



HAL
open science

Interleaved Weighted Round-Robin: A Network Calculus Analysis

Seyed Mohammadhossein Tabatabaee, Jean-Yves Le Boudec, Marc Boyer

► **To cite this version:**

Seyed Mohammadhossein Tabatabaee, Jean-Yves Le Boudec, Marc Boyer. Interleaved Weighted Round-Robin: A Network Calculus Analysis. International Teletraffic Congress, Sep 2020, Osaka, Japan. hal-03007964

HAL Id: hal-03007964

<https://hal.science/hal-03007964>

Submitted on 16 Nov 2020

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.

Interleaved Weighted Round-Robin: A Network Calculus Analysis

Seyed Mohammadhossein Tabatabaee

EPFL

Lausanne, Switzerland

hossein.tabatabaee@epfl.ch

Jean-Yves Le Boudec

EPFL

Lausanne, Switzerland

jean-yves.leboudec@epfl.ch

Marc Boyer

ONERA/DTIS, University of Toulouse

F-31055 Toulouse, France

Marc.Boyer@onera.fr

Abstract—Weighted Round-Robin (WRR) is often used, due to its simplicity, for scheduling packets or tasks. With WRR, a number of packets equal to the weight allocated to a flow can be served consecutively, which leads to a bursty service. Interleaved Weighted Round-Robin (IWRR) is a variant that mitigates this effect. We are interested in finding bounds on worst-case delay obtained with IWRR. To this end, we use a network calculus approach and find a strict service curve for IWRR. The result is obtained using the pseudo-inverse of a function. We show that the strict service curve is the best obtainable one, and that delay bounds derived from it are tight (i.e., worst-case) for flows of packets of constant size. Furthermore, the IWRR strict service curve dominates the strict service curve for WRR that was previously published. We provide some numerical examples to illustrate the reduction in worst-case delays caused by IWRR compared to WRR.

I. INTRODUCTION

Weighted Round-Robin (WRR) is a scheduling algorithm that is often used for scheduling tasks, or packets, in real-time systems or communication networks. The capacity is shared between several clients or queues by giving each of them a weight, which is a positive integer, and by providing more service to those with larger weights. Specifically, every queue is visited one after the other, and when a queue i with weight w_i has an emission opportunity, it sends w_i packets, or less if fewer packets are present. The advantage of WRR is that it is fair and simple. However, the service is bursty because up to w_i packets can be served consecutively for queue i , which can cause a large worst-case waiting time for other queues. Interleaved Weighted Round-Robin (IWRR) mitigates this effect [1]. With IWRR, a queue i with weight w_i has w_i emission opportunities per round and can send up to one packet at every emission opportunity. In contrast, with WRR, it has one emission opportunity per round and can send up to w_i packets at every emission opportunity. Hence, IWRR spreads out emission opportunities of each queue in a round, which is expected to result in a smoother service and lower worst-case delays. There exist several versions of IWRR; we focus on the simplest one, where queue i has emission opportunities in the first w_i cycles within a round (see Section III for a formal description of IWRR and Section IV for WRR variants).

We are interested in delay bounds for the worst case, as is typical in the context of deterministic networking. To this end, a standard approach is network calculus. Specifically,

with network calculus, the service offered to a flow of interest by a system is abstracted by means of a service curve. A bound on the worst-case delay is obtained by combining the service curve with an arrival curve for the flow of interest. An arrival curve is a constraint on the amount of data that the flow of interest can send; such a constraint is necessary to the existence of a finite delay bound. The exact definitions are recalled in Section II.

The network calculus approach was applied to WRR in [2, Sec. 8.2.4], where a *strict* service curve is obtained. As explained in Section II, a strict service curve is a special case of a service curve hence can be used to derive delay (and backlog) bounds. Our first contribution is to obtain a strict service curve for IWRR. Compared to WRR, the interleaving in IWRR makes the analysis more difficult, and the method of proof in [2] cannot easily be extended. To circumvent this difficulty, we rely heavily on the method of pseudo-inverse, recalled in Section II. As expected, the IWRR strict service curve dominates that of WRR, hence the resulting delay bounds for IWRR are always less than or equal to those for WRR.

The strict service curve enables us to obtain delay bounds by using network calculus, but such bounds might not always be tight, i.e., they might not always be equal to worst-cases. This is because the strict service curve is an abstraction of the system. Our second contribution is to show that, for flows with packets of constant sizes, the strict service curve obtained for IWRR provides tight delay bounds. We show that the same result holds for the existing strict service curve of WRR. Extending such results to flows with packets of variable sizes is left for further study.

The strict service curve obtained for IWRR has some description complexity, see also Fig. 3. Therefore, we provide simplified lower bounds that can be used, at the expense of sub-optimality, when analytic, closed-form expressions are important.

After giving some necessary background on network calculus and the lower-pseudo inverse technique in Section II, we describe our system model in Section III. We describe the state of the art in Section IV. In Section V, we present our strict service curve for IWRR, the proof of which we present in Section VI. In Section VII, we show that both the IWRR and WRR strict service curves are the best possible and that

they give tight delay bounds for a flow with constant packet sizes. We use numerical examples to illustrate the worst-case latency improvement of IWRR over WRR obtained with our method in Section VIII. Due to lack of space, proofs of results other than Theorem 1 are in [3].

II. BACKGROUND

We use the framework of network calculus [2], [4], [5]. A flow is represented by a cumulative arrival function $R \in \mathcal{F}$, where \mathcal{F} denotes the set of wide-sense increasing functions $f: \mathbb{R}^+ \mapsto \mathbb{R}^+ \cup \{+\infty\}$ and $R(t)$ is the number of bits observed on the flow between times 0 and t . We say that a flow R has $\alpha \in \mathcal{F}$ as arrival curve if for all $s \leq t$, $R(t) - R(s) \leq \alpha(t - s)$. A frequently used arrival curve is $\alpha = \gamma_{r,b}$, defined by $\gamma_{r,b}(t) = rt + b$ for $t > 0$ and $\gamma_{r,b}(t) = 0$ for $t = 0$ (token bucket arrival curve, with rate r and burst b). An arrival curve α can always be assumed to be sub-additive, i.e., to satisfy $\alpha(s + t) \leq \alpha(s) + \alpha(t)$ for all s, t .

For two functions f and g in \mathcal{F} , the min-plus convolution is defined by $(f \otimes g)(t) = \inf_{0 \leq s \leq t} \{f(t-s) + g(s)\}$. An example of min-plus convolution used in this paper is illustrated in Fig. 1.

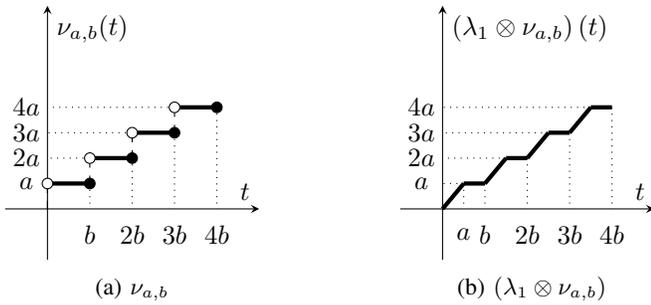


Fig. 1: Left: the stair function $\nu_{a,b} \in \mathcal{F}$ defined for $t \geq 0$ by $\nu_{a,b}(t) = a \lceil \frac{t}{b} \rceil$. Right: min-plus convolution of $\nu_{a,b}$ with the function $\lambda_1 \in \mathcal{F}$ defined by $\lambda_1(t) = t$ for $t \geq 0$, when $a \leq b$. The discontinuities are smoothed, and replaced with a unit slope.

Consider a system S and a flow through S with input and output functions R and R^* and let $\beta \in \mathcal{F}$. We say that the system S offers β as a service curve to the flow if $R^* \geq R \otimes \beta$, which often means that for every $t \geq 0$ there exists some $s \leq t$ such that $R^*(t) \geq R(s) + \beta(t - s)$ [2, Sec. 3.2.2]. We say that system S offers a *strict* service curve $\beta \in \mathcal{F}$ to the flow if $R^*(t) - R^*(s) \geq \beta(t - s)$ whenever $(s, t]$ is a backlogged period (i.e., $R^*(\tau) > R(\tau)$ for all τ such that $s < \tau \leq t$). If β is a strict service curve, then it is a service curve, but the converse is not always true [4, Section 1.3]. A frequently used service curve is the rate-latency function $\beta_{r,T}$ that is the function in \mathcal{F} defined by $\beta_{r,T}(t) = r[t - T]^+$, where we use the notation $[x]^+ = \max\{x, 0\}$. Saying that a system offers a service curve $\beta_{r,T}$ to a flow expresses that the flow is guaranteed a service rate r , except for possible interruptions that might impact the delay by at most T . Saying that a system offers a *strict* service curve $\beta_{r,T}$ to a flow expresses that the flow is guaranteed a service rate r , except for possible interruptions that might not exceed T in total per backlogged

period. A strict service curve β can always be assumed to be super-additive, i.e., to satisfy $\beta(s + t) \geq \beta(s) + \beta(t)$ for all s, t (otherwise, it can be replaced by its super-additive closure [2, Prop. 5.6]).

Assume that a flow, constrained by arrival curve α , traverses a system that offers a service curve β to the flow and that respects the ordering of the flow (FIFO per-flow). The delay of the flow is upper bounded by $h(\alpha, \beta)$ (horizontal deviation), defined by

$$h(\alpha, \beta) = \sup_{t \geq 0} \{ \inf_{d \geq 0} \{ \alpha(t) \leq \beta(t + d) \} \} \quad (1)$$

Our technique of proof uses the lower pseudo-inverse. The lower pseudo-inverse f^\downarrow of a function $f \in \mathcal{F}$ is defined by

$$f^\downarrow(y) = \inf \{ x | f(x) \geq y \} = \sup \{ x | f(x) < y \} \quad (2)$$

We use the following property from [6, Sec. 10.1]:

$$\forall x, y \in \mathbb{R}^+, y \leq f(x) \Rightarrow x \geq f^\downarrow(y) \quad (3)$$

III. SYSTEM MODEL

We consider a weighted round-robin subsystem that serves n input flows, has one queue per flow, and uses a weighted round-robin algorithm (described later) to arbitrate between flows. The weighted round-robin subsystem is itself placed in a larger system, and can compete with other queuing subsystems. For example, consider the case of a constant-rate server with several priority levels, without preemption, and where the weighted round-robin subsystem is at a priority level that is not the highest. Assuming some arrival curve constraints for the higher priority traffic, the service received by the entire weighted round-robin subsystem can be modelled using a strict service curve [2, Section 8.3.2].

This motivates us to assume that the aggregate of all flows in the weighted round-robin subsystem receives a strict service curve, say $\beta \in \mathcal{F}$ that we call “aggregate strict service curve”. If the weighted round-robin subsystem has exclusive access to a transmission line of rate c , then $\beta(t) = ct$ for $t \geq 0$. We assume that $\beta(t)$ is finite for every (finite) t and, without loss of generality, we assume β to be super-additive. Furthermore, we need an additional technical assumption, primarily for establishing the tightness result: we assume that β is Lipschitz-continuous, i.e., there exists a constant $K > 0$ such that $\frac{\beta(t) - \beta(s)}{t - s} \leq K$ for all $0 \leq s < t$; this does not appear to be a restriction as the rate at which data is served has a physical limit.

The arbitration algorithm assumed in this paper is IWRR, shown in Algorithm 1. When a packet of flow i enters the weighted round-robin subsystem, it is put into queue i . The weight of flow i is w_i . IWRR runs an infinite loop of *rounds*. In one round, each queue i has w_i emission opportunities; one packet can be sent during one emission opportunity. The inner loop defines a *cycle*, where each queue is visited but only those with a weight not smaller than the cycle number have an emission opportunity. The `send` instruction is assumed to be the only one with a non-null duration. Its actual duration depends on the packet size but also on the amount of service

available to the entire weighted round-robin subsystem. See Figure 2 for an illustration.

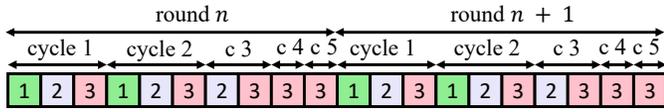


Fig. 2: Emission opportunities on two successive rounds for IWRR with three flows and $w_1 = 2, w_2 = 3, w_3 = 5$. Mind that this is not the temporal behaviour: each opportunity can lead to an empty interval if the queue is empty at this time. Furthermore, the duration of each non-empty interval depends on the packet size and the aggregate service available (we do not assume constant rate service).

Algorithm 1 Interleaved Weighted Round-Robin

Input: Integer weights $w_1 \leq w_2 \leq \dots \leq w_n$

```

1:  $w_{\max} = \max\{w_1, \dots, w_n\}$ 
2: while True do ▷ A round starts.
3:   for  $C \leftarrow 1$  to  $w_{\max}$  do ▷ A cycle starts.
4:     for  $i \leftarrow 1$  to  $n$  do
5:       if  $C \leq w_i$  then
6:         if (not empty( $i$ )) then
7:           ▷ A service for queue  $i$ .
8:             print(now, $i$ );
9:             send(head( $i$ ));
10:            removeHead( $i$ );
11:          end if
12:        end if
13:      end for
14:    end for ▷ A cycle finishes.
15:  end while ▷ A round finishes.

```

Here, we use the context of communication networks, but the results equally apply to real-time systems: Simply map flow to task, packet to job, packet size to job execution time and strict service curve to “delivery curve” [7], [8].

IV. STATE OF THE ART

One of the first use of round-robin scheduling in the network context appeared in [9], with a fairness objective, i.e., a fair way to share the bandwidth between sessions. It is also mentioned in [10] as a way to implement “fair queueing”.

The term “Weighed Round Robin” was coined in [1] as a generalisation of round-robin to share the bandwidth “in proportion to prescribed weights” in the context of ATM (i.e., with constant size packets). Two versions of the algorithm are presented in [1]. The former is presented in Algorithm 1: at cycle C (with C between 1 and w_{\max}), only flows with weight $w_i \geq C$ can emit one packet. We call this version IWRR. The latter version assumes that there exists for each flow i a bit-list of length w_{\max} , $o_i \in \{0, 1\}^{w_{\max}}$, such that $w_i = \sum_{k=1}^{w_{\max}} o_i[k]$. A flow i can emit a packet at cycle C only if $o_i[C] = 1$. A strategy is given to build these vectors in [1] and is refined with fairness objectives in [11]. Call LIWRR (list-based IWRR) this version.

IWRR is modified into WRR/SB in [12] to enable some flow to send slightly more packets than permitted in a cycle, and to decrease accordingly at the next cycle.

As mentioned in Section I, plain WRR (which we simply call “WRR”) enables each flow i to send up to w_i packets every time it is selected [13]. A “Multiclass WRR” is also defined in [13]. Surprisingly, the authors of [13] were not aware of [1] and have re-invented LIWRR. Note that even if WRR was designed for packets of constant size, it has been applied in network of variable size packets such as Ethernet [14, Sec. 8.6, Sec. 8.6.8.3, Sec. 37], in request balancing in cloud infrastructures [15], in the LinuxVirtualServer scheduling [16], in network of chip [17], and so on. In fact, looking for expression “weighted round-robin” in the title or abstracts of papers index by Scopus returns more than 400 entries (March 2020), and Google references more than 4000 patents with this expression (March 2020). Unfortunately, when authors refer to WRR, they often do not explicit which version of WRR it is.

A WRR server is also a latency-rate server, with latency and rates given in [18] for packets of constant size. The latency result is generalised to LIWRR in [19]. Even if the notion of latency-rate server is very close to the one of a service curve $\beta_{r,T}$ in network calculus, both notions are slightly different, and results cannot be directly imported from one theory to the other [20]. In [17], the authors consider a Network on Chip (NoC), with WRR arbitration at the flit level. A flit is the elementary data unit of the NoC, one flit is sent per CPU/NoC cycle. Assuming that the weights are such that packets are never fragmented by the arbiter, a strict service curve β_{R_i, T_i} for flow i is found, with $R_i = \frac{w_i}{\sum_k w_k}$, $T_i = \sum_{j \neq i} w_j$.

WRR arbitration in an Ethernet switch is also considered in [21], with the assumption that all flows of an output ports have the same constant packet size. It then computes, in the network calculus framework, a residual service with service curve β_{R_i, T_i} with $R_i = \frac{w_i}{\sum_k w_k} C$, $T_i = \frac{\sum_{j \neq i} w_j}{C}$, where C is the link rate. We assume that the missing packet size in the T_i term was a typo. This network calculus result on conventional WRR arbitration in Ethernet is refined in [22], considering packets of variable size, leading to residual service with strict service curve β_{R_i, T_i} with $R_i = \frac{w_i l_i^{\min}}{w_i l_i^{\min} + \sum_{j \neq i} w_j l_j^{\max}} C$ and $T_i = \frac{\sum_{j \neq i} w_j l_j^{\max}}{C}$ (cf. eq. (1) and (2) in [22]) where l_i^{\min} , l_i^{\max} are, respectively, lower and upper bounds on the size of the packets in the flow i . It refines this result by subtracting the part of the bandwidth not used by interfering flows (considering their arrival curves).

Observe that computing a residual service with a $\beta_{R,T}$ curve is pessimistic as it assumes that, once the worst latency is payed, each packet is served with the long-term residual rate. Whereas, in reality, each packet, when it is selected for emission, is transmitted at full link speed up to completion. A residual service for the conventional WRR with a curve that is an alternation of full services and plateaus is given in [2, Sec. 8.2.4]. This effect of “full speed up to completion” can also be captured when computing the local delay of a server

with $\beta_{R,T}$ service curve [23].

V. STRICT SERVICE CURVES FOR IWRR

Our first result is a strict service curve for IWRR that, as we show in Section VII, is the best possible. We compare it to WRR and also give simpler, lower approximations.

Theorem 1 (Strict Service Curve of IWRR). *Let S be a server shared by n flows that uses IWRR as explained in Section III, with weight w_i for flow i . Recall that the server offers a strict service curve β to the aggregate of the n flows. For any flow i , l_i^{\min} [resp. l_i^{\max}] is a lower [resp. upper] bound on the packet size.*

Then S offers to every flow i a strict service curve β_i given by $\beta_i(t) = \gamma_i(\beta(t))$ with

$$\gamma_i = \lambda_1 \otimes U_i \quad (4)$$

$$U_i(x) \stackrel{\text{def}}{=} \sum_{k=0}^{w_i-1} \nu_{l_i^{\min}, L_{\text{tot}}} \left([x - \psi_i(k l_i^{\min})]^+ \right) \quad (5)$$

$$L_{\text{tot}} = w_i l_i^{\min} + \sum_{j, j \neq i} w_j l_j^{\max} \quad (6)$$

$$\psi_i(x) \stackrel{\text{def}}{=} x + \sum_{j, j \neq i} \phi_{i,j} \left(\left\lfloor \frac{x}{l_i^{\min}} \right\rfloor \right) l_j^{\max} \quad (7)$$

$$\begin{aligned} \phi_{i,j}(x) &\stackrel{\text{def}}{=} \left\lfloor \frac{x}{w_i} \right\rfloor w_j + [w_j - w_i]^+ \\ &\quad + \min(x \bmod w_i + 1, w_j) \end{aligned} \quad (8)$$

In the above, $\nu_{a,b}$ is the stair function, λ_1 is the unit rate function and \otimes is the min-plus convolution, all are described in Fig. 1.

Furthermore, β_i is super-additive.

The proof is in Section VI. See Fig. 3 for some illustration of β_i . Observe that γ_i in (4) is the strict service curve obtained when the aggregate strict service curve is $\beta = \lambda_1$ (i.e., when the aggregate is served at a constant, unit rate). In the common case where β is equal to a rate-latency function, say $\beta_{c,T}$, we have $\beta_i(t) = \gamma_i(c(t-T))$ for $t \geq T$ and $\beta_i(t) = 0$ for $t \leq T$, namely, β_i is derived from γ_i by a rescaling of the x axis and a right-shift.

As mentioned in Section II, any strict service curve that is not super-additive can be replaced by its super-additive closure. The last statement in the theorem guarantees that this does not occur here.

We now compare to WRR. The best known service curve for (non-interleaved) WRR is given in [2, Sec. 8.2.4] and is

$$\beta'_i(t) = (\lambda_1 \otimes \nu_{q_i, L_{\text{tot}}}) \left([\beta(t) - Q_i]^+ \right) \quad (9)$$

with $q_i = w_i l_i^{\min}$ and $Q_i = \sum_{j, j \neq i} w_j l_j^{\max}$. In Section VII, we show that $\beta'_i(t)$ is indeed the best possible strict service curve for WRR. Furthermore, it is dominated by the strict service curve for IWRR:

Theorem 2. *With the assumptions in Theorem 1 and in (9):*

$$\beta'_i \leq \beta_i \quad (10)$$

The proof is in [3]. Fig. 3 illustrates how the strict service curve for IWRR improves on that for WRR, by providing a smoother, and generally larger, service.

The service curve found in Theorem 1 is the best possible one but has a complex expression. If there is interest in a simpler expression, any lower bounding function is a strict service curve; in particular, the strict service curve β'_i for WRR is also a valid, though suboptimal, strict service curve for IWRR. There is often interest in service curves that are rate-latency functions. Observe that, if the aggregate service curve β is a rate-latency function, then replacing γ_i by a rate-latency lower-bounding function also yields a rate-latency function for β_i , and vice-versa. Therefore, we are interested in rate-latency functions that lower bound γ_i .

Among all of such these, there is not a single best one, as some have a smaller latency while others have a larger rate. We say that a rate-latency function $\beta_{r,T}$ that lower bounds γ_i is non-dominated if there is no other rate latency function $\beta_{r',T'}$ that lower bounds γ_i and dominates $\beta_{r,T}$, i.e., such that $r' \geq r$ and $T' \leq T$. The following result gives all such non-dominated rate-latency functions. Let $r = \frac{q_i}{L_{\text{tot}}} = \frac{w_i l_i^{\min}}{L_{\text{tot}}}$, $r_{w_i-1} = 1$, and

$$r_k = \frac{l_i^{\min}}{\psi_i((k+1)l_i^{\min}) - \psi_i(k l_i^{\min})}, \quad 0 \leq k < w_i - 1 \quad (11)$$

$$k^* = \min\{0 \leq k < w_i \mid r_k \geq r\} \quad (12)$$

$$r_k^* = \min(r_k, r), \quad 0 \leq k \leq k^* \quad (13)$$

$$T_k^* = \psi_i(k l_i^{\min}) - \frac{k l_i^{\min}}{r_k^*}, \quad 0 \leq k \leq k^* \quad (14)$$

Theorem 3. *With the assumptions in Theorem 1 and the definitions (11)-(14), a rate-latency function $\beta_{r,T}$ lower bounds γ_i and is non-dominated if and only if $(r, T) = (r_k^*, T_k^*)$ for some integer k with $0 \leq k \leq k^*$. Among all such rate-latency functions, the one with lowest latency is $\beta_{r_0^*, T_0^*}$ and the one with largest rate is $\beta_{r_{k^*}^*, T_{k^*}^*}$.*

The proof is in [3]. Fig 3 illustrates $\beta_{r_0^*, T_0^*}$ and $\beta_{r_{k^*}^*, T_{k^*}^*}$ in some examples. Observe that $k \mapsto r_k^*$ is wide-sense increasing with k for $0 \leq k \leq k^*$, but the values of $(r_{k^*}^*, T_{k^*}^*)$ are not necessarily all distinct. It can also happen that $k^* = 0$ (as in the top panel of Fig. 3), in which case there is one optimal rate-latency service curve. In general, however, this does not happen and a simple lower bounding approximation can be obtained with $\max(\beta_{r_0^*, T_0^*}, \dots, \beta_{r_{k^*}^*, T_{k^*}^*})$. When β is a rate-latency function, this provides a convex piecewise linear function, which has several good properties [2, Sec. 4.2].

VI. PROOF OF THEOREM 1

The idea of proof is as follows. We consider a backlogged period $(s, t]$ of flow of interest i , and let p be the number of packets of flow i that are entirely served during this period. For every other flow j , the number of packets that are entirely served is upper bounded by a function of p , given in Lemma 3. Also, p is upper bounded by a function of the amount of service received by flow i in Lemma 5. Combining these two

Illustration of Theorems 1, 2, and 3,

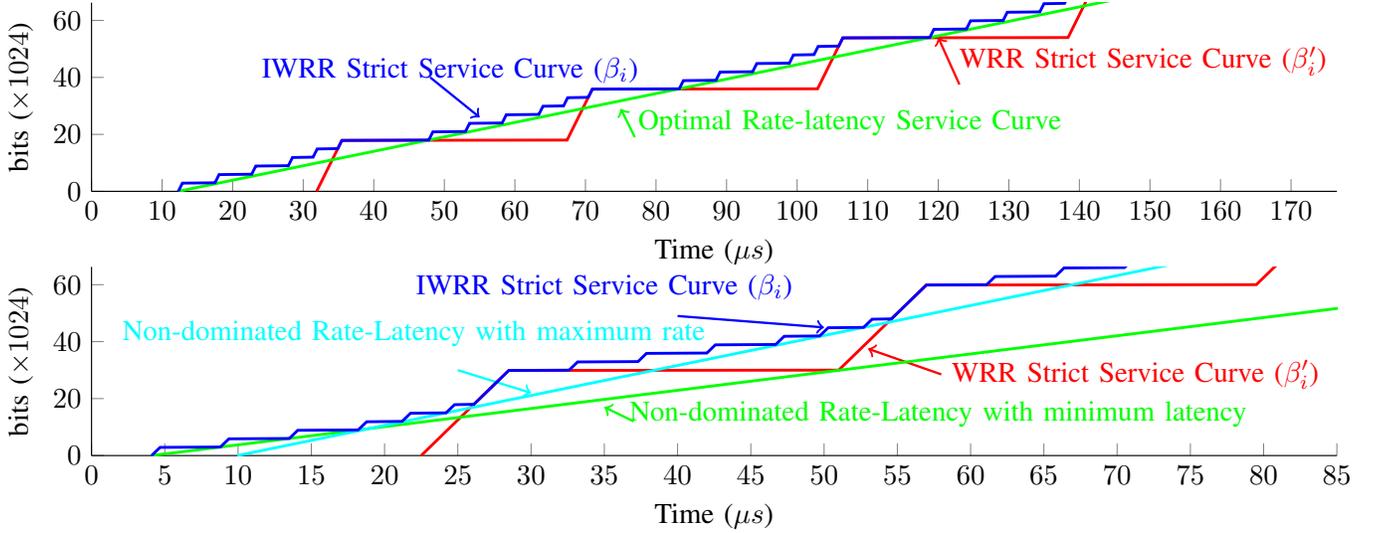
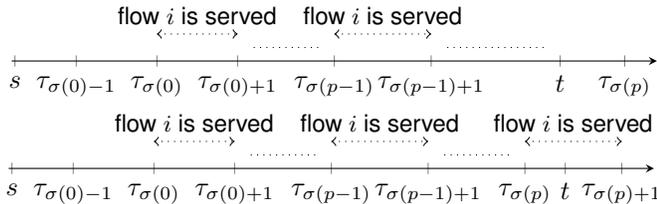


Fig. 3: Strict service curves obtained in Section V for an example with four input flows, weights = $\{4, 6, 7, 10\}$, $l^{\min} = \{4096, 3072, 4608, 3072\}$ bits, $l^{\max} = \{8704, 5632, 6656, 8192\}$ bits and $\beta(t) = ct$ with $c = 10\text{Mb/s}$ (i.e., the aggregate of all flows is served at a constant rate). The figure shows the IWRR service curve β_i and the WRR strict service curve β'_i for two of the flows; it also shows the non-dominated rate-latency strict service curves $\beta_{r_0^*, T_0^*}$ and $\beta_{r_{k^*}^*, T_{k^*}^*}$ of Theorem 3 (in the top panel both are equal).

gives an implicit inequality for the total amount of service (27). By using the technique of pseudo-inverse, this inequality is inverted and provides a lower bound for the amount of service received by the flow of interest.

A. Key Variables and Basic Properties

Let $(s, t]$ be a backlogged period of flow i . Let (τ_k, fl_k) be couples of (instant, flow), printed at line 8 of Algorithm 1. Note that $\tau_k < \tau_{k+1}$ as the send instruction has a non-null duration (because the aggregate service curve β is Lipschitz continuous). Let $\sigma(0), \sigma(1), \dots$ be the sequence of service opportunities for flow i at or after s , i.e., $\sigma(0) = \min\{m \mid \tau_m \geq s, fl_m = i\}$ and $\sigma(k) = \min\{m \mid \tau_m > \tau_{\sigma(k-1)}, fl_m = i\}$. The k th service opportunity for flow i occurs at time $\tau_{\sigma(k-1)}$; we say that it is “complete” if $\tau_{\sigma(k-1)+1} \leq t$, i.e., the interval taken by this service is entirely in $[s, t]$. Let $p \geq 0$ be the number of complete service opportunities. Observe that it is possible that $p = 0$, and it might happen that $\tau_{\sigma(p)} < t$ or $\tau_{\sigma(p)} \geq t$.



In each service of flow i , during a backlogged period, it sends one packet with a length $\geq l_i^{\min}$, thus, for all $k = 0 \dots p$, $R_i^*(\tau_{\sigma(k+1)}) - R_i^*(\tau_{\sigma(k)}) \geq l_i^{\min}$, therefore

$$R_i^*(\tau_{\sigma(p)}) - R_i^*(\tau_{\sigma(0)}) \geq pl_i^{\min} \quad (15)$$

B. Amount of Service to Other Flows

In order to upper bound the number of emission opportunities for another flow j , we first find an expression, in Lemma 1, for the number of emission opportunities for flow j between two consecutive emission opportunities for flow i . Lemma 2 then finds an upper bound on the number of emission opportunities for flow j in $(s, \tau_{\sigma(p)})$, as a function of the cycle number (variable C in Algorithm 1) at $\tau_{\sigma(0)}$. Lastly, Lemma 3 maximizes the previous upper bound over all values of C .

Lemma 1. *The number of emission opportunities for flow $j \neq i$ between two consecutive emission opportunities for flow i , given that the latter emission opportunity for flow i occurs at cycle C , is equal to $q_{i,j}(C) \stackrel{\text{def}}{=}$*

$$\begin{cases} 0 & \text{if } 1 < C \leq w_i \text{ and } w_j < C \\ 1 & \text{if } 1 < C \leq w_i \text{ and } w_j \geq C \\ [w_j - w_i]^+ + 1 & \text{if } C = 1 \end{cases} \quad (16)$$

Proof. According to Algorithm 1, flow i has emission opportunities only in the first w_i cycles of each round. Both emission opportunities are either in the same round (Case 1) or in two consecutive rounds (Case 2). As C is the cycle number for the second emission opportunity for flow i , Case 1 can occur only when $1 < C \leq w_i$, and Case 2 can occur when $C = 1$. For Case 1, we further differentiate between $w_j < C$ and $w_j \geq C$.

Case 1a: $1 < C \leq w_i$ and $w_j < C$: Queue j does not have an emission opportunity in cycle C because $w_j < C$. Also, we must have $w_j < w_i$, thus queue j does not have an emission opportunity after i in cycle $C - 1$. Hence, $q_{i,j}(C) = 0$.

Case 1b: $1 < C \leq w_i$ and $w_j \geq C$: If $w_j > w_i$, then queue j has an emission opportunity after queue i in cycle $C - 1$. If

$w_j = w_i$, then queue j has an emission opportunity before i in cycle C , or after i in cycle $C - 1$. Else, $C \leq w_j < w_i$ and queue j has an emission opportunity in cycle C , before i . In all cases, $q_{i,j}(C) = 1$.

Case 2: $C = 1$: The first emission opportunity for i is in the last cycle of a round that includes i (cycle w_i). If $w_j > w_i$, then queue j has an emission opportunity in the rest of cycle w_i and also has emission opportunities during the next $(w_j - w_i)$ cycles of the last round. In this case, $q_{i,j}(C) = w_j - w_i + 1$, which is also the value in the last line of (16). Else if $w_j = w_i$, queue j has an emission opportunity before i in this cycle or after i in cycle w_i of the first round, thus $q_{i,j}(C) = 1$, which is also the value in the last line of (16). Else, $w_j < w_i$ and queue j has an emission opportunity before i in this cycle. Here too, $q_{i,j}(C) = 1$, the value in the last line of (16). \square

Lemma 2. *The number of emission opportunities for flow $j \neq i$ in $(s, \tau_{\sigma(p)})$, for any backlogged period $(s, t]$ of flow i with p complete services, given that the first service starts at cycle number C (cycle number at time $\tau_{\sigma(0)}$) is upper bounded by*

$$q'_{i,j}(C, p) \stackrel{\text{def}}{=} \sum_{k=0}^p q_{i,j}((C + k - 1) \bmod w_i + 1) \quad (17)$$

Also, let $C'(p)$ be the cycle number at $\tau_{\sigma(p)}$. Then,

$$C'(p) = (C + p - 1) \bmod w_i + 1 \quad (18)$$

Proof. By induction on p .

Base Case: $p = 0$

In this case, $q'_{i,j}(C, 0)$ is the number of emission opportunities for flow j between two consecutive emission opportunities for flow i that by Lemma 1, is equal to $q_{i,j}(C)$. As $1 \leq C \leq w_i$, $(C - 1) \bmod w_i + 1 = C$ thus $q_{i,j}(C) = q_{i,j}((C - 1) \bmod w_i + 1)$. This shows (17). Also, by definition, $C'(0) = C$; using again $(C - 1) \bmod w_i + 1 = C$ shows that (18) holds.

Induction step:

We assume that (17) and (18) hold for $p - 1$, and we want to show that they also hold for p .

First, let's prove (18). There are two possible cases: (a) if $0 \leq C'(p - 1) < w_i$, then both $(p - 1)$ th and p th emission opportunities occur in the same round, thus $C'(p) = C'(p - 1) + 1$. By the induction hypothesis, $(C + p - 2) \bmod w_i + 1 < w_i$, i.e., $(C + p - 2) \bmod w_i < w_i - 1$. Note that, for any integer x

$$(x+1) \bmod w = \begin{cases} (x \bmod w) + 1 & \text{if } (x \bmod w) < w - 1 \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

By using (19), we obtain that $C'(p)$ is given by (18) as required. (b) In the second case, $C'(p - 1) = w_i$ then the next emission opportunity occurs in the first cycle of the next round, thus $C'(p) = 1$. Here too, applying (19) shows that $C'(p)$ is given by (18) as required.

Then, we prove (17). Let N be the number of emission opportunities for flow j in $[s, \tau_{\sigma(p)})$. N is the sum of N_1 , the number of emission opportunities in $[s, \tau_{\sigma(p-1)})$, and N_2 , the

number of emission opportunities in $(\tau_{\sigma(p-1)}, \tau_{\sigma(p)})$. By the induction hypothesis, $N_1 \leq q'_{i,j}(C, p - 1)$. Also, by Lemma 1, we have $N_2 \leq q_{i,j}(C'(p))$. Thus, by using (18) which was just shown to also hold for p , we obtain

$$N \leq \sum_{k=0}^{p-1} q_{i,j}((C + k - 1) \bmod w_i + 1) + q_{i,j}((C + p - 1) \bmod w_i + 1) \quad (20)$$

where the right-hand side is equal to $q'_{i,j}(C, p)$ as required. \square

Lemma 3. *For any backlogged period $(s, t]$ of flow i with p complete services, the number of emission opportunities for flow $j \neq i$ in $(s, \tau_{\sigma(p)})$ is upper bounded by $\phi_{i,j}(p)$, defined in (8).*

Proof. Lemma 2 gives the number of emission opportunities for flow $j \neq i$ in $(s, \tau_{\sigma(p)})$, for any backlogged period $(s, t]$ of flow i with p complete services, when the first service starts at cycle number C (cycle number at time $\tau_{\sigma(0)}$). To obtain the lemma, we maximize this result over C . We show the following properties.

(P1) For any integer $C \in [1, w_i]$,

$$\sum_{k=0}^{w_i-1} q_{i,j}((C + k - 1) \bmod w_i + 1) = w_j \quad (21)$$

The mapping $k \mapsto (C + k - 1) \bmod w_i + 1$ is one-to-one from $\{0, \dots, w_i - 1\}$ onto $\{1, \dots, w_i\}$, thus the left-hand side of (21) is equal to $\sum_{k=1}^{w_i} q_{i,j}(k)$ that as we show now, is equal to w_j . First, we have $q_{i,j}(1) = [w_j - w_i]^+ + 1$. Also, $q_{i,j}(k) = 1$ when $k > 1$ and $w_j \geq k + 1$. Thus, $\sum_{k=2}^{w_i} q_{i,j}(k) = \min(w_i - 1, w_j - 1)$ and finally the left-hand side is equal to $[w_j - w_i]^+ + \min(w_i - 1, w_j - 1) + 1$, which is equal to w_j .

(P2) For any integers $C \in [1, w_i]$ and $p \geq 0$, $q'_{i,j}(C, p) =$

$$\left\lfloor \frac{p}{w_i} \right\rfloor w_j + \sum_{k=0}^{p \bmod w_i} q_{i,j}((C + k - 1) \bmod w_i + 1) \quad (22)$$

$q_{i,j}$ is a periodic function with period w_i . By (P1), the sum over one complete period is w_j . Also, we can write $p = \left\lfloor \frac{p}{w_i} \right\rfloor w_i + p \bmod w_i$. Thus, we have $\left\lfloor \frac{p}{w_i} \right\rfloor$ complete rounds, and the sum in (22) is the remainder.

(P3) $q_{i,j}$ is a wide-sense decreasing function. This means that for any integer $k \in [1, w_i]$, $q_{i,j}(k + 1) \leq q_{i,j}(k)$. If $k = 1$, this follows from $q_{i,j}(1) \geq 1$ and $q_{i,j}(2) \leq 1$. Else if $k \leq w_j < k + 1$, then $q_{i,j}(k + 1) = 0$ and $q_{i,j}(k) = 1$. Else, they are equal. Hence, in all cases the property holds.

(P4) For any integer $C \in [1, w_i]$ and $p \geq 0$,

$$q'_{i,j}(C, p) \leq q'_{i,j}(1, p) \quad (23)$$

By using (P2), we should show that $\sum_{k=0}^{p \bmod w_i} q_{i,j}((C + k - 1) \bmod w_i + 1)$ is upper bounded by $\sum_{k=0}^{p \bmod w_i} q_{i,j}(k \bmod w_i + 1)$. Note that here we have $k \bmod w_i = k$. Both sides are the sum of $a \stackrel{\text{def}}{=} p \bmod w_i + 1$ unique elements of the set $\{q_{i,j}(k)\}_{k \in [1, w_i]}$. By (P3), the

right-hand side is the maximum sum of a unique elements of this set.

(P5) For any integer $p \geq 0$,

$$q'_{i,j}(1,p) = \phi_{i,j}(p) \quad (24)$$

We apply (P2) with $C = 1$ to compute $q'_{i,j}(1,p)$. Then, the sum in the right-hand side of (22) is equal to $\sum_{k=0}^{p \bmod w_i} q_{i,j}(k+1)$, as $k \bmod w_i = k$. Then, by using the same argument after (21), it is equal to $[w_j - w_i]^+ + 1 + \min(p \bmod w_i, w_j - 1)$, which, by (8), is precisely $\phi_{i,j}(p)$.

The lemma then follows directly from (P4) and (P5). \square

Lemma 4. For every flow $j \neq i$,

$$R_j^*(t) \leq R_j^*(\tau_{\sigma(p)}) \quad (25)$$

Proof. If $t \leq \tau_{\sigma(p)}$, the result follows from R_j^* being wide-sense increasing. Else, we have $t > \tau_{\sigma(p)}$; this implies that flow i is served during $[\tau_{\sigma(p)}, t]$; thus for any other flow j , $R_j^*(t) = R_j^*(\tau_{\sigma(p)})$. \square

C. Amount of Service to Flow of Interest

Lemma 5. The number of complete services, p , of flow of interest, i , in $(s, t]$ is upper bounded by:

$$p \leq \left\lfloor \frac{R_i^*(t) - R_i^*(s)}{l_i^{\min}} \right\rfloor \quad (26)$$

Proof. First, $R_i^*(s) \leq R_i^*(\tau_{\sigma(0)})$, as $s \leq \tau_{\sigma(0)}$ and R_i^* is wide-sense increasing. Second, consider the two cases in VI-A. If $t \geq \tau_{\sigma(p)}$, the property holds. Else, the scheduler in not serving flow i in $[\tau_{\sigma(p-1)+1}, \tau_{\sigma(p)})$, thus, $R_i^*(t) = R_i^*(\tau_{\sigma(p)})$. Hence, in both cases $R_i^*(t) \geq R_i^*(\tau_{\sigma(p)})$. By (15), $R_i^*(t) - R_i^*(s) \geq pl_i^{\min}$. Then, observe that p is integer. \square

D. Total Amount of Service

Lemma 6. For any backlogged period $(s, t]$ of the flow of interest i ,

$$\beta(t-s) \leq \psi_i(R_i^*(t) - R_i^*(s)) \quad (27)$$

where ψ_i is defined in (7).

Proof. As the interval $(s, t]$ is a backlogged period, by the definition of the strict service curve for the aggregate of flows, $\beta(t-s) \leq \sum_j R_j^*(t) - R_j^*(s)$. We upper bound $R_j^*(t)$ for all $j \neq i$ by applying Lemma 4,

$$\beta(t-s) \leq (R_i^*(t) - R_i^*(s)) + \sum_{j,j \neq i} R_j^*(\tau_{\sigma(p)}) - R_j^*(s) \quad (28)$$

Each flow j has at most $\phi_{i,j}(p)$ emission opportunities during $(s, \tau_{\sigma(p)})$ (Lemma 3) and can send at most one packet of maximum size in each. Thus,

$$\beta(t-s) \leq (R_i^*(t) - R_i^*(s)) + \sum_{j,j \neq i} \phi_{i,j}(p) l_j^{\max} \quad (29)$$

Also, Lemma 5 finds an upper bound on p . Thereby,

$$\begin{aligned} \beta(t-s) &\leq (R_i^*(t) - R_i^*(s)) \\ &+ \sum_{j,j \neq i} \phi_{i,j} \left(\left\lfloor \frac{R_i^*(t) - R_i^*(s)}{l_i^{\min}} \right\rfloor \right) l_j^{\max} \end{aligned} \quad (30)$$

where the right-hand side is equal to $\psi_i(R_i^*(t) - R_i^*(s))$. \square

E. Lower Pseudo-inverse of ψ_i

Our next step is to invert (27) by computing the lower-pseudo inverse of ψ_i . As the calculus of pseudo inverses applies to wide-sense increasing functions, we first show:

Lemma 7. ψ_i , defined in (7), is wide-sense increasing.

Proof. It is sufficient to show that $\phi_{i,j}$, defined in (8), is a wide-sense increasing function. For any non-negative integers x and y such that $y \leq x$, we can write $x = kw_i + (x \bmod w_i)$ and $y = k'w_i + (y \bmod w_i)$, where k, k' are non-negative integers. We must have $k \leq k'$. If $k = k'$, we know that $(y \bmod w_i \leq x \bmod w_i)$ and $\left\lfloor \frac{x}{w_i} \right\rfloor = \left\lfloor \frac{y}{w_i} \right\rfloor$. Hence, $\phi_{i,j}(y) \leq \phi_{i,j}(x)$. Else, $k > k'$ and $\left\lfloor \frac{x}{w_i} \right\rfloor > \left\lfloor \frac{y}{w_i} \right\rfloor$. Thereby, $\phi_{i,j}(x)$ is at least one w_j larger than $\phi_{i,j}(y)$. Hence, $\phi_{i,j}(y) < \phi_{i,j}(x)$. \square

Lemma 8. Let $g_0, g_1, \dots, g_k, \dots$ be a non-negative sequence such that $g_{k+1} - g_k \geq 1$. The sequence can be extended to a function in \mathcal{F} by $g(x) = g_{\lfloor x \rfloor}$ and let g^\downarrow be its lower pseudo-inverse, so that $g^\downarrow(y) = k \in \mathbb{N} \Leftrightarrow g_k \leq y < g_{k+1}$. Define $f \in \mathcal{F}$ by $f(x) = g_{\lfloor x \rfloor} + x \bmod 1$. Then, $f^\downarrow = \lambda_1 \otimes g^\downarrow$.

Proof. Observe that convolving g^\downarrow with λ_1 consists in smoothing the unit steps with a slope of 1 (Fig. 1). Thus $(\lambda_1 \otimes g^\downarrow)(y) = k + y - g_k$ whenever $g_k \leq y \leq g_{k+1}$ and $(\lambda_1 \otimes g^\downarrow)(y) = k + 1$ whenever $g_k + 1 \leq y \leq g_{k+1}$.

Also, f is piecewise linear and can be inverted in closed form on every interval where it is linear. A direct calculation gives $f^\downarrow(y) = k + y - g_k$ whenever $g_k \leq y \leq g_{k+1}$ and $f^\downarrow(y) = k + 1$ whenever $g_k + 1 \leq y \leq g_{k+1}$. \square

Lemma 9. Let $f \in \mathcal{F}$ and $l, m > 0$. Define $h \in \mathcal{F}$ by $h(x) = mf(\frac{x}{l})$. Then, for all $y \geq 0$, $h^\downarrow(y) = lf^\downarrow(\frac{y}{m})$.

Proof. Let $B(f, y) \stackrel{\text{def}}{=} \{x \geq 0, h(x) \geq y\}$ so that $f^\downarrow(y) = \inf B(y, f)$. Observe that $x \in B(h, y) \Leftrightarrow \frac{x}{l} \in B(f, \frac{y}{m})$. \square

Lemma 10. Let $a \in \mathcal{F}$ and $l > 0$. Define $b \in \mathcal{F}$ by $b(x) = lf(\frac{x}{l})$. Then, for all $x \geq 0$, $(\lambda_1 \otimes b)(x) = l(\lambda_1 \otimes a)(\frac{x}{l})$.

Proof. Do the change of variable $u = lv$ in the expansion $(\lambda_1 \otimes b)(x) = \inf_{0 \leq u \leq x} (u + b(x-u))$ and obtain $(\lambda_1 \otimes b)(x) = \inf_{0 \leq v \leq \frac{x}{l}} (lv + a(\frac{x}{l} - v)) = l(\lambda_1 \otimes a)(\frac{x}{l})$. \square

We can now compute the lower-pseudo inverse of ψ_i . First, define the sequence g by $g_k = \frac{1}{l_i^{\min}} \psi_i(kl_i^{\min})$. As in Lemma 8, g can be extended to a piecewise constant function whose lower-pseudo inverse, g^\downarrow , can be directly computed:

$$g^\downarrow(x) = \frac{1}{l_i^{\min}} \sum_{k=0}^{w_i-1} \nu_{l_i^{\min}, L_{\text{tot}}} \left(l_i^{\min} [x - g_k]^+ \right) \quad (31)$$

Second, observe that for all $x \geq 0$, $\psi_i(x) = \psi_i(\lfloor \frac{x}{l_i^{\min}} \rfloor l_i^{\min}) + x \bmod l_i^{\min}$. Define f and h from g as in Lemmas 8 and 9 with $l = m = l_i^{\min}$, so that $h = \psi_i$. Apply Lemmas 8 and 9 and obtain $\psi_i^\downarrow(x) = l_i^{\min} (\lambda_1 \otimes g^\downarrow)(\frac{x}{l_i^{\min}})$. Now apply Lemma 10 with $a = g^\downarrow$, $l = l_i^{\min}$, and $b = U_i$ to obtain

$$\psi_i^\downarrow = \lambda_1 \otimes U_i \quad (32)$$

F. Proof of Theorem 1

Proof. Lemma 6 gives, in (27), an upper bound on the total amount of service as a function of the service received by the flow of interest. We invert (27) by the lower-pseudo inverse technique in (3) and obtain $R_i^*(t) - R_i^*(s) \geq \psi_i^l(\beta(t-s))$. The lower-pseudo inverse of ψ_i is given by (32), thus

$$R_i^*(t) - R_i^*(s) \geq (\lambda_1 \otimes U_i)(\beta(t-s)) = \beta_i(t-s) \quad (33)$$

Lastly, we need to prove that β_i is super-additive. This follows from the tightness result in Theorem 4 (the proof of which is independent of rest of this proof). Indeed, the super-additive closure $\bar{\beta}_i$ of β_i is also a strict service curve, and $\bar{\beta}_i(t) \geq \beta_f(t)$ for all t [2, Prop. 5.6]. By Theorem 4, we also have $\bar{\beta}_i(t) \leq \beta_i(t)$ for all t , hence $\bar{\beta}_i = \beta_i$. \square

VII. TIGHTNESS

We first show that the strict service curve we have obtained is the best possible. Proofs of results in this Section are in [3].

A. Tightness of Strict Service Curve

Theorem 4 (Tightness of the IWRR Service Curve). *Consider a weighted round-robin subsystem that uses the IWRR scheduling algorithm, as defined in Section III. Assume the following system parameters are fixed: the number of input flows, the weight w_j allocated to every flow j , the bounds on packet sizes l_j^{\min} and l_j^{\max} for every flow j , and the strict service curve β for the aggregate of all flows. Let i be the index of one of the flows.*

Assume that $b_i \in \mathcal{F}$ is a strict service curve for flow i in any system that satisfies the specifications above. Then $b_i \leq \beta_i$ where β_i is given in Theorem 1.

Interestingly, we obtain a similar result for WRR. Recall that β'_i is the strict service curve for flow i , described in (9), which was obtained in [2, Sec. 8.2.4].

Theorem 5 (Tightness of the WRR Service Curve). *Theorem 4 is also valid if we replace IWRR with WRR. Specifically, using WRR as a scheduling policy, β'_i is the largest possible strict service curve for flow i .*

B. Tightness of Delay Bounds With Constant Packet Sizes

Having obtained the best-possible strict service curve does not guarantee that the delay bounds derived from it are tight, i.e., are worst-case delays. This is because a service curve is only an abstraction of the system; and we have obtained a strict service curve, and non-strict service curves might provide better results. However, we show that, for flows of packets of constant size, we do obtain tight delay bounds. We show that it holds for IWRR and for WRR.

Recall that a delay bound requires the knowledge of an arrival curve α_i for the flow of interest. If this flow generates only packets of length l , then α_i can be assumed to be a multiple of l and sub-additive. A delay bound for this flow is then equal to $h(\alpha_i, \beta_i)$ (see (1)).

Theorem 6 (Tightness of Delay Bound for IWRR with Constant Packet Size). *Consider a system, as in Theorem 4,*

with the additional assumption that, for the flow of interest i , $l_i^{\min} = l_i^{\max} = l$.

Let $\alpha_i \in \mathcal{F}$ be a sub-additive function that is an integer multiple of l , and assume that flow i has α_i as arrival curve. The network calculus delay bound is tight, i.e, there exists a trajectory where the delay of one packet of flow i is equal to $h(\alpha_i, \beta_i)$.

Theorem 7 (Tightness of Delay Bound for WRR with Constant Packet Size). *Theorem 6 is also valid for the WRR policy.*

VIII. NUMERICAL EXAMPLES

To compare IWRR and WRR worst-case delays, we provide some numerical examples. First, we consider a system of 8 input flows f_1, \dots, f_8 with respective weights $\{22, 27, 28, 30, 30, 34, 41, 45\}$ and $l^{\min} = l^{\max} = l = 7119$ bit. Let the aggregate service, β , be a constant bit rate of 10 Mb/s. For every flow i , we compute the IWRR and WRR strict service curves β_i, β'_i . Then, for every i , we generate $N = 1000$ leaky-bucket arrival curves γ_{r, b_k} , $k = 1 \dots N$, with rate $r = 0.5$ Mb/s and burst b_k picked uniformly at random in $[1, 20]$ packets. Then, we use $\alpha_i^k = \lceil \frac{\gamma_{r, b_k}}{l} \rceil l$ to satisfy the conditions of Theorems 6 and 7 and to compute $d_i^k = h(\alpha_i^k, \beta_i)$ and $d_i^k = h(\alpha_i^k, \beta'_i)$. Fig. 4 gives the box-and-whisker plots of the $d_i^k - d_i^k$ series. The median of WRR delay bounds d_i^k are also provided to illustrate the improvement.

Second, we repeated the same study for $M = 10000$ sets of system parameters. For each system, we choose the weights of 8 flows by picking them uniformly at random between 10 and 50, and we pick a packet length l uniformly at random between 64 to 1522 bytes. For each experiment, we call flow 1 the flow with the smallest weight, flow 2 with second smallest weight, and so on. As the scale of delay bounds depends on the choices of weights and the packet length, the $d_i^k - d_i^k$ series are divided by d_i^m , the median of WRR delay bounds for flow i . Fig. 4 gives the box-and-whisker plots of the $\frac{d_i^k - d_i^k}{d_i^m}$ series. Using IWRR improves worst-case delays, as expected, and the improvement is larger for flows with larger weights.

IX. CONCLUSION

IWRR is a variant of WRR with the same long-term rate and the same complexity. We have provided a residual strict service curve for IWRR and have showed that it is the best possible one under general assumptions. For flows with packets of constant size, we have showed that the delay bounds derived from it are worst-case. We have proved that IWRR worst-case delay is not greater than WRR and shown on experiments that the gain is significant (20%-60%) in practice, which speaks in favour of using IWRR as a replacement to WRR. Our result assumes that the aggregate of all IWRR queues receives a strict service curve guarantee, and we find a strict service curve guarantee for every IWRR queue. Therefore, our results apply to hierarchical schedulers. In future research, we plan improve the results with supplementary hypotheses on flows, considering arrival curves and packet size distribution, with “packet curves” [24].

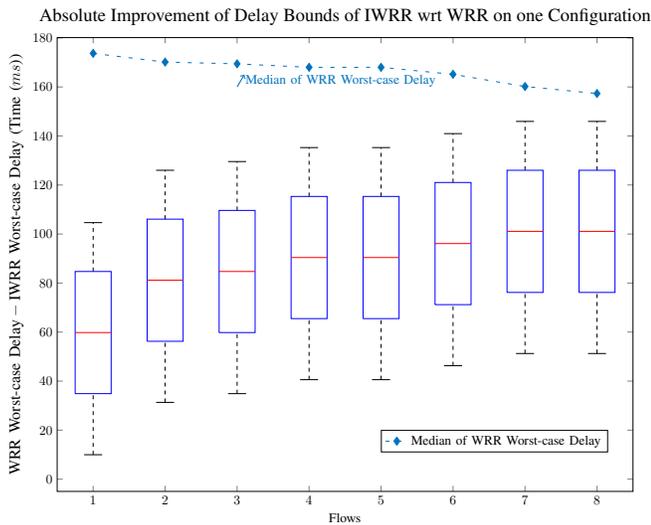


Fig. 4: Box-and-whisker plots of difference between WRR and IWRR delay bounds with weights $\{22, 27, 28, 30, 30, 34, 41, 45\}$ and $l = 7119$ bit with random arrival curves. Median WRR delay bounds are also provided.

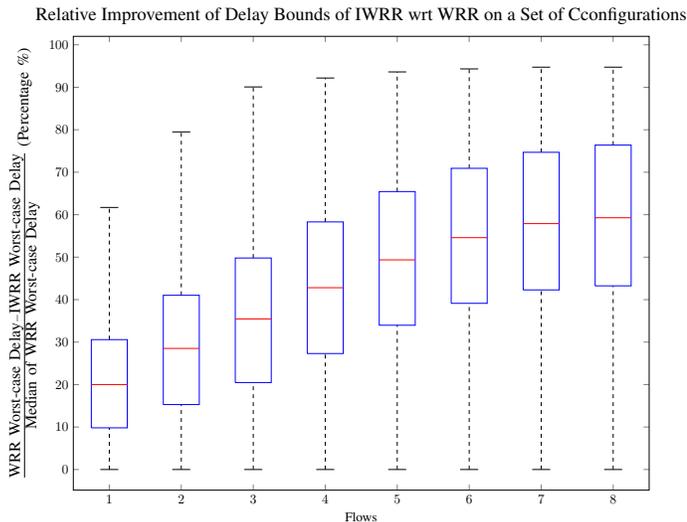


Fig. 5: Box-and-whisker plots of difference between WRR and IWRR delay bounds normalized to the median of WRR delay bounds, for several systems with weights picked uniformly at random in $[10, 50]$, assigned to flow by increasing order, and a packet length picked uniformly at random in $[64, 1522]$ bytes.

REFERENCES

- [1] M. Katevenis, S. Sidiropoulos, and C. Courcoubetis, "Weighted round-robin cell multiplexing in a general-purpose ATM switch chip," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 8, pp. 1265–1279, 1991.
- [2] A. Bouillard, M. Boyer, and E. Le Corronc, *Deterministic Network Calculus: From Theory to Practical Implementation*. Wiley-ISTE.
- [3] S. M. Tabatabaee, J.-Y. L. Boudec, and M. Boyer, "Interleaved weighted round-robin: A network calculus analysis," 2020. [Online]. Available: <https://arxiv.org/pdf/2003.08372.pdf>
- [4] J.-Y. Le Boudec and P. Thiran, *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*. Springer Science & Business Media, 2001, vol. 2050.
- [5] C. S. Chang, *Performance Guarantees in Communication Networks*. New York: Springer-Verlag, 2000.
- [6] J. Liebeherr, "Duality of the max-plus and min-plus network calculus," *Foundations and Trends in Networking*, vol. 11, no. 3-4, pp. 139–282, 2017.
- [7] D. B. Chokshi and P. Bhaduri, "Modeling fixed priority non-preemptive scheduling with real-time calculus," in *2008 14th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, Aug 2008, pp. 387–392.
- [8] L. Thiele, S. Chakraborty, and M. Naedele, "Real-time calculus for scheduling hard real-time systems," in *2000 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. 4, May 2000, pp. 101–104 vol.4.
- [9] E. L. Hahne and R. G. Gallager, "Round robin scheduling for fair flow control in data communication networks," in *Proc. of the IEEE Int. Conf. on Communications (ICC 86)*, June 1986.
- [10] J. Nagle, "On packet switches with infinite storage," *Communications, IEEE Transactions on*, vol. 35, no. 4, pp. 435–438, Apr 1987.
- [11] Y.-T. Wang, T.-P. Lin, and K.-C. Gan, "An improved scheduling algorithm for weighted round-robin cell multiplexing in an ATM switch," in *Proceedings of the International Conference on Communications (SUPERCOMM'94)*, May 1994, pp. 1032–1037 vol.2.
- [12] H. Shimonishi, M. Yoshida, R. Fan, and H. Suzuki, "An improvement of weighted round robin cell scheduling in ATM networks," in *Proc. of the IEEE Global Telecommunications Conference (GLOBECOM 97)*, vol. 2, Nov 1997, pp. 1119–1123 vol.2.
- [13] H. M. Chaskar and U. Madhoo, "Fair scheduling with tunable latency: A round robin approach," in *Proc of the IEEE Global Telecommunications Conference (GLOBECOM'99)*, vol. 2. IEEE, 1999, pp. 1328–1333.
- [14] "IEEE standard for local and metropolitan area networks – bridges and bridged networks," IEEE, IEEE Standard 802.1Q, 2018.
- [15] D. B. L.D. and P. V. Krishna, "Honey bee behavior inspired load balancing of tasks in cloud computing environments," *Applied Soft Computing*, vol. 13, no. 5, pp. 2292 – 2303, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1568494613000446>
- [16] Wensong, "Weighted round-robin scheduling, documentation of the linuxvirtualserver knowledge base," http://kb.linuxvirtualserver.org/wiki/Weighted_Round-Robin_Scheduling, 2005.
- [17] Y. Qian, Z. Lu, and W. Dou, "Analysis of worst-case delay bounds for best-effort communication in wormhole networks on chip," in *Proc. of the 3rd ACM/IEEE International Symposium on Networks-on-Chip (NoCS 2009)*. IEEE, 2009, pp. 44–53.
- [18] D. Stiliadis and A. Varma, "Latency-rate servers: A general model for analysis of traffic scheduling algorithms," *IEEE/ACM Trans. Netw.*, vol. 6, no. 5, pp. 611–624, October 1998. [Online]. Available: <http://dx.doi.org/10.1109/90.731196>
- [19] S. Nananukul, "Latency of weighted round-robin scheduler," *Electronics Letters*, vol. 39, no. 2, pp. 256–257, 2003.
- [20] Y. Jiang, "Relationship between guaranteed rate server and latency rate server," *Computer Networks*, vol. 43, no. 3, pp. 307 – 315, 2003. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128603002767>
- [21] J.-P. Georges, T. Divoux, and E. Rondeau, "Network calculus: application to switched real-time networking," in *Proc. of the 5th Int. ICST Conf. on Performance Evaluation Methodologies and Tools*, ser. VALUETOOLS '11. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011, pp. 399–407. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2151688.2151733>
- [22] A. Soni, X. Li, J.-L. Scharbarg, and C. Fraboul, "WCTT analysis of avionics switched ethernet network with WRR scheduling," in *Proc. of the 26th International Conference on Real-Time Networks and Systems (RTNS)*. ACM, 2018, pp. 213–222.
- [23] E. Mohammadpour, E. Stai, and J.-Y. Le Boudec, "Improved delay bound for a service curve element with known transmission rate," *IEEE Networking Letters*, pp. 1–1, 2019. [Online]. Available: <http://infoscience.epfl.ch/record/267840>
- [24] A. Bouillard, N. Farhi, and B. Gaujal, "Packetization and packet curves in network calculus," in *Performance Evaluation Methodologies and Tools (VALUETOOLS), 2012 6th International Conference on*. IEEE, 2012, pp. 136–137.