of the IR cannot be explained by any network-calculus service-curve model.
- For any IR that processes at least four flows, we prove that any context-agnostic service curve for an individual flow is upper bounded by a constant.
- We exhibit a strict service curve of the IR and a function that upper bounds any other context-agnostic strict service curve.
- For any IR that processes at least four flows, we show that the long-term rate of any context-agnostic service curve for the aggregate is upper bounded by three times the rate of the per-flow contract.

The paper is organized as follows. We provide the background on network-calculus service curves and regulators in Section II. We discuss the related work in Section III and provide the system model in Section IV. We then analyze the role of the FIFO assumption in the shaping-for-free property of the IR in Section V. Afterward, we discuss the context-agnostic service curves of the IR in Section VI. We provide our conclusive remarks in Section VII.

II. BACKGROUND

In time-sensitive networks, performance metrics such as flows' end-to-end latencies have to be bounded in the worst case, not in average. The network-calculus framework [5], [6], [7] can provide such performance bounds. It describes the data traffic with cumulative functions, such as R^A , where $R^A(t)$ is the amount of data that cross the observation point A between an arbitrary time reference 0 and t . Cumulative functions belong to $\mathfrak{F}_0 = \{f : \mathbb{R}^+ \rightarrow \mathbb{R}^+ | f(0) = 0\}$.

A. Network-Calculus Service Curves

A causal network system S offers a service curve β if (a) β is wide-sense increasing and (b) for any input cumulative function $R^A(t)$, the resulting output traffic $R^B(t)$ verifies

$$\forall t \geq 0, \quad R^B(t) \geq (R^A \otimes \beta)(t) \quad (1)$$

where \otimes describes the min-plus convolution (Table I). Common service curves are of the form *rate latency* $\beta_{R,T} : t \mapsto R \cdot [t-T]^+$ with rate R and latency T , where $[\cdot]^+ = \max(\cdot, 0)$.

Some network systems provide stronger guarantees through a *strict* service curve. A causal network system S offers a strict service curve β^{strict} if (a) β^{strict} is wide-sense increasing and (b) during any interval $]s, t]$ in which the system is never empty (a so-called *backlogged period*), the output R^B verifies

$$\forall t \geq s \geq 0, \quad R^B(t) - R^B(s) \geq \beta^{\text{strict}}(t - s) \quad (2)$$

Such a system then also offers β^{strict} as a service curve: β^{strict} also verifies Inequation (1) [5, Prop. 1.3.5]. We say that




Fig. 1. Notations of Definition 1. Z offers the *fluid service curve* β if it can be modelled as the concatenation of Z' followed by a packetizer, where Z' offers the service curve β .

a regulator offers a *context-agnostic* service curve [resp., strict service curve] if (1) [resp., (2)] holds for any packetized input $R^A(t)$, without any other assumption on the upstream systems.

Reciprocally, a wide-sense increasing function $\alpha \in \mathfrak{F}_0$ is an arrival curve for the traffic at A if

$$\forall t \geq s \geq 0, \quad R^A(t) - R^A(s) \leq \alpha(t - s) \quad (3)$$

We also note $R^A \sim \alpha$ and say that the traffic is α -constrained. Common arrival curves are of the form *leaky bucket* $\gamma_{r,b}$ with a rate r and a burst b : $\forall t > 0, \gamma_{r,b}(t) = rt + b$.

Given some arrival-curve and service-curve constraints, network-calculus results provide delay and backlog bounds at a network element. A common approach for computing end-to-end performance bounds in time-sensitive networks consists in obtaining an arrival-curve model for each flow and a service-curve model for each network element. Service-curve models for most IEEE TSN mechanisms can be found in [8], [9].

B. The Packetizer and Fluid Service Curves

In packet-switching time-sensitive networks, the stream of data at an observation point A can either be fluid (*e.g.*, on the transmission links) or packetized (packet-by-packet, *e.g.*, within the switches). A packetizer transforms a fluid stream into a packetized stream by releasing the packet's bits only when the last bit is received. It does not increase the end-to-end latency bounds [5, Thm. 1.7.5].

When a system Z with packetized input and output can be split into a fluid service-curve element followed by a packetizer, we say that Z offers a *fluid* service curve.

Definition 1 (Fluid service curve). Consider a function $\beta \in \mathfrak{F}_0$ and a system Z with packetized input and output. We say that Z offers β as a *fluid service curve* if there exists a system Z' that offers the service curve β such that Z can be realized as the concatenation of Z' , followed by a packetizer (Figure 1).

C. Individual Service Curve for a Flow

In time-sensitive networks, the service modeled by the service curves of Sections II-A and II-B is shared between the flows of the aggregate \mathcal{F} . In (1), the cumulative arrival function R^A of the aggregate at A is the sum of the individual arrival functions $\{R_f^A\}_{f \in \mathcal{F}}$ for each flow f in the aggregate \mathcal{F} : $\forall t \in \mathbb{R}^+, R^A(t) = \sum_{f \in \mathcal{F}} R_f^A(t)$.

We say that a system S offers to flow g the *individual* service curve β_g if (a) β_g is wide-sense increasing, and (b) for any cumulative function R_g^A of the flow g at the input A of S , the cumulative function R_g^B of g at its output B verifies

$$\forall t \geq 0, \quad R_g^B(t) \geq (R_g^A \otimes \beta_g)(t) \quad (4)$$




Fig. 2. Two flavors of traffic regulators. With PFRs, we need one PFR per flow. In contrast, the IR uses a single FIFO queue to shape several flows.

D. Traffic Regulators and their Shaping-For-Free Properties

Traffic regulators are hardware elements placed before a multiplexing stage to remove the increased burstiness due to interference with other flows in previous hops. They enable the computation of guaranteed latency bounds in networks with cyclic dependencies [3], [10], [11]. They come in two flavors.

A per-flow regulator (PFR) is a causal, lossless, FIFO system configured for a unique flow f with a shaping curve σ_f (Figure 2a). It stores the packets of f in order of arrival and releases the head-of-line (HOL) packet at the earliest time such that the resulting output has σ_f as an arrival curve. In a network with multiple flows, there is one PFR per flow.

The interleaved regulator (IR) is a causal, lossless, and FIFO system that processes an aggregate $\mathcal{F} = \{f_1, f_2, \dots\}$ of several flows, each one with its own shaping curve $(\sigma_{f_1}, \sigma_{f_2}, \dots)$, see Figure 2b). It stores all the packets of the aggregate \mathcal{F} in order of arrival into a single FIFO queue and only looks at the head-of-line (HOL) packet. The HOL packet p is released as soon as doing so does not violate the configured shaping curve for the associated flow f_i : Packet p can either be immediately released (if the resulting traffic for f_i at the IR's output is σ_{f_i} -constrained) or delayed to the earliest date that ensures that f_i is σ_{f_i} -constrained at the IR's output. This delay depends on the shaping curve σ_{f_i} for the associated flow f_i and the history of departure dates for previous packets of the same flow. During this delay, any other packet p' in the queue is blocked by the HOL packet p , even if p' belongs to another flow f_j and even if p' could be immediately released without violating the shaping curve σ_{f_j} for its flow f_j .

Traffic regulators can delay individual packets, but there exist specific conditions in which they do not increase the worst-case latency bounds of the flows. This fundamental *shaping-for-free* property is central in the analysis of time-sensitive networks with traffic regulators. It slightly differs for the two flavors.

Theorem 1 (Shaping-for-free property of the PFR [12, Thm. 3]). Consider a flow f with input arrival curve α_f that crosses in sequence a causal system S followed by a PFR (Figure 3a). If the PFR is configured with $\sigma_f \geq \alpha_f$ and if S is FIFO for f , then the worst-case delay $W_{f,S+\text{PFR}}$ of f through the concatenation is equal to the worst-case delay $W_{f,S}$ of the flow through the previous system S .




Fig. 3. Shaping-for-free properties of the traffic regulators. For the PFR, the system S only needs to be FIFO-per-flow. For the IR, S must be FIFO for the aggregate.

Theorem 2 (Shaping-for-free property of the IR [12, Thm. 4]). Consider an aggregate $\mathcal{F} = \{f_1, f_2, \dots\}$ with input arrival curves $\{\alpha_f\}_{f \in \mathcal{F}}$ that crosses in sequence a causal system S followed by an IR (Figure 3b). If the IR is configured with $\forall i, \sigma_{f_i} \geq \alpha_{f_i}$ and if S is FIFO for the aggregate, then the worst-case delay $W_{S+\text{IR}}$ of the aggregate \mathcal{F} through the concatenation is equal to the worst-case delay W_S of the aggregate through the previous system S only.

Theorems 1 and 2 exhibit two fundamental differences. First, Theorem 2 only ensures that the worst-case delay of the aggregate is not increased, whereas Theorem 1 guarantees that the worst-case delay of the individual flow is preserved. Within an aggregate, the first bound can be larger than the latter, *e.g.*, when the flows have different packet sizes. Second, Theorem 1 only requires the previous system S to be FIFO for each flow individually (FIFO-per-flow), whereas the same system is required to be globally FIFO for Theorem 2.

III. RELATED WORK ON THE MODELING OF TRAFFIC REGULATORS

In time-sensitive networks with traffic regulators, end-to-end latency bounds for the flows are obtained by combining the shaping-for-free property for traffic regulators with service-curve-based network-calculus results for other systems. This differentiated treatment of network elements (traffic regulators vs. other systems with service-curve models) restrains the choice of the end-to-end analysis method. Methods based on total-flow analysis (TFA) [13, §3.2] can be adapted to networks with traffic regulators [3], [11]. Other approaches, such as single-flow analysis (SFA) [13, §3.3], pay multiplexing only once (PMOO) [14] and flow prolongation [15] provide tighter end-to-end latency bounds than TFA in several types of networks [16], but they heavily rely on service-curve models. In addition, service-curve models provide continuity and differentiability properties, which allows for synthesizing network designs from the performance requirements, as shown by Geyer and Bondorf in [17]. Hence, a need exists for obtaining service-curve models for all elements of time-sensitive networks, including traffic regulators such as PFRs and IRs.

The per-flow regulator (PFR) was introduced under the name *packetized greedy shaper* in [5, §1.7.4]. Le Boudec and Thiran proved in [5, §1.7.4] that if σ_f is concave and such that $\lim_{t \rightarrow 0^+} \sigma_f(t)$ is larger than the maximum packet size of f , then the PFR offers σ_f as a fluid service curve [5, Thm.

1.7.3]. This model proves Theorem 1. In this paper, we say that the curve σ_f *explains the shaping-for-free property* of the PFR and we formally define this notion in Section V-B. Due to its network-calculus service-curve model, the behavior of the PFR can also be studied in situations where Theorem 1 does not apply. In [18], the consequence of redundancy mechanisms – that can affect the FIFO property – is studied, and end-to-end latency bounds are obtained for flows in networks with redundancy mechanisms and PFRs. In [12, §IV.A], Le Boudec also provides an input-output characterization of the PFR. This type of model does not rely on the concept of service curve but describes the PFR’s output packet sequence as a function of the input packet sequence.

The interleaved regulator (IR) was introduced by Specht and Samii under the name *Urgency-Based Scheduler* [10]. As opposed to the PFR, its shaping-for-free property was proved without the concept of service curves: with a trajectorial approach in [10] and with an input-output characterization in [12, §V]. The equivalence between the theoretical model of the IR and the TSN implementation (*Asynchronous Traffic Shaping*, [2]) was proved by Boyer in [19], who also provides a second input-output characterization [19, §3.3]. The only useful service curves that are known for the IR are only valid when the IR is placed in a context that meets the conditions of Theorem 2. A first context-dependent service curve is provided in [3, §IV.A.1] and then slightly improved in [4, §III.B.1]. In contrast, the only context-agnostic service curve known for the IR is the trivial function $t \mapsto 0$. In [20], Hamscher mentions the first conjecture on a non-trivial context-agnostic service curve for the IR and uses a linear-programming approach for hardening their conjecture pending formal proof. The conjecture was not shared (and, to our knowledge, has not been published at the time of this writing). However, the presentation triggered discussions on whether the IR’s behavior could be captured by context-agnostic service curves. These discussions motivated this paper.

Hence, two questions remain open: Beyond the function $\beta : t \mapsto 0$, what other context-agnostic service curves does the IR provide? Do any of them explain Theorem 2? We address these two questions in this paper.

IV. SYSTEM MODEL AND NOTATIONS

We consider an asynchronous packet-switching time-sensitive network that contains traffic regulators. We focus on a particular traffic regulator within this network. It can either be a per-flow regulator (PFR) that processes a single flow $\mathcal{F} = \{f\}$ with shaping curve σ_f or an interleaved regulator (IR) that processes an aggregate \mathcal{F} with leaky-bucket shaping curves $\{\sigma_f\}_{f \in \mathcal{F}} = \{\gamma_{r_f, b_f}\}_{f \in \mathcal{F}}$. We focus on the subset of flows \mathcal{F} that cross the regulator. We model any other network elements (queues, schedulers, switching fabrics, transmission links, ...) or sequence of network elements crossed by the flows \mathcal{F} between their sources, the regulator, and their destinations as black-box systems. Each system has a traffic input and a traffic output and is only assumed to be causal and lossless: it neither produces nor loses any data internally. Data

TABLE I
NOTATIONS

Common Operators		
$a \vee b$	$= \max(a, b)$	Maximum of a and b .
$a \wedge b$	$= \min(a, b)$	Minimum of a and b .
$[c]^+$	$= \max(0, c)$	
$\lfloor x \rfloor$	$= \max\{n \in \mathbb{Z} n \leq x\}$	Floor function.
$f \otimes g$	$t \mapsto \inf_{s \leq t} f(s) + g(t - s)$	Min-plus convolution
$\bar{\otimes} f, g$	$t \mapsto \sup_{s \leq t} f(s) + g(t - s)$	Max-plus convolution
$\bar{\otimes} f, g$	$t \mapsto \inf_{u > 0} f(t + u) + g(u)$	Max-plus deconvolution
\mathfrak{F}_0	$= \{f : \mathbb{R}^+ \rightarrow \mathbb{R}^+ f(0) = 0\}$	Set of curves
Common Curves		
$\gamma_{r, b}$	$t \mapsto \begin{cases} 0 & \text{if } t = 0 \\ rt + b & \text{if } t > 0 \end{cases}$	Leaky-bucket curve.
$\beta_{R, T}$	$t \mapsto R[t - T]^+$	Rate-latency curve.
Flows		
$f \in \mathcal{F}$	A flow f in the set of flows \mathcal{F}	
$\{\sigma_f\}_{f \in \mathcal{F}}$	A set of shaping curves for the flows \mathcal{F}	
L_f^{\min}, L_f^{\max}	Minimum [resp., maximum] packet size of flow f	
Trajectory Description		
x	A trajectory: Description of all the events in the network	
M	An observation point	
\mathcal{M}^x	Packet sequence at M in trajectory x	
$R^{x, M}$	Cumulative function of the aggregate...	
[resp., $R_f^{x, M}$]	... [resp., of f] at M in Trajectory x .	
$R^{x, M} \sim \alpha$	$R^{x, M}$ is constrained by α , Equation (3)	
Parameters of the Spring adversary (Section V-A)		
I	Expected spacing for same-flow packets after the IR	
d	Maximum delay in the Spring-controlled system S_1	
ϵ	Margin (minimum packet spacing after S_1)	
τ	Period of the six-packet-long profile	



Fig. 4. Input [resp., output] observation point B [resp., D] for a regulator.

is produced at the flows’ sources and consumed at the flows’ destinations.

A *trajectory* x is a description of all the events in the network (packet arrival, packet departure). It is *acceptable* if all known constraints are satisfied. For an observation point M , we denote by $R^{x, M}$ [resp., $R_f^{x, M}$] the cumulative function of the aggregate [resp., of the flow $f \in \mathcal{F}$] at observation point M in trajectory x . If the stream is packetized at M , we call \mathcal{M}^x the packet sequence that describes the packets’ arrival date, size, and associated flow at M in trajectory x .

For an input packet sequence \mathcal{B}^x at the input B of the regulator (Figure 4), we use the equivalent input-output characterizations of traffic regulators from [12] and [19] to obtain the output packet sequence \mathcal{D}^x at the output D of the regulator.

We list the notations in Table I. Each result in this paper is associated with an intuition of the proof, and the formal proofs are available in [21].

V. LIMITS OF THE SHAPING-FOR-FREE PROPERTY FOR THE INTERLEAVED REGULATOR

The shaping-for-free property is a strong attribute of the interleaved regulator (IR). However, it is context dependent: It makes assumptions on the context in which the IR is placed. In this section, we investigate the limits of these assumptions.




Fig. 5. The network \mathcal{N}_1 and the Spring-generated Trajectory 1 that yields unbounded latencies in the IR when S_1 is not assumed FIFO.

First, we observe that Theorem 2 requires the upstream system to be FIFO. In Section V-A, we prove that removing this assumption makes the IR unstable: it can yield unbounded latencies. We then prove in Section V-B that there exists no service-curve model of the IR that can explain Theorem 2.

A. Instability of the IR when Placed after a Non-FIFO System

In this subsection, we discuss the role of the FIFO assumption in Theorem 2. When removed, we prove that the IR can yield unbounded latencies. Specifically, we prove

Theorem 3 (Instability of the IR after a non-FIFO system). Consider an IR that processes three or more flows with the same leaky-bucket shaping curve for the first three flows: $\forall f_i \in \{f_1, f_2, f_3\}, \sigma_{f_i} = \gamma_{r,b}$ with $r > 0$ and b greater than the maximum packet size of f_1, f_2, f_3 . For any $W > 0$, there exists a system S_1 and a source ϕ (Figure 5), such that:

- 1/ each flow f_i is σ_{f_i} -constrained at the source ϕ ,
- 2/ S_1 is causal, lossless and FIFO-per-flow (but globally non-FIFO),
- 3/ when the system S_1 is placed after the source as in Figure 5, then the delay of each flow within S_1 is upper-bounded by W ,
- 4/ when the IR is placed after S_1 as in Figure 5, then the delay of any flow within the IR is not bounded.

The proof of Theorem 3 in [21, Appendix B.A] relies on an adversarial traffic generation that we call ‘‘Spring’’. Spring is an adversary that knows the values of b, r and W in Theorem 3 and controls the source ϕ and the system S_1 of Figure 5 such that Properties 1/ to 4/ of Theorem 3 hold. It defines the constants I, d, ϵ and τ as follows

$$I \triangleq \frac{b}{r}; 0 < d < \min(I, W); 0 < \epsilon < \min(I - d, \frac{d}{3}); \tau \triangleq 3I + 3\epsilon - d \quad (5)$$

Intuition: Trajectory 1 generated by Spring is illustrated in Figure 5. All packets have the size b . The far-left timeline shows the packet sequence \mathcal{A}^1 for the three flows at the output of the Spring-controlled source. A sequence of six packets is repeated with period τ . Only the period $k = 0$ is shown.

The dotted arrows that lead to the second timeline highlight each packet’s delay in the Spring-controlled system S_1 and the resulting packet sequence \mathcal{B}^1 . The main property of Trajectory 1 is that the first packet of the dash-dotted red flow f_2 and the second packet of the solid blue flow f_1 have exchanged their order at B compared to their order at A. This is because the former suffers a delay d through S_1 , but the latter does not suffer any delay. Note that the Spring-controlled system S_1 is not FIFO but remains causal, lossless, and FIFO-per-flow with a delay bound $d < W$.

The dotted arrows that link the second to the third timeline describe the behavior of the IR (not controlled by Spring) when provided with the input sequence \mathcal{B}^1 . For example, the first packet of f_1 is immediately released by the IR because the network was previously empty. However, the second solid blue packet of f_1 is shaped (delayed) by the IR because releasing it would violate the $\gamma_{r,b}$ shaping constraint for f_1 at the output of the IR. This packet is released as soon as doing so does not violate the $\gamma_{r,b}$ constraint, *i.e.*, I seconds after the previous packet. Because of this, the first dash-dotted red packet of the flow f_2 is blocked by the head-of-line (HOL). And the second packet of f_2 is shaped and delayed to ensure a distance of I from the previous packet of f_2 .

As a result, it takes $3I$ seconds for the IR to output the six packets of the first period, whereas they entered the IR within τ seconds. As $\tau < 3I$, we can generate a constant build-up of delay and backlog in the IR by repeating the six-packet-long profile every τ seconds.

B. The Shaping-for-Free Property of the Interleaved Regulator Cannot be Explained by a Service Curve

Theorem 3 shows that the IR does not provide any context-agnostic delay guarantees as a stand-alone network element. In contrast, if a system Z offers a context-agnostic service curve β that explains a context-dependent property (*e.g.*, shaping-for-free), then β continues to hold when Z is placed in a context that differs from the assumptions of the context-dependent property. β can be used to compute the consequences of the deviation from the assumptions and their resulting penalties on performance bounds. For such a system, slight deviations from the assumptions should lead to small delay penalties.

This is the case for the PFR, for which we can find a service-curve model that explains its shaping-for-free property. We formally define this notion as follows.

Definition 2 (A curve explains the shaping-for-free property). Consider a set of flows \mathcal{F} and a set of shaping curves $\sigma = \{\sigma_f\}_{f \in \mathcal{F}}$. We say that $\beta^\sigma \in \mathfrak{F}_0$ explains the shaping-for-free property if and only if: For any causal, lossless and FIFO systems Z' and S , if Z' offers β^σ as a service curve, then the worst-case delay of the aggregate between A and B (Figure 6)




Fig. 6. Notations of Definition 2. β^σ explains the shaping-for-free property if any system Z' that offers β^σ as a service curve does not increase the worst-case delay of the flows when placed after any FIFO system S .




Fig. 7. Application of Proposition 1 to prove Theorem 1. (a) A PFR placed in the conditions of Theorem 1. (b) The equivalent model as per Proposition 1.

over all the trajectories $X = \{x | \forall f \in \mathcal{F}, R_f^{A,x} \sim \sigma_f\}$ equals the worst-case delay between A and C over the same set X .

As per this definition, a function β^σ explains the shaping-for-free property if any system Z' that offers β^σ as a service curve does not increase the worst-case delay of the aggregate \mathcal{F} when placed in a context that meets the assumptions of Theorems 1 and 2 (*i.e.*, placed after a FIFO system S with flows that are initially constrained by their shaping curves). For the PFR, we have a positive result:

Proposition 1. Consider a PFR that shapes a single flow $\mathcal{F} = \{f\}$ with a concave shaping curve σ_f such that $\lim_{t \rightarrow 0} \sigma_f(t) \geq L_f^{\max}$. Then the PFR offers the fluid service curve $\beta^\sigma = \sigma_f$ that explains the shaping-for-free property.

The formal proof in [21, Appendix B.B] directly derives from [5, Thm. 1.7.3]. Note that Proposition 1 contains two statements: (1) σ_f is a fluid service curve of the PFR. (2) σ_f explains the shaping-for-free (Definition 2).

Let us discuss why these two statements prove Theorem 1. Consider a causal, lossless, and FIFO system S and a PFR configured with σ_f placed after S (Figure 7a). By combining the first statement of Proposition 1 with Definition 1, the PFR can be realized as the concatenation of Z' followed by a packetizer (Figure 7b), where Z' is a causal, lossless, and FIFO system that offers $\beta^\sigma = \sigma_f$ as a service curve. We then combine the second statement of Proposition 1 with Definition 2. We obtain that if σ_f is an arrival curve for f at the input of S , then the worst-case delay of the flow f through S equals the worst-case delay of the flow through the concatenation of S and Z' . Finally, the packetizer does not increase the worst-case latency bounds [12, Thm. 1.7.1], which proves Theorem 1.

As opposed to the PFR, we prove that no fluid service curve can explain the shaping-for-free property of the IR:

Theorem 4. An IR that processes at least three flows with the same leaky-bucket shaping curve does not have any fluid service curve that explains its shaping-for-free property.

To prove Theorem 4, we rely on the following lemma, that we prove in [21, Appendix B.C].

Lemma 1. If β^σ explains the shaping-for-free (Definition 2), then $\beta^\sigma \geq \sum_{f \in \mathcal{F}} \sigma_f$

By reusing Spring (from the proof of Theorem 3), we then exhibit a trajectory that shows that a function larger than $\sum_{f \in \mathcal{F}} \sigma_f$ cannot be a fluid service curve of the IR. The formal proof of Theorem 4 is in [21, Appendix B.D].

VI. SERVICE CURVES OF THE INTERLEAVED REGULATOR

In the previous section, we use the sensitivity of the interleaved regulator (IR) to the FIFO assumption in Theorem 2 to prove that the IR has no fluid service curve that can explain its shaping-for-free property.

In this section, we show that the IR still offers a family of non-trivial context-agnostic service curves. In particular, we exhibit a non-bounded strict service curve for the IR (Theorem 5). This strict-service-curve model is of interest for understanding the behavior of the IR in situations that differ from the shaping-for-free property. It can be used to model the IR in service-curve-oriented analysis like SFA and PMOO.

However, any service curves of the IR are also fluid service curves of the IR, and we know from Theorem 4 that they must be weak because they cannot explain its shaping-for-free property. Indeed, their long-term rate is upper bounded: For an IR that processes more than four flows, the long-term rate of any of its service curves is upper bounded by three times the rate enforced for a single flow, as we show in Theorem 7.

For an IR that processes at least four flows, we also prove that any individual service curve β_g offered to a single flow g is upper bounded by its minimum packet size L_g^{\min} (Theorem 6).

The section is organized as follows. First, we obtain a strict service curve of the IR (Theorem 5) by using the input/output models of [12], [19]. Then, we use the Spring trajectory from Section V to obtain upper bounds on the individual service curve of the IR for a single flow (Theorem 6). Last, we use Theorem 6 to upper-bound the long-term rate of any context-agnostic service curve of the IR for the aggregate (Theorem 7).

A. A Strict Service Curve for the Aggregate

Even though Spring's Trajectory from Section V-A generates a constant build-up of delay in the IR, we can observe in Figure 5 that the IR continuously outputs two packets every I seconds. In fact, we can find a minimum output rate whenever the IR is non-empty. This shows that the IR offers a strict service curve as defined in Section II-A:

Theorem 5 (IR strict service curve). Consider an IR that processes an aggregate of flows \mathcal{F} with leaky-bucket shaping curves: $\forall f \in \mathcal{F}, \sigma_f = \gamma_{r_f, b_f}$ with $r_f > 0$ and $b_f \geq L_f^{\max}$, where L_f^{\min} [resp., L_f^{\max}] is the minimum [resp., maximum] packet size of f . Define

$$L^{\min} = \min_{f \in \mathcal{F}} L_f^{\min} \quad I^{\max} = \max_{f \in \mathcal{F}} \frac{L_f^{\max}}{r_f} \quad (6)$$

Then the staircase curve $\beta_{sc} : t \mapsto \lfloor t/I^{\max} \rfloor \cdot L^{\min}$ (with $\lfloor \cdot \rfloor$ the floor function, Table I) and the rate-latency curve $\beta_{R,T}$




Fig. 8. Three different strict service curves of the IR (Theorem 5). The step dashed blue function β_0 is the first curve obtained in the theorem's proof. Its super-additive closure, the solid red function β_{sc} is then also a strict service curve, as well as the dash-dotted green rate-latency curve $\beta_{\frac{L^{\min}}{I^{\max}}, I^{\max}}$.

and $T = I^{\max}$ and $R = L^{\min}/I^{\max}$ (Figure 7) are context-agnostic strict service curves of the IR for the aggregate.

To prove this result, we consider a non-empty IR. The output time of the head-of-line packet is given by the input/output characterizations of [12], [19]. It can be upper-bounded. From this we obtain that the dashed-blue curve β_0 in Figure 8 is a strict service curve of the IR, and so is its super-additive closure [7, Prop 5.6], defined as the function

$$\beta_0 \vee (\beta_0 \otimes \beta_0) \vee ((\beta_0 \otimes \beta_0) \otimes \beta_0) \vee \dots \quad (7)$$

where \vee is the maximum and \otimes is the max-plus convolution (Table I). The computation of (7) gives β_{sc} , shown in solid red in Figure 8. Any wide-sense increasing curve that remains below the β_{sc} service curve is also a strict service curve of the IR [7, Prop 5.6]. This is the case for the rate-latency service curve $\beta_{L^{\min}/I^{\max}, I^{\max}}$ shown with a dash-dotted green line in Figure 8. The formal proof of Theorem 5 is in [21, Appendix B.E].

Application to the Situation of Section V-A. Figure 9 shows the cumulative arrival function R^B at the input of the IR as well as the cumulative departure function R^D at the output of the IR, in Spring's trajectory¹ described in Section V-A and Figure 5. We also provide the arrival curve α^B of the aggregate at the input of the IR.

In Section V-A, all packets of the aggregate have the same size L and all three flows have the same leaky-bucket shaping curve $\sigma_{f_1} = \sigma_{f_2} = \sigma_{f_3} = \gamma_{r,b}$ with $b = L$. The application of Theorem 5 gives that $\beta_{\frac{L}{T}, I} = \beta_{r, \frac{b}{r}}$ is a context-agnostic strict service curve of the IR for the aggregate. In Figure 5, we place this curve in green at the beginning of the backlog period. We confirm that when the IR is non-empty, the cumulative output

¹With the notations of (5), the parameters used in Figure 9 are: $W = 0.86I$, $d = 0.85I$, $\epsilon = 0.05I$.




Fig. 9. Application of Theorem 5 to the Spring trajectory of Section V-A. The dotted red [resp., blue] curve is the cumulative function at the input [resp., output] of the IR. The dash-dotted orange leaky-bucket curve is an arrival curve of the aggregate at the input of the IR. The solid green rate-latency curve is a strict service curve of the IR, as proved by Theorem 5.

R^B is larger than the strict service curve. Also, it is clear that the horizontal deviation between the arrival curve α^B and the service curve $\beta_{\frac{L}{T}, I}$ is not bounded. This means that the network-calculus theory cannot provide a latency bound with this service-curve model, which is consistent with Theorem 3.

In Figure 9, the rate $\frac{L}{T} = r$ of the green service curve $\beta_{\frac{L}{T}, I}$ does not follow the long-term rate $\frac{2L}{T} = 2r$ of the output cumulative function R^D . However, no better rate for the rate-latency context-agnostic strict service curve can be achieved:

Proposition 2 (Upper Bound on the Strict Service Curve). Consider an IR that processes an aggregate \mathcal{F} with leaky-bucket shaping curves $\{\sigma_f\}_{f \in \mathcal{F}} = \{\gamma_{r_f, b_f}\}_{f \in \mathcal{F}}$. Consider a curve β^{strict} and assume that β^{strict} is a context-agnostic strict service curve of the IR. Then

$$\forall t \geq 0, \quad \beta^{\text{strict}}(t) \leq \min_{f \in \mathcal{F}} \sigma_f(t) \quad (8)$$

In particular, if $\beta^{\text{strict}} = \beta_{R, T}$ is a rate-latency curve, then $R \leq \min_{f \in \mathcal{F}} r_f$.

To prove this result, we consider a set of trajectories $\{x_f\}_{f \in \mathcal{F}}$. For each flow $f \in \mathcal{F}$, the trajectory x_f is obtained by having only flow f send packets of size b_f to the IR at twice the frequency allowed by its shaping curve. Then the backlog of the IR quickly becomes non-empty, but the cumulative output of the IR, $R^{x_f, D} = R_f^{x_f, D}$ is constrained by the shaping curve σ_f : $\forall 0 \leq s \leq t, R^{x_f, D}(t) - R^{x_f, D}(s) \leq \sigma_f(t - s)$. Combined with the definition of a strict service curve (2), this gives $\beta^{\text{strict}} \leq \sigma_f$. This is valid for all trajectories $\{x_f\}_{f \in \mathcal{F}}$, hence the result. The formal proof is in [21, Appendix B.F].

Figure 10 shows the areas of proven, possible, and impossible strict service curves of an IR that processes the flows with the same leaky-bucket shaping curve $\gamma_{r,b}$, assuming that all the packets have the size b . Any wide-sense increasing function that remains in the green area is a strict service curve of the




Fig. 10. Areas of proven, possible, and impossible strict service curves for an IR that processes the flows with the same shaping curve $\gamma_{r,b}$ and if all packets have size b .

IR, as proven by Theorem 5. In contrast, any function that enters the red area cannot be a strict service curve of the IR, as proven by Proposition 2.

Incidentally, any wide-sense increasing function that remains within the green area is also a service curve of the IR. So far, as opposed to the strict-service-curve property, we have not obtained any limit for the service curve (1). To obtain this limit, we first need to consider the *individual* service curve offered to any flow by the IR, which we analyze in the following subsection.

B. Upper-Bound on the Individual Service Curve

Consider an IR that processes an aggregate \mathcal{F} and take a flow $g \in \mathcal{F}$. In this section, we are interested in the service the IR guarantees to g , i.e., in an *individual* service curve of the IR for g .

If each flow $f \in \mathcal{F} \setminus \{g\}$ enters the IR with an arrival curve α_f^B that is equal or smaller than its shaping curve σ_f , then none of the packets of the flows $\mathcal{F} \setminus \{g\}$ is ever delayed by the IR. In this case, the IR acts as a PFR for g and provides an individual fluid service curve σ_g for g .

In a more likely setting, though, the IR reshapes flows that exhibit an input arrival curve (strictly) larger than their configured shaping curve. In such a case, and if the IR processes more than four flows, no useful individual service curve for g can be obtained:

Theorem 6 (Upper-Bound on the Individual Service Curve). Consider an IR that processes an aggregate \mathcal{F} of at least four flows and with the same leaky-bucket shaping curve for at least three of them: $\forall f_i \in \{f_1, f_2, f_3\}, \sigma_{f_i} = \gamma_{r,b}$. Consider a flow $g \in \mathcal{F} \setminus \{f_1, f_2, f_3\}$ and assume that for each $f_i \in \{f_1, f_2, f_3\}, \gamma_{r,b_i}$ is an arrival curve for f_i at the input B of the IR (Figure 4), with $b_1 > b, b_2 \geq b, b_3 \geq b$ (permutating the indexes if required). Last, consider a curve $\beta_g \in \mathfrak{F}_0$ that can depend on $\{\sigma_f\}_{f \in \mathcal{F}}$ and on $\{\alpha_h^B\}_{h \in \mathcal{F} \setminus \{g\}}$.

If β_g is an individual service curve of the IR for the flow g , then β_g is upper-bounded by g 's minimum packet size L_g^{\min} .

Theorem 6 shows that any individual-service-curve model of the IR for a flow g can only guarantee that one single packet of g will ever cross the IR over the entire network's lifetime. Hence, no useful context-agnostic service curve exists to model the service offered to a single flow by the IR (each flow is likely to send many packets).

To prove Theorem 6, we reuse the Spring trajectory of Theorem 3 for the three flows f_1, f_2, f_3 . This trajectory creates a constant build-up of delay and backlog inside the IR. We then consider a fourth flow, g , and its first packet of size L . On one hand, if β_g is an individual service curve of the IR for the flow g , then the delay of the first packet of g through the IR can be upper bounded by $\inf\{t \in \mathbb{R}^+ | \beta_g(t) \geq L\}$. On the other hand, the first packet of g can suffer a delay as large as desired within the Spring trajectory: We send it when the accumulated delay in the IR is large enough. The combination of the two above observations provides the result. The formal proof is in [21, Appendix B.G].

C. Limits on the Aggregate Service Curve

Let us go back to the analysis of the service offered to the aggregate, by focusing on (non-strict) service-curve models. One consequence of Theorem 6 is that the long-term rate of the service curve for the aggregate is upper-bounded by three times the rate of a single contract.

Theorem 7 (Maximum long-term rate of any service curve). Consider an IR that processes an aggregate \mathcal{F} of at least four flows and with the same leaky-bucket shaping curve for at least three of them: $\forall f_i \in \{f_1, f_2, f_3\}, \sigma_{f_i} = \gamma_{r,b}$. Consider a curve $\beta \in \mathfrak{F}_0$ that can depend on $\{\sigma_f\}_{f \in \mathcal{F}}$. If the IR offers β as a context-agnostic service curve, then

$$\liminf_{t \rightarrow +\infty} \frac{\beta(t)}{t} \leq 3r \quad (9)$$

To prove Theorem 7, we pick one flow $g \in \mathcal{F} \setminus \{f_1, f_2, f_3\}$ and a traffic arrival for $h \in \mathcal{F} \setminus \{g\}$ that meets the requirements of Theorem 6. The IR is a FIFO system. Hence, if β is a service curve of the IR, then for any $\theta \in \mathbb{R}^+$, the curve β_g^θ defined by

$$\beta_g^\theta : t \mapsto \left[\beta(t) - \sum_{f \in \mathcal{F} \setminus \{g\}} \alpha_f(t - \theta) \right]^+ \cdot \mathbb{1}_{\{t > \theta\}} \quad (10)$$

verifies Inequation (4) [5, Prop. 6.4.1]. Note that β_g^θ may not be wide-sense increasing, thus not an individual arrival curve for g . Inspired by [7, §5.2.1], we resolve this by considering the curve $\beta_g^\theta \bar{\circ} \mathbf{0} : t \mapsto \inf_{s \geq t} \beta_g^\theta(s)$, where $\bar{\circ}$ is the max-plus deconvolution (Table I) and $\mathbf{0}$ is the zero function $t \mapsto 0$. The function $\beta_g^\theta \bar{\circ} \mathbf{0}$ is wide-sense increasing, smaller than β_g^θ , thus an individual service curve for g . From Theorem 6, we obtain $\forall \theta \geq 0, \forall t \geq 0, (\beta_g^\theta \bar{\circ} \mathbf{0})(t) \leq L_g^{\min}$. The left-hand side of this inequation is an infimum, but the result is valid for any $\theta \geq 0, t \geq 0$. Hence, we can derive a bound on the long-term rate of β . The formal proof is in [21, Appendix B.H].

With Theorem 7, we can conclude that for an IR that processes more than four flows, no useful context-agnostic

TABLE II
CONTEXT-AGNOSTIC SERVICE CURVES FOR AN IR THAT PROCESSES AT
LEAST FOUR FLOWS WITH THE SAME LEAKY-BUCKET SHAPING CURVE
 $\gamma_{r,b}$, ASSUMING ALL PACKETS HAVE SIZE b .

Service-curve type	Curve exhibited in this paper	Limit exhibited in this paper
Service curve β	$\beta_{r, \frac{b}{r}}$	$\liminf_{t \rightarrow +\infty} \frac{\beta(t)}{t} \leq 3r$
Strict service curve β^{strict}	$\beta_{r, \frac{b}{r}}$	$\forall t \geq 0, \beta^{\text{strict}}(t) \leq \gamma_{r,b}(t)$
Individual service curve $\beta_g, \forall g \in \mathcal{F}$	None ($t \mapsto 0$)	$\forall t \geq 0, \beta_g(t) \leq L_g^{\min}$

service curve exists to model the IR. Indeed, the aggregate of the four flows can exhibit a sustained rate of four times the rate of a single-flow contract, whereas a context-agnostic service curve can only guarantee a long-term service rate of three times this value. Hence, Theorem 7 concludes our search of context-agnostic service-curve models for the IR and the results of Section VI are summarized in Table II.

VII. CONCLUSION

Network calculus is a framework for obtaining worst-case performance bounds of time-sensitive networks, as required for their validation. Most of the mechanisms standardized by the time-sensitive networking (TSN) task group of the IEEE enjoy a network-calculus service-curve model published in the literature. The interleaved regulator (IR), standardized as *asynchronous traffic shaping* (ATS) in TSN, is an exception. Its shaping-for-free property is instrumental in designing and analyzing time-sensitive networks but was proved without network-calculus service curves. The existence of a service-curve model that explains the IR's behavior and its shaping-for-free property remained an open question. If such a model existed, network engineers could use the IR outside of the shaping-for-free requirements and still compute end-to-end performance bounds with service-curve-oriented tools.

In this paper, we settled the question: Network-calculus service curves cannot explain the behavior of the IR. We show that the IR still offers non-trivial functions as (strict) service curves, but (a) none of them can explain the shaping-for-free property of the IR and (b) these curves are too weak to be helpful and cannot offer any delay guarantee in most cases. Consequently, performance bounds cannot be obtained with service-curve-oriented approaches when the IR is used in a context that differs from the shaping-for-free requirements, *e.g.*, after a non-FIFO system. We prove that these bounds do not even exist: The IR can yield unbounded latencies after a non-FIFO system.

VIII. ACKNOWLEDGMENTS

The authors thank Anja Hamscher for the initial discussions on the topic and the reviewers for their valuable feedback.

REFERENCES

- [1] J. Farkas, "TSN Basic Concepts," p. 20, Nov. 2018. [Online]. Available: <https://www.ieee802.org/1/files/public/docs2018/detnet-tsn-farkas-tsn-basic-concepts-1118-v01.pdf>
- [2] "IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks - Amendment 34:Asynchronous Traffic Shaping," *IEEE Std 802.1Qcr-2020*, pp. 1–151, Nov. 2020.
- [3] E. Mohammadpour, E. Stai, M. Mohiuddin, and J.-Y. Le Boudec, "Latency and Backlog Bounds in Time-Sensitive Networking with Credit Based Shapers and Asynchronous Traffic Shaping," in *2018 30th International Teletraffic Congress (ITC 30)*, vol. 02, Sep. 2018, pp. 1–6.
- [4] L. Zhao, P. Pop, and S. Steinhorst, "Quantitative Performance Comparison of Various Traffic Shapers in Time-Sensitive Networking," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2899–2928, Sep. 2022.
- [5] J.-Y. Le Boudec and P. Thiran, *Network Calculus*, ser. Lecture Notes in Computer Science, G. Goos, J. Hartmanis, and J. van Leeuwen, Eds. Berlin, Heidelberg: Springer, 2001, vol. 2050. [Online]. Available: <http://link.springer.com/10.1007/3-540-45318-0>
- [6] C.-S. Chang, *Performance Guarantees in Communication Networks*. Springer London, 2000. [Online]. Available: <https://doi.org/10.1007/978-1-4471-0459-9>
- [7] A. Bouillard, M. Boyer, and E. Le Corronc, *Deterministic Network Calculus*. John Wiley & Sons, Inc., Oct. 2018. [Online]. Available: <https://doi.org/10.1002/9781119440284>
- [8] L. Maile, K.-S. Hielscher, and R. German, "Network Calculus Results for TSN: An Introduction," in *2020 Information Communication Technologies Conference (ICTC)*, May 2020, pp. 131–140.
- [9] L. Zhao, P. Pop, Z. Zheng, H. Daigmore, and M. Boyer, "Latency Analysis of Multiple Classes of AVB Traffic in TSN with Standard Credit Behavior using Network Calculus," May 2020. [Online]. Available: <http://arxiv.org/abs/2005.08256>
- [10] J. Specht and S. Samii, "Urgency-Based Scheduler for Time-Sensitive Switched Ethernet Networks," in *2016 28th Euromicro Conference on Real-Time Systems (ECRTS)*, Jul. 2016, pp. 75–85.
- [11] L. Thomas, J.-Y. Le Boudec, and A. Mifdaoui, "On Cyclic Dependencies and Regulators in Time-Sensitive Networks," in *2019 IEEE Real-Time Systems Symposium (RTSS)*, Dec. 2019, pp. 299–311.
- [12] J.-Y. Le Boudec, "A Theory of Traffic Regulators for Deterministic Networks With Application to Interleaved Regulators," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2721–2733, Dec. 2018.
- [13] J. B. Schmitt and F. A. Zdarsky, "The DISCO network calculator: A toolbox for worst case analysis," in *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*, ser. Valuetools '06. New York, NY, USA: Association for Computing Machinery, Oct. 2006, pp. 8–es. [Online]. Available: <https://doi.org/10.1145/1190095.1190105>
- [14] J. B. Schmitt, F. A. Zdarsky, and I. Martinovic, "Improving Performance Bounds in Feed-Forward Networks by Paying Multiplexing Only Once," in *14th GIITG Conference - Measurement, Modelling and Evaluation of Computer and Communication Systems*, Mar. 2008, pp. 1–15.
- [15] S. Bondorf, "Better bounds by worse assumptions — Improving network calculus accuracy by adding pessimism to the network model," in *2017 IEEE International Conference on Communications (ICC)*, May 2017, pp. 1–7.
- [16] S. Bondorf, P. Nikolaus, and J. B. Schmitt, "Quality and Cost of Deterministic Network Calculus: Design and Evaluation of an Accurate and Fast Analysis," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 1, no. 1, pp. 16:1–16:34, Jun. 2017. [Online]. Available: <https://doi.org/10.1145/3084453>
- [17] F. Geyer and S. Bondorf, "Network Synthesis under Delay Constraints: The Power of Network Calculus Differentiability," in *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, May 2022, pp. 1539–1548.
- [18] L. Thomas, A. Mifdaoui, and J.-Y. Le Boudec, "Worst-Case Delay Bounds in Time-Sensitive Networks With Packet Replication and Elimination," *IEEE/ACM Transactions on Networking*, pp. 1–15, 2022.
- [19] M. Boyer, "Equivalence between the theoretical model and the standard algorithm of Asynchronous Traffic Shaping," Sep. 2022. [Online]. Available: <https://hal.science/hal-03788302>
- [20] A. Hamscher, "Using Mathematical Programming to Harden Conjectures on Service Curves," Lausanne, Switzerland, Sep. 2022. [Online]. Available: https://www.youtube.com/watch?v=eJwOrMIL5Q&list=PLGrWRLGd9yS_nezfKdxK1x-e3yNt1krwj&index=19
- [21] L. Thomas and J.-Y. Le Boudec, "Network-Calculus Service Curves of the Interleaved Regulator," May 2023. [Online]. Available: <http://arxiv.org/abs/2305.18036>