Interleaved Weighted Round-Robin: A Network Calculus Analysis

ITC 32, September 2020, Osaka, Japan

5th Workshop on Network Calculus (WoNeCa-5)
Virtual Event on October 9th, 2020

Hossein Tabatabaee (EPFL, Switzerland),

Jean-Yves Le Boudec (EPFL, Switzerland), and

Marc Boyer (ONERA, France)

Weighted Round-Robin (WRR)

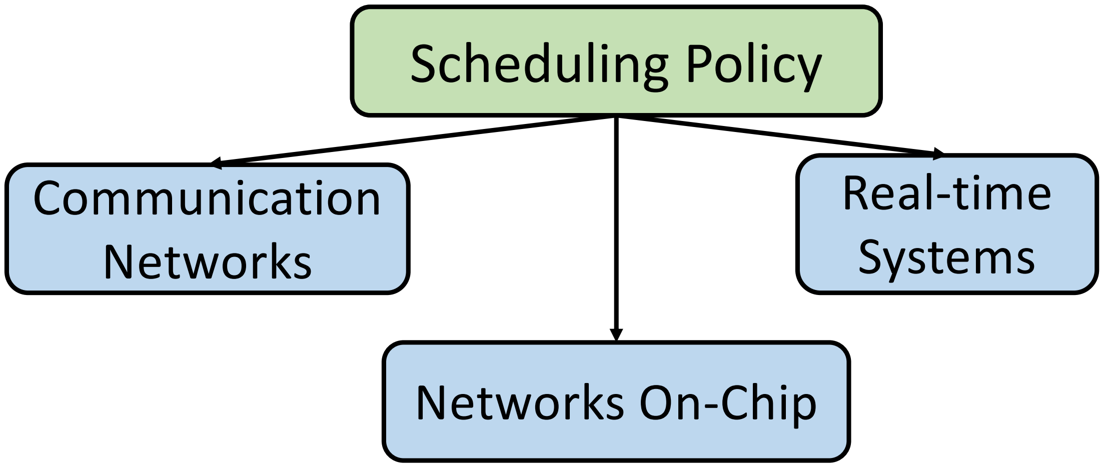

WRR Behavior

Each queue is given a weight.

↓

An infinite loop of rounds to visit queues.

↓

When a queue with weight w is visited.

↓

It can send up to w packets

↓

Simple

Example of WRR Scheduling

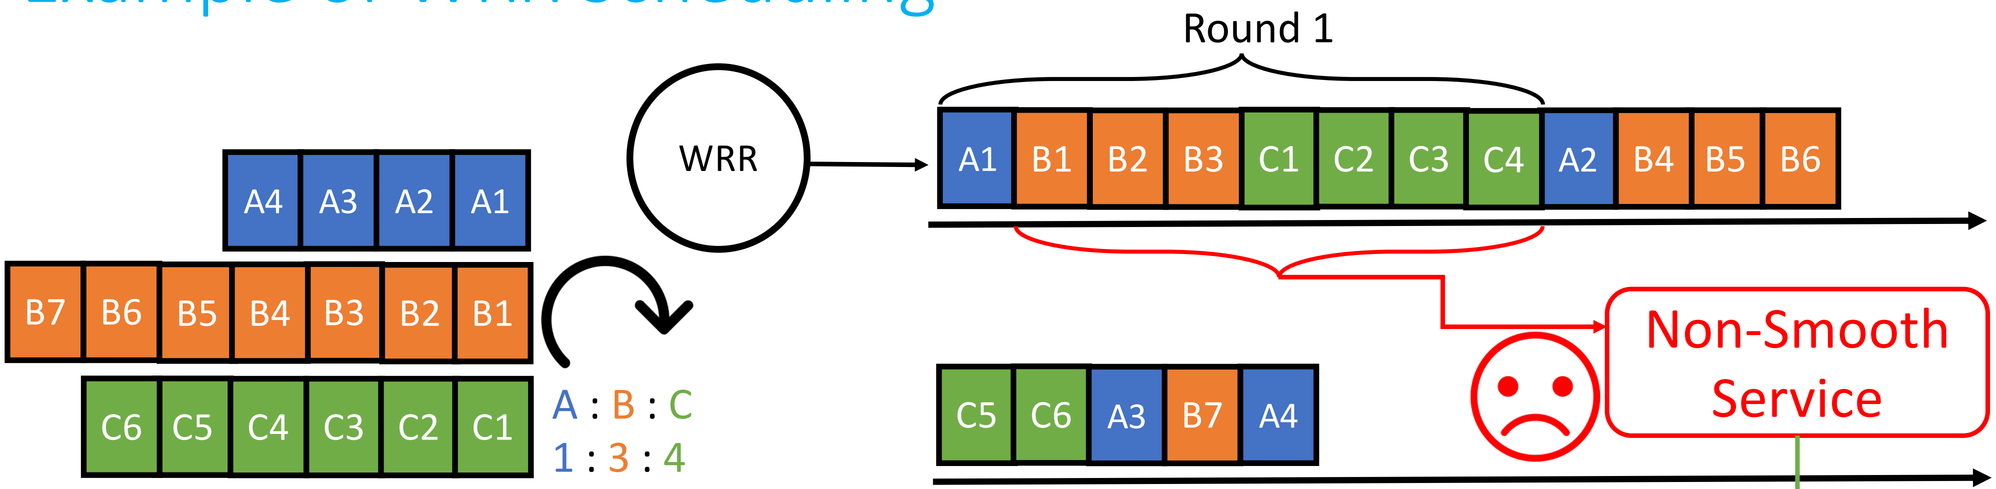

Interleaved Weighted Round-Robin (IWRR)

A Variant of WRR

↓

A queue with weight w is visited w times in each round.

↓

It can send up to one packet in each visit.

M. Katevenis, S. Sidiropoulos and C. Courcoubetis,
"Weighted round-robin cell multiplexing in a general-purpose ATM switch chip,"
in *IEEE Journal on Selected Areas in Communications*

IWRR Offers a Smoother Service.

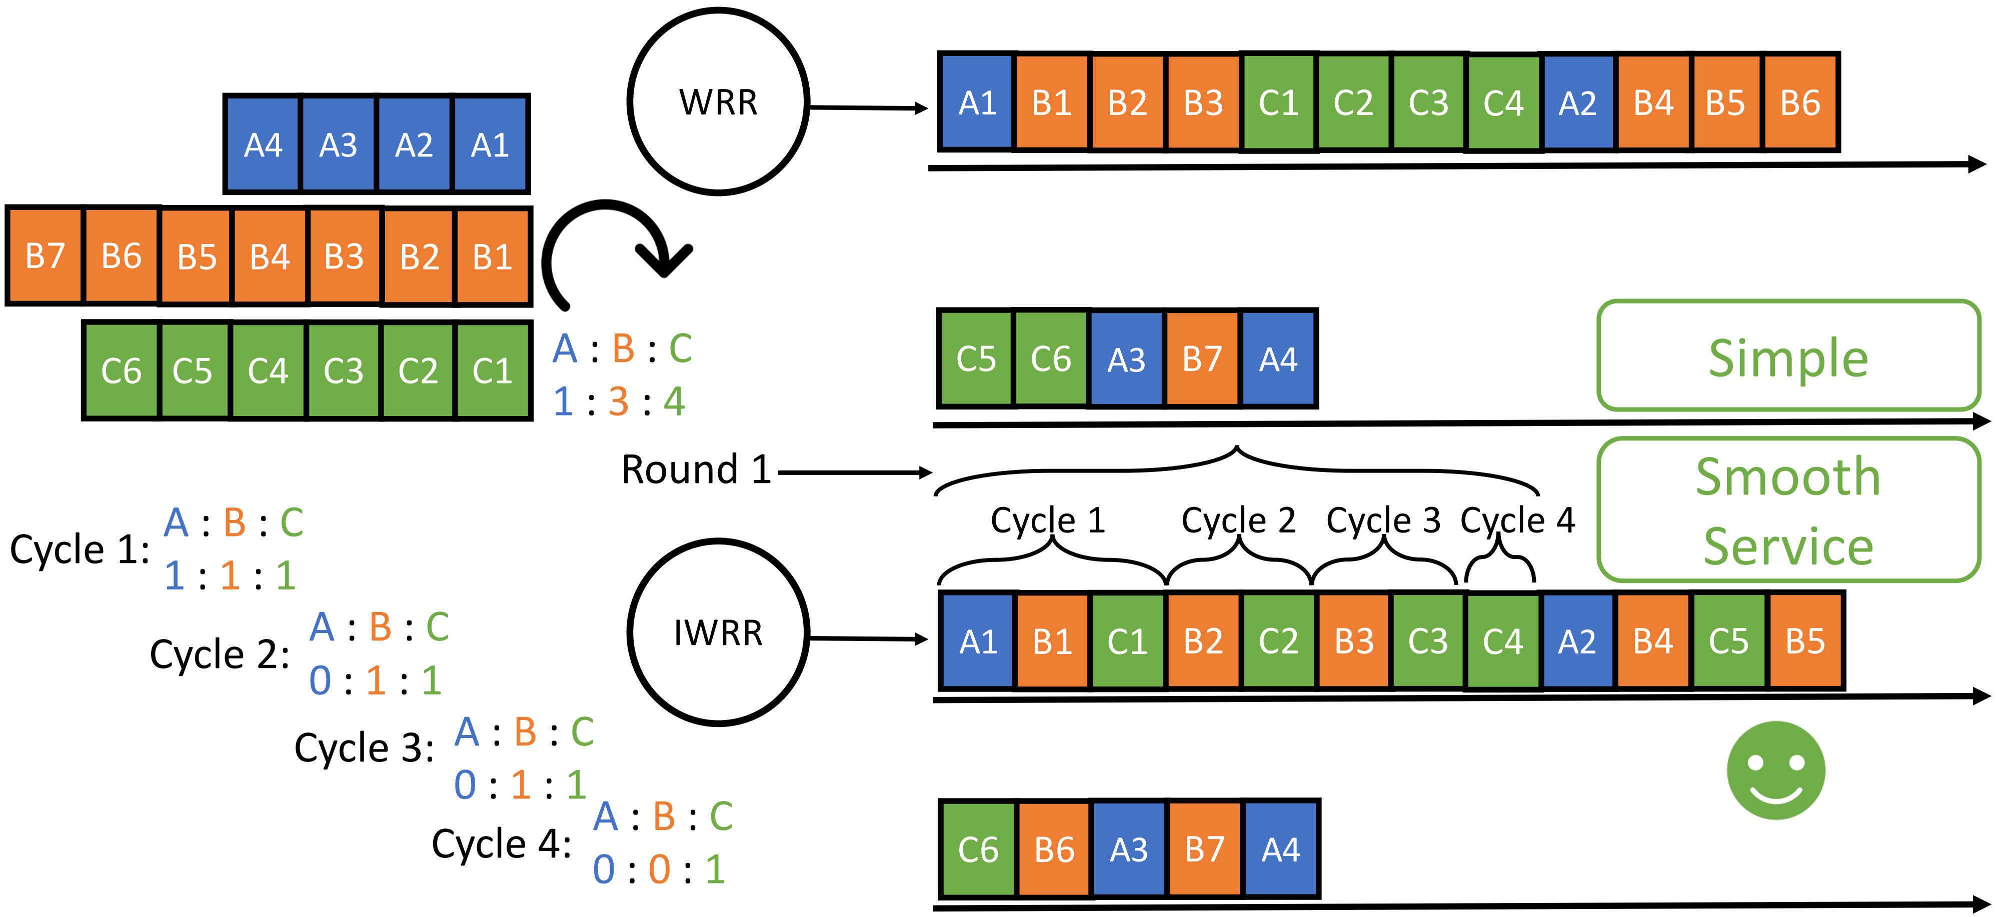

Does IWRR Reduce Worst-case Delays?

We expect that IWRR would reduce the worst-case delays.

↓ ?

How can we capture that?

↓

Using the *Network Calculus* framework.

↓

Finding a service curve

Table of Contents

State-of-the-art: WRR Strict Service Curve

Theorem[Bouillard, Boyer, Le Corronc 2018, Section 8.2.4]**:**
The strict service curve guaranteed to flow i is

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

$$q_i = w_i l_i^{\min},\ Q_i = \sum_{j, j \neq i} w_j l_j^{\max},\ L_{\text{tot}} = q_i + Q_i$$

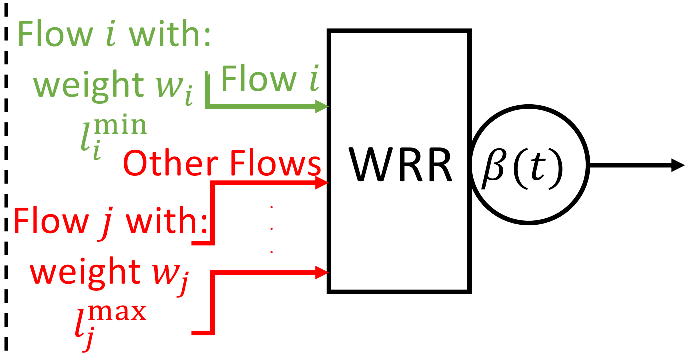

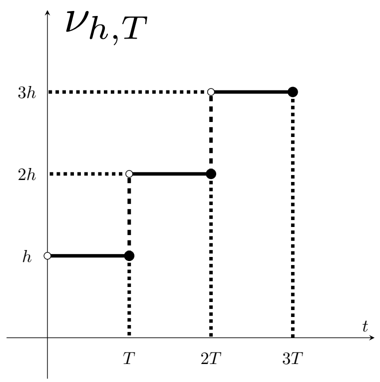

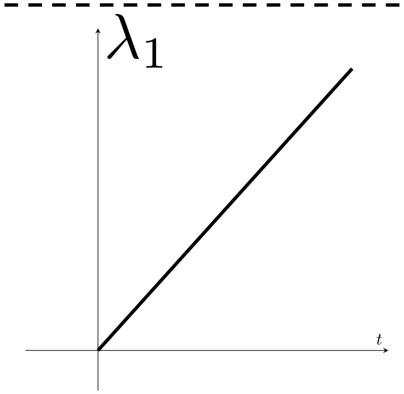

Min-plus convolution:
$(f \otimes g)(t) = \inf_{s \leq t} \{ f(s) + g(t-s) \}$

illustration?

Example WRR Strict Service Curve

$$\text{weight } w_i = 7$$
$$l_i^{\min} = 9 * 512 \text{ bit}$$
$$\text{weights } w_j = \{4, 6, 10\}$$
$$l_j^{\max} = \{17, 11, 16\} * 512 \text{ bit}$$
$$\beta(t) = ct \text{ with } c = 10Mbps$$

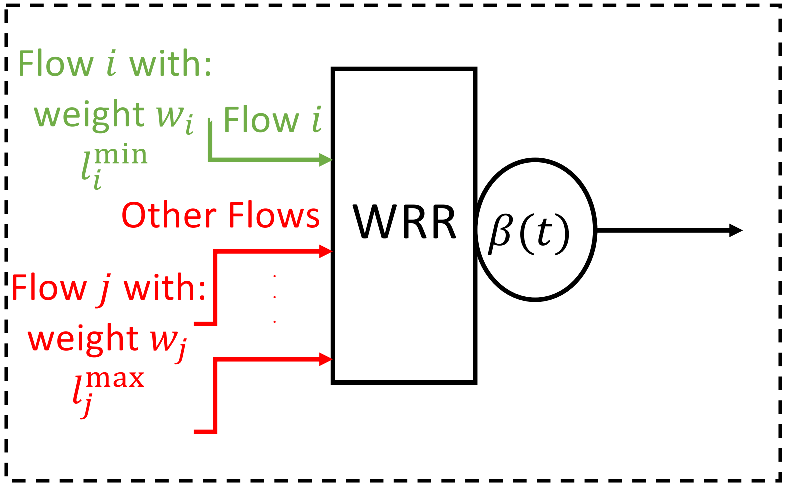

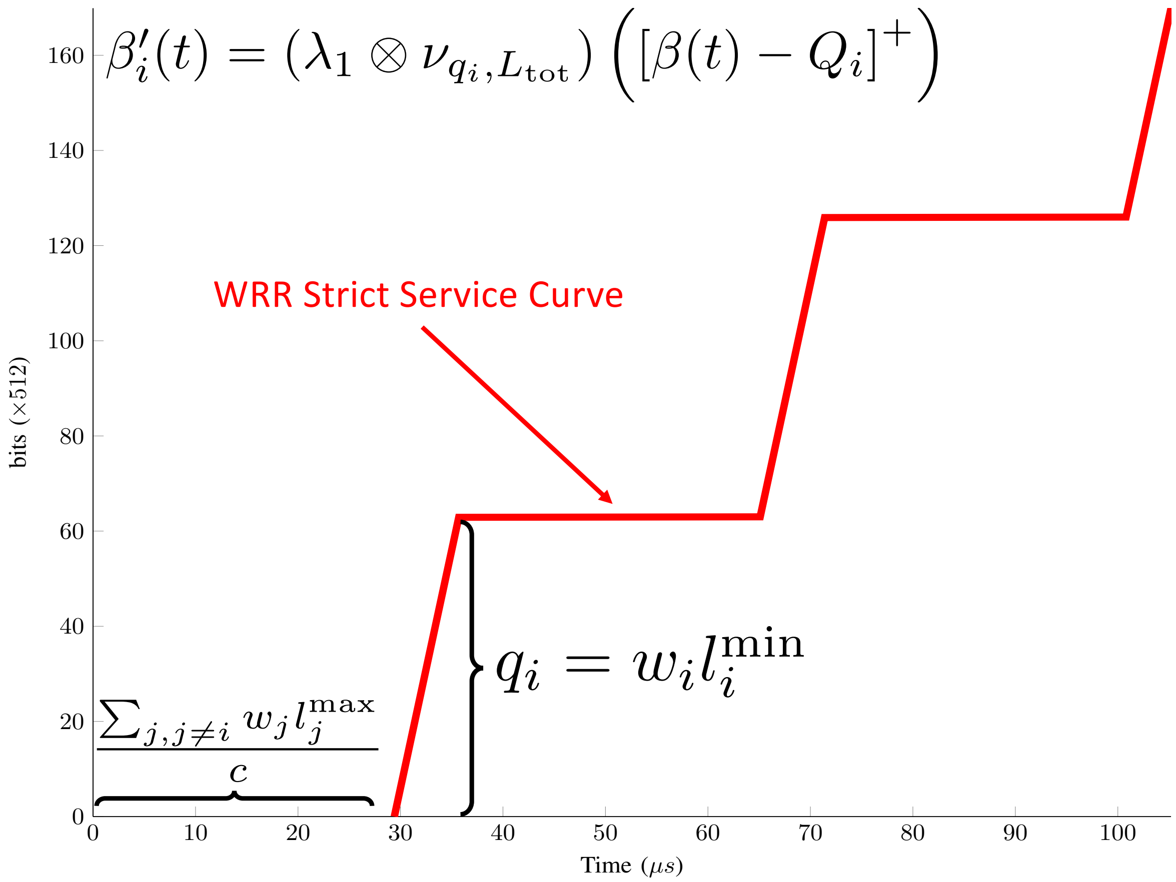

New Result: IWRR Strict Service Curve

Theorem 1: The strict service curve guaranteed to flow i is

$$\beta_i(t) = (\lambda_1 \otimes U_i)(\beta(t))$$

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

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

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

$$\phi_{i,j}(x) \stackrel{def}{=} \left\lfloor \frac{x}{w_i} \right\rfloor w_j + [w_j - w_i]^+ + \min(x \mod w_i + 1, w_j)$$

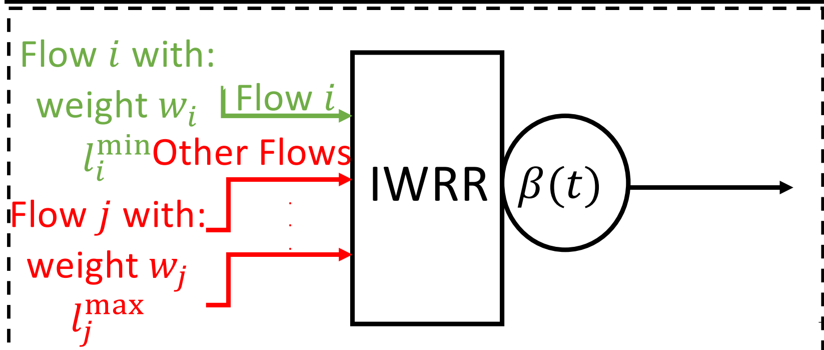

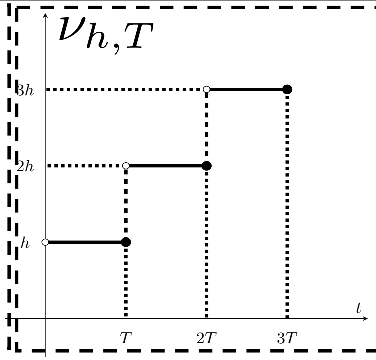

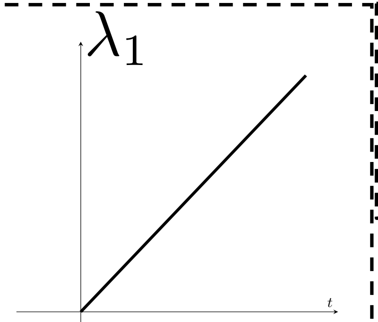

Min-plus convolution:
$(f \otimes g)(t) = \inf_{s \leq t}\{f(s) + g(t-s)\}$

illustration?

Example IWRR Strict Service Curve

$$\text{weight } w_i = 7$$
$$l_i^{\min} = 9 * 512 \text{ bit}$$
$$\text{weights } w_j = \{4, 6, 10\}$$
$$l_j^{\max} = \{17, 11, 16\} * 512 \text{ bit}$$
$$\beta(t) = ct \text{ with } c = 10Mbps$$

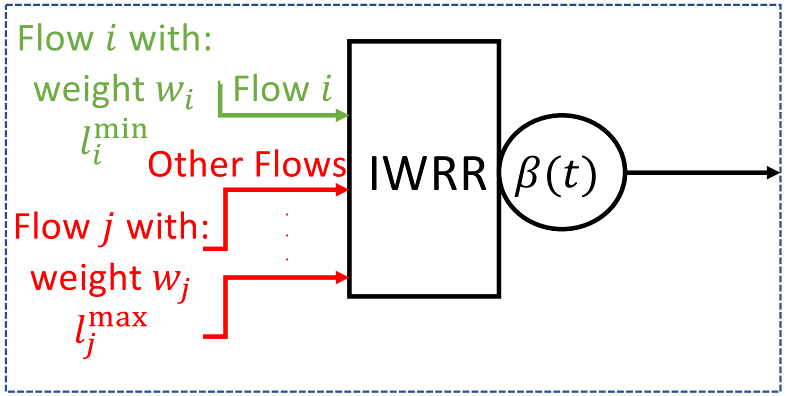

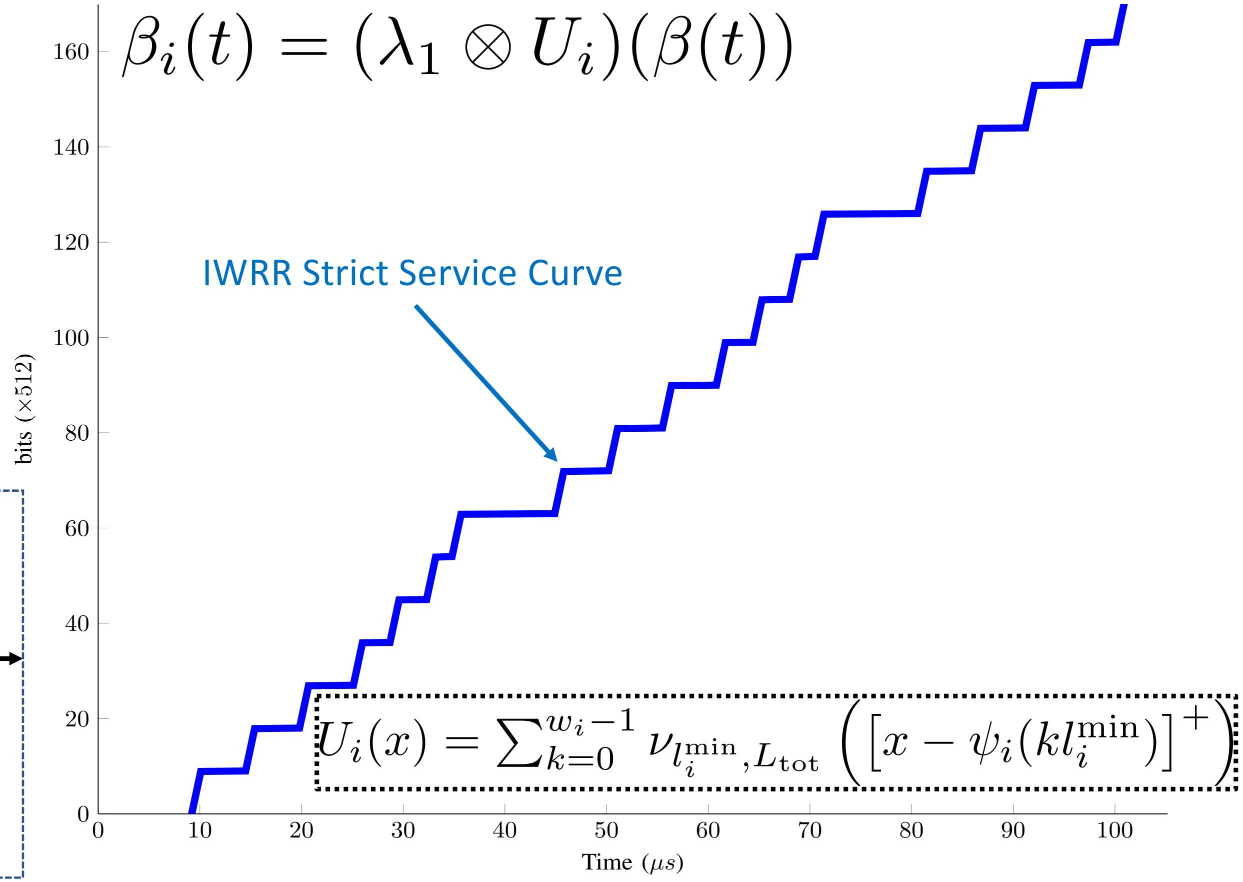

Sketch of Proof of IWRR Strict Service Curve

- $(\boldsymbol{s}, \boldsymbol{t}]$: flow of interest i is backlogged.
- $\boldsymbol{p}$: number of complete services of flow i.
- number of services for other flow $j \le \boldsymbol{\phi_{i,j}}\ (\boldsymbol{p})$

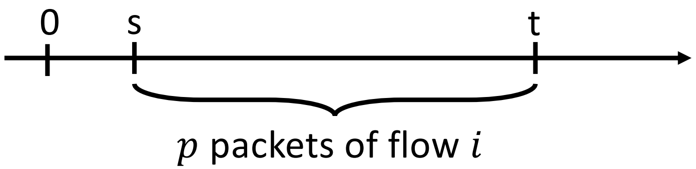

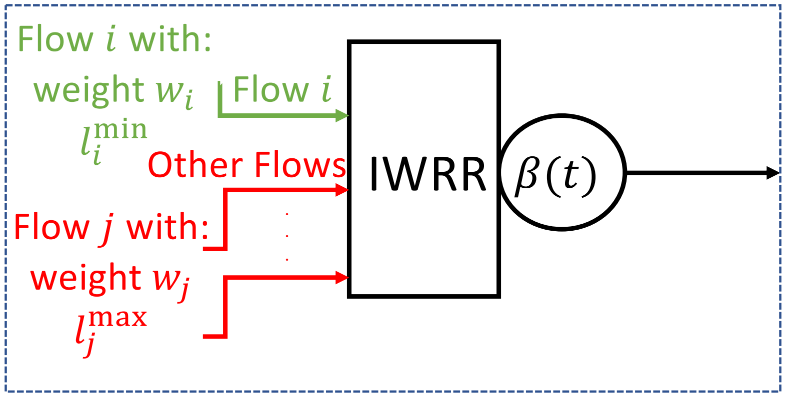

$$\beta(t-s) \le (R_i^*(t) - R_i^*(s)) + \sum_{j,j\neq i} R_j^*(t) - R_j^*(s)$$ (β is the strict service curve of server)

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

$$\le \underbrace{(R_i^*(t) - R_i^*(s)) + \sum_{j,j\neq i} \phi_{i,j}\left(\left\lfloor \frac{R_i^*(t) - R_i^*(s)}{l_i^{\min}} \right\rfloor\right) l_j^{\max}}_{\psi_i(R_i^*(t)-R_i^*(s))}$$

Lower Pseudo Inverse Technique

Definition[Jörg Liebeherr 2017, Section 10.1] **:** $f^{\downarrow}(y) = \inf\{x|f(x) \geq y\} = \sup\{x|f(x) < y\}$

Example:

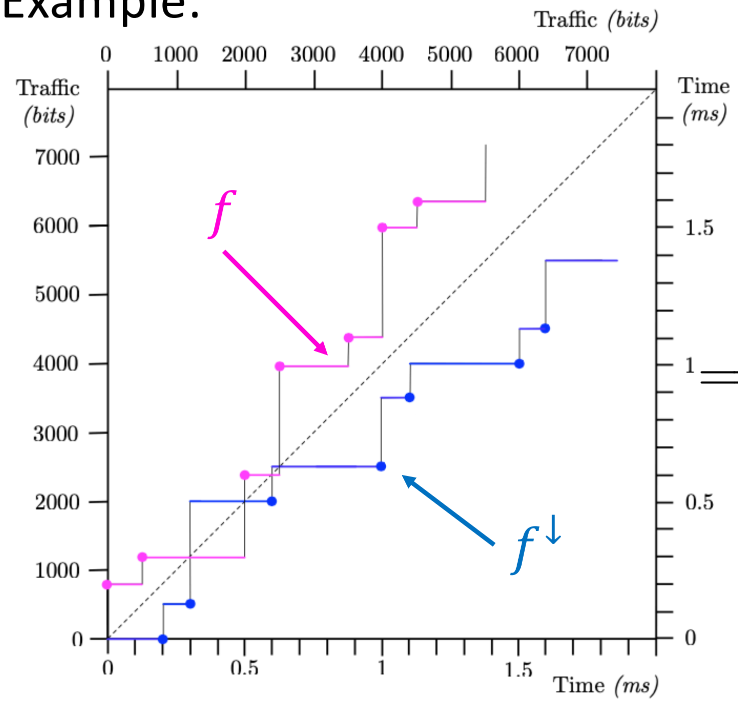

[Jörg Liebeherr 2017, Section 10.1]

- Sketch of Proof of IWRR Strict Service Curve (cont.)

Property: $y \leq f(x) \Longrightarrow x \geq f^{\downarrow}(y)$

$\beta(t-s) \leq \psi_i(R_i^*(t) - R_i^*(s))$(Last step of last slide)

$$\Longrightarrow R_i^*(t) - R_i^*(s) \geq \psi_i^{\downarrow}(\beta(t-s)) = \underbrace{(\lambda_1 \otimes U_i)(\beta(t-s))}_{\text{IWRR strict service curve for flow } i}$$

IWRR strict service curve for flow i

$\square$

Re-Cap 1

1. Strict Service Curve for IWRR.
   - We find a novel strict service curve for IWRR.
     - Using lower-pseudo inverse technique.

2. Comparison to WRR.

   Strict Service Curves (IWRR Vs. WRR)?

3. Tightness.

IWRR Strict Service Curve Always Improves Compared to WRR.

Theorem 2: The IWRR strict service curve is larger than or equal to the WRR strict service curve for each flow i

$$\underbrace{\beta'_i(t) = (\lambda_1 \otimes \nu_{q_i,L_{\text{tot}}})\left([\beta(t) - Q_i]^+\right)}_{\text{WRR Strict Service curve}} \le \underbrace{\beta_i(t) = (\lambda_1 \otimes U_i)(\beta(t))}_{\text{IWRR Strict Service curve}}$$

⟹Delay bounds for IWRR are less than or equal to delay bounds for WRR.

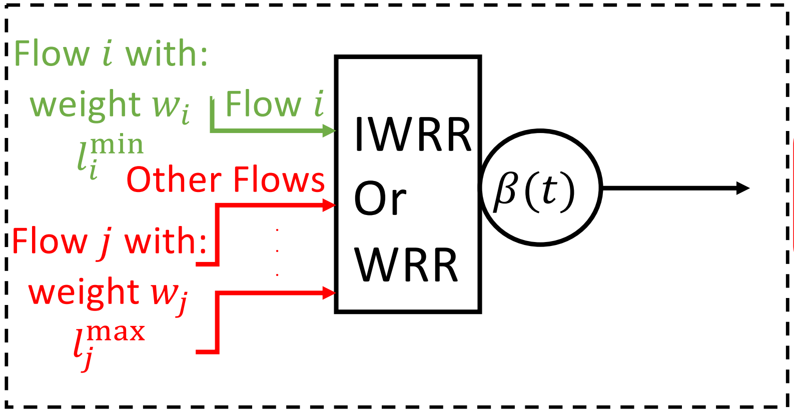

illustration?

Strict Service Curve of IWRR Is always Larger Than or Equal to WRR

$$\text{weight } w_i = 7$$

$$l_i^{\min} = 9 * 512 \text{ bit}$$

$$\text{weights } w_j = \{4, 6, 10\}$$

$$l_j^{\max} = \{17, 11, 16\} * 512 \text{ bit}$$

$$\beta(t) = ct \text{ with } c = 10Mbps$$

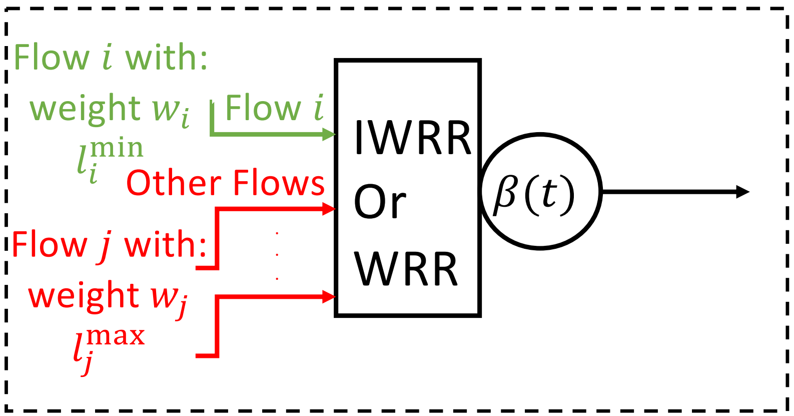

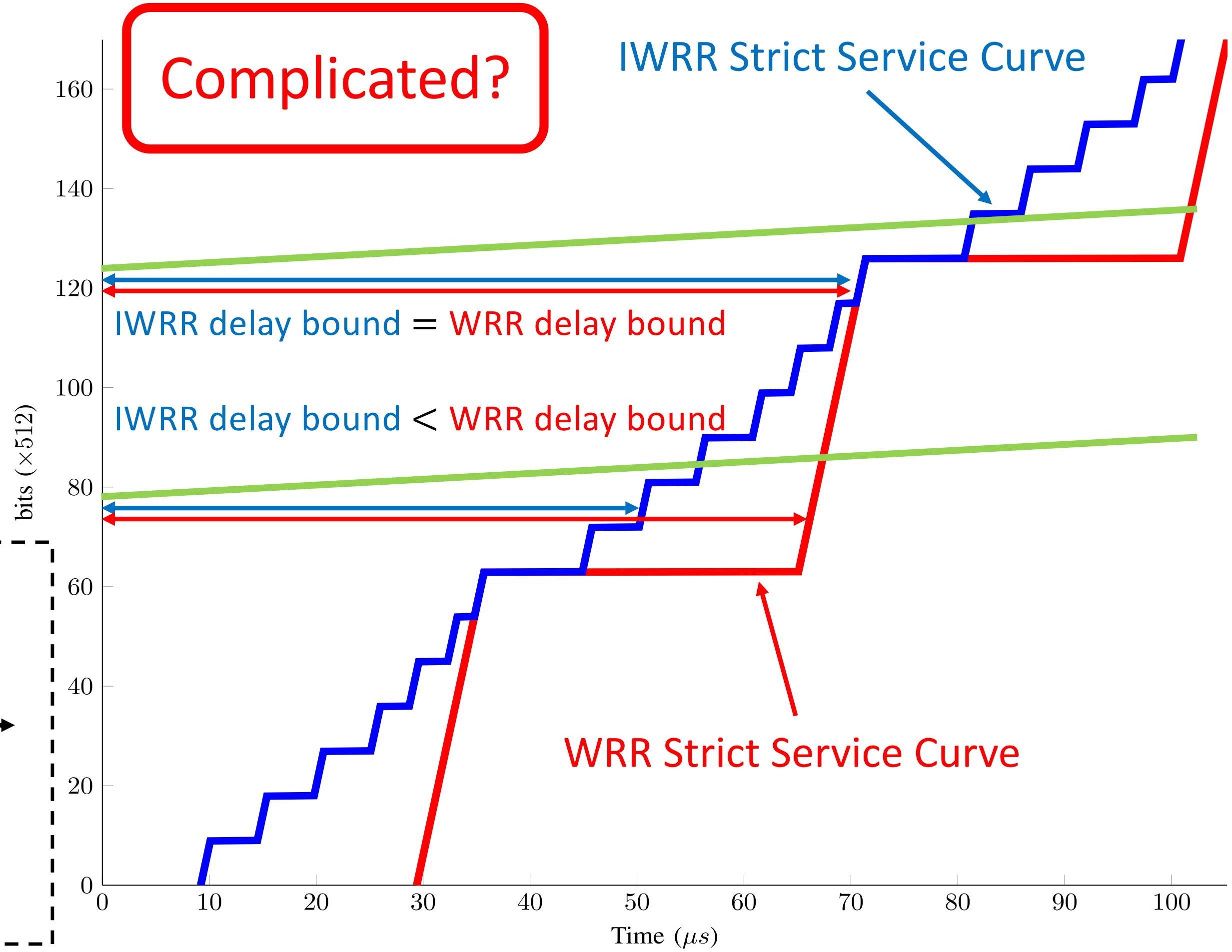

All Non-dominated Rate-latency Strict Service Curves for IWRR.

Theorem 3: All non-dominated rate-latency service curves guaranteed to flow i, among them $\beta_{r_0^*,T_0^*}$ is the one with the **minimum latency** and $\beta_{r_{k^*}^*,T_{k^*}^*}$ is the one with the **maximum rate**:

$$r_0^* = \frac{l_i^{\min}}{\psi_i(l_i^{\min}) - \psi_i(0)} \text{ and } T_0^* = \psi_i(0)$$

$$k^* = \min\{0 \le k < w_i \mid \frac{l_i^{\min}}{\psi_i((k+1)l_i^{\min}) - \psi_i(kl_i^{\min})} \ge \frac{q_i}{L_{\text{tot}}}\}$$

$$r_{k^*}^* = \frac{q_i}{L_{\text{tot}}} \text{ and } T_{k^*}^* = \psi_i(k^* l_i^{\min}) - \frac{k^* l_i^{\min}}{r_{k^*}^*}$$

illustration?

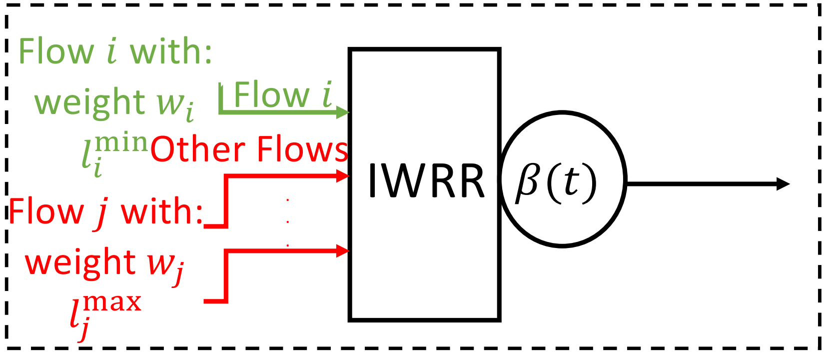

If the strict service curve of the aggregate is a rate-latency ($\beta(t) = \beta_{c,T}$) , then $\beta_{r_0^*,T_0^*}(\beta(t))$ and $\beta_{r_{k^*}^*,T_{k^*}^*}(\beta(t))$ are also rate-latency.

Example of Non-dominated Rate-latency IWRR Strict Service Curve

Possible case 1: One optimal non-dominated rate-latency:

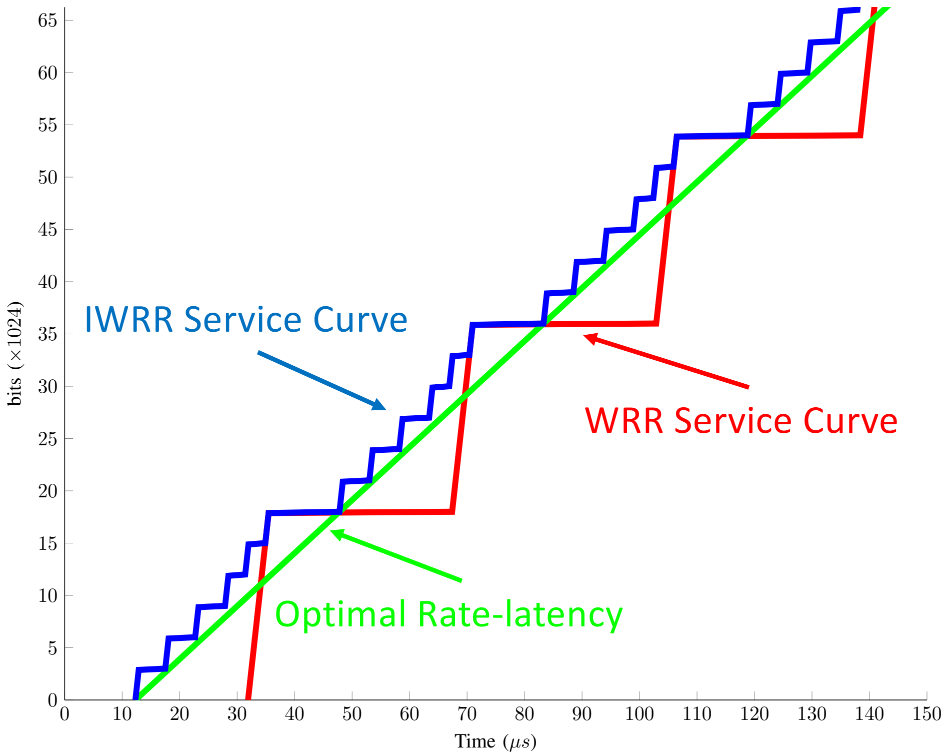

Possible case 2: Set of non-dominated rate-latency service curves:

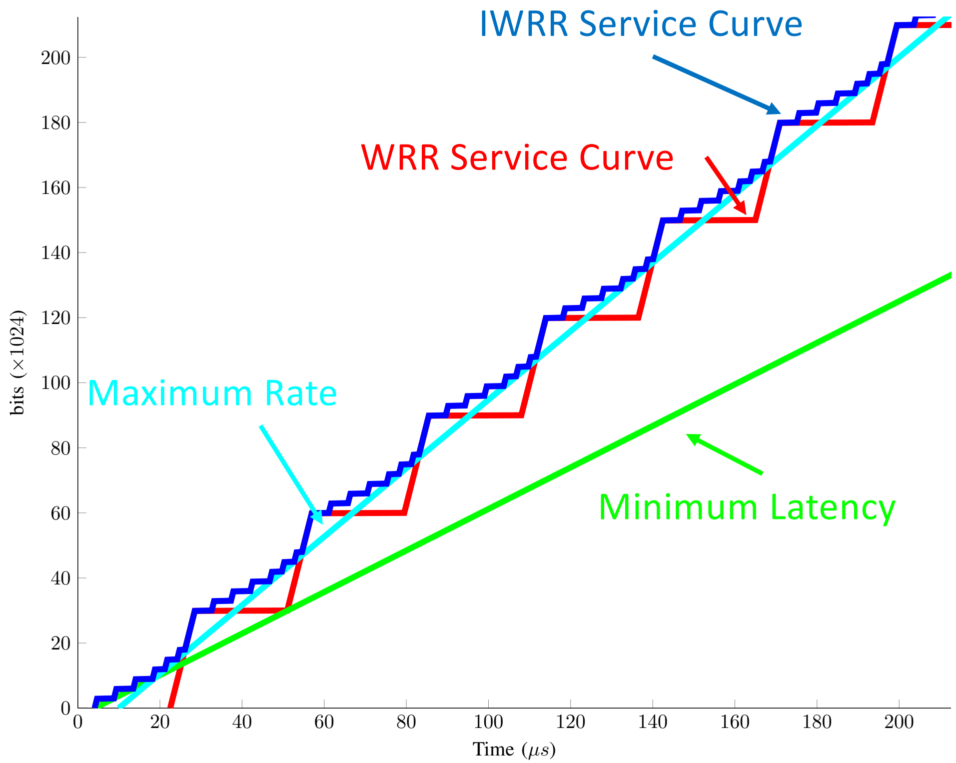

Re-Cap 2.1

1. Strict Service Curve for IWRR.
   - We find a novel strict service curve for IWRR.
     - Using lower-pseudo invers technique.
   - We find all non-dominated rate-latency service curves.

2. Comparison to WRR.
   - IWRR Strict Service Curve always improves compared to WRR.

3. Tightness.

Better Strict Service Curve?

Our Strict Service Curve Is the Best Possible.

Theorem 4: β_i is the **largest** strict service curve that can be given for flow i.

Specifically, we show for all flows i, there exists a trajectory scenario such that:

$$\exists s \geq 0,\ (s, s+\tau] \text{ is backlogged for flow } i$$
$$\text{and } R_i^*(s+\tau) - R_i^*(s) = \beta_i(\tau)$$

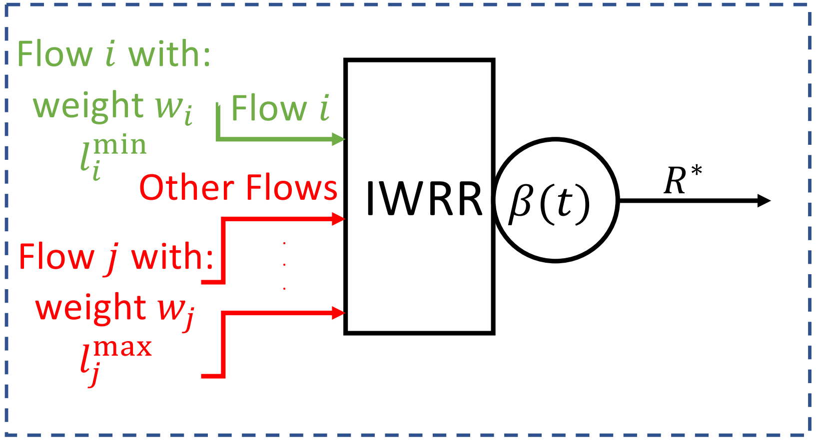

The same is also valid for the WRR strict service curve. (Theorem 5)

Re-Cap 2.2

We have found the largest possible strict service curve.

↓ ?

Automatically implies best delay bounds?

↓

No!

Service curve is only an abstraction of the service. The true thing is IWRR.

There might be a better service curve that is not strict.

Tight delay bounds for constant packet size.

Obtained Delay Bound Is **Tight,** for Constant Packet Size.

Theorem 6:

Flow i only generates packets of size l ($l_i^{\min} = l_i^{\max}$). Then, for every integer (multiple of l) arrival curve α_i^l, the network calculus delay bound is **tight** (i.e., worst case).

The same is also valid for the WRR strict service curve. (Theorem 7)

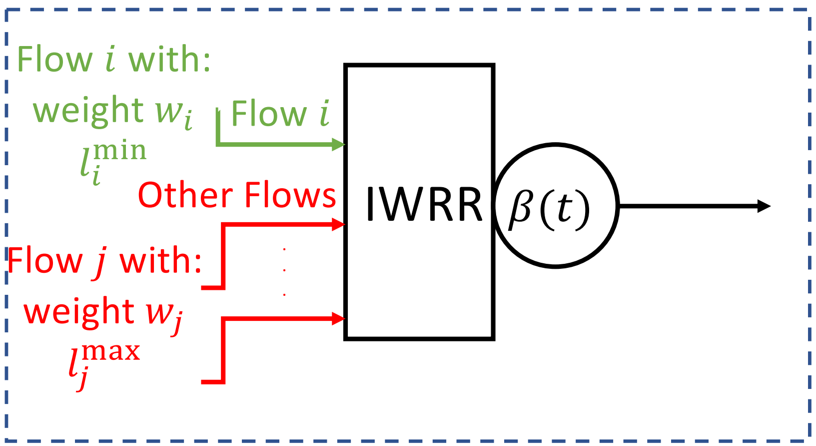

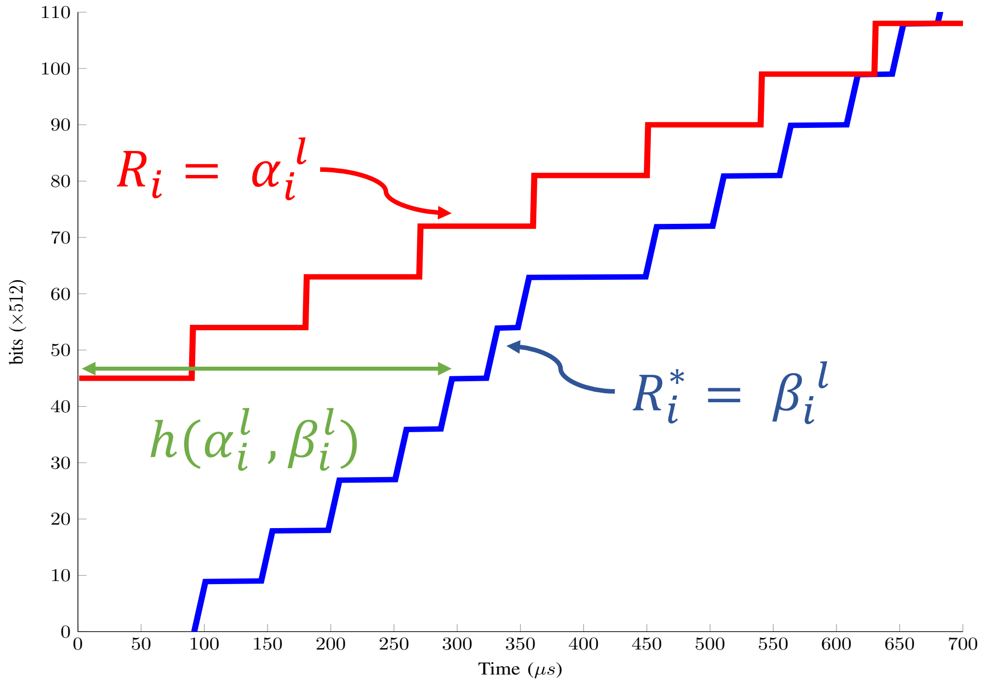

Re-Cap 3

1. Strict Service Curve for IWRR.
   - We find a novel strict service curve for IWRR.
   - We find all non-dominated rate-latency service curves.

2. Comparison to WRR. Numerical Example?
   - IWRR Strict Service Curve always improves compared to WRR.

3. Tightness.
   - Our Strict Service Curve is the best possible one.
   - For a flow with constant packet size, we show that:
     - Obtained delay bound is **tight**.

Improvement of Delay Bounds with IWRR.

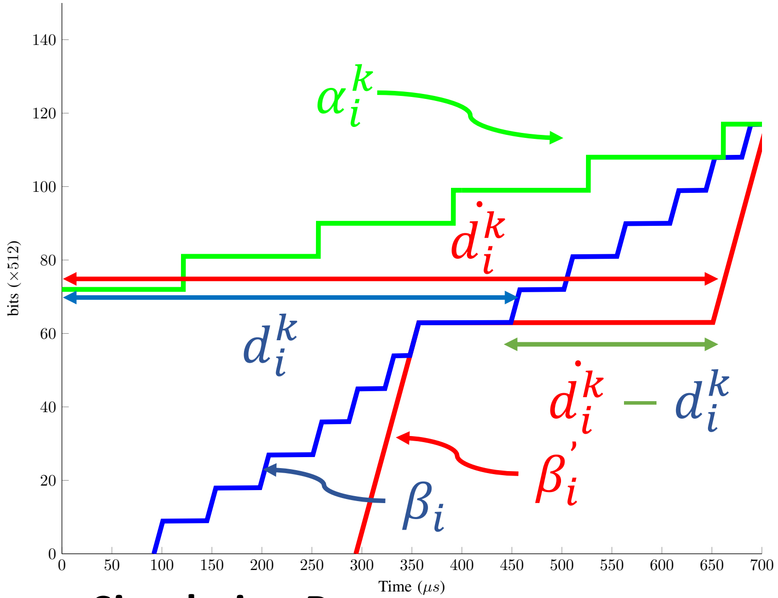

Simulation Parameters:
8 input flows.
Weights = {22,27,28,30,34,41,45}
$l^{\max} = l^{\min} = 7119$ bit.
Constant bit rate server, 10 Mbps.

Token-bucket Arrival Curves:
Initial bursts: uniformly [1,20] packets.
A rate of 0.5 Mbps.

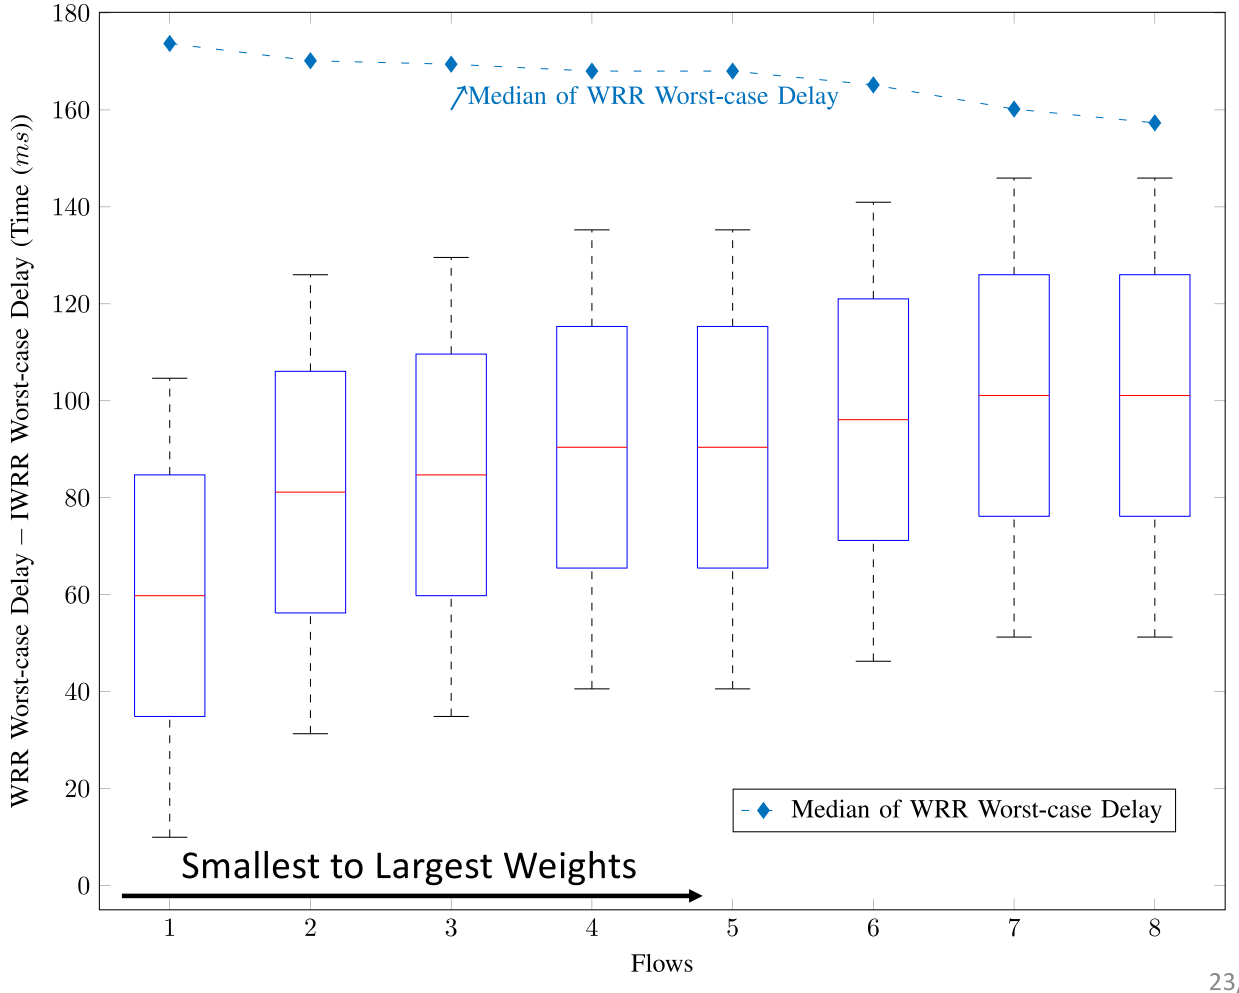

Better Improvement for Flows with Larger Weights.

$$\text{Improvement Ratio} = \frac{\text{Improvement}}{\text{Median of WRR worst\_case delay}}$$

Simulation Parameters:
8 input flows.
Weights $\in [10\,,50]$.
$l^{\max} = l^{\min} \in [64\,,1522]$ byte.
Uniformly random.
Constant bit rate server, 10 Mbps.

Token-bucket Arrival Curves:
Initial bursts: uniformly [1,20] packets.
A rate of 0.5 Mbps.

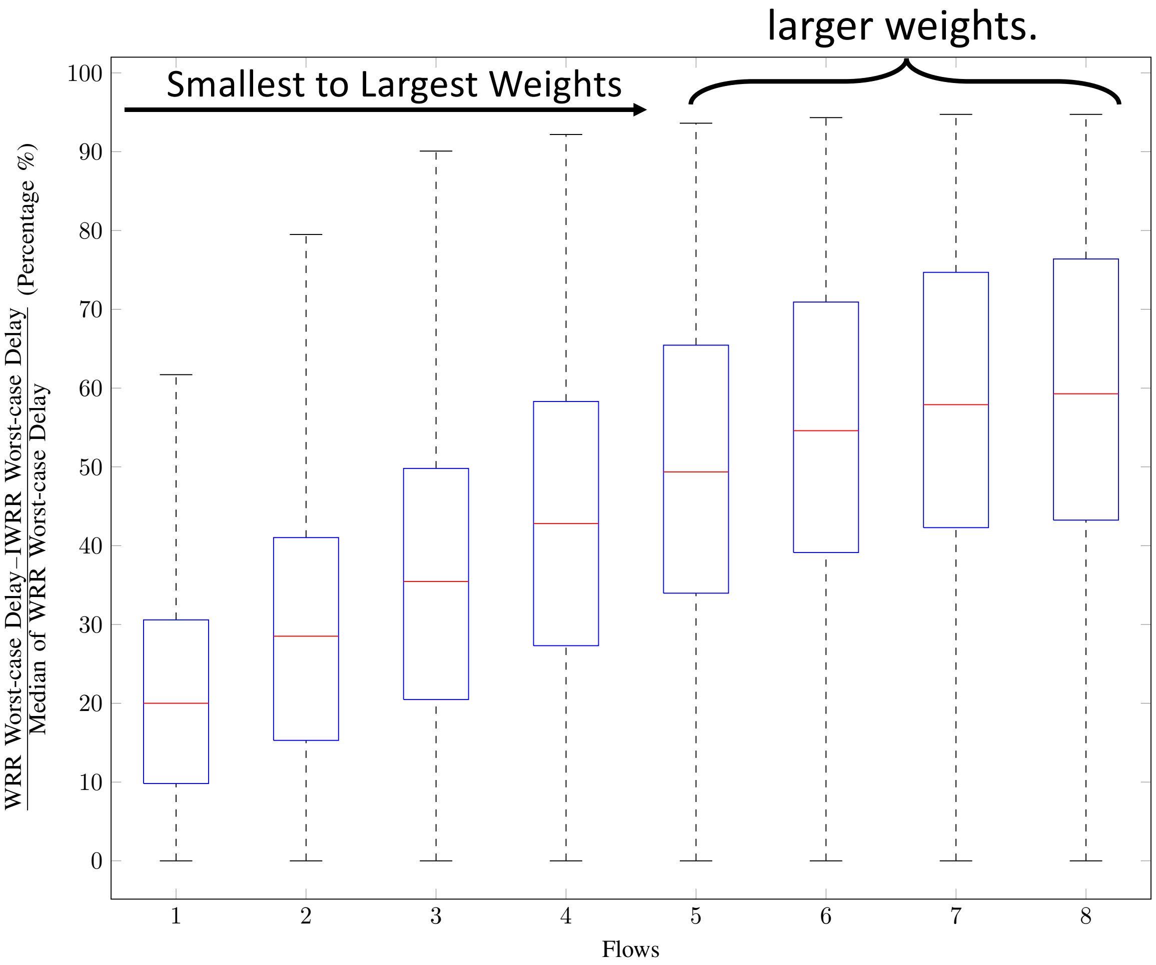

Conclusion

- We find a novel strict service curve for IWRR.
  - Using lower-pseudo inverse technique.
- We show that:
  - It is the best possible one.
  - It always improves compared to WRR.
- We also find all non-dominated rate-latency service curves.
- We show that obtained delay bound is **tight,** for a flow with constant packet size.

hossein.tabatabaee@epfl.ch

References

- M. Katevenis, S. Sidiropoulos and C. Courcoubetis, "Weighted round-robin cell multiplexing in a general-purpose ATM switch chip," in *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 8, pp. 1265-1279, Oct. 1991, doi: 10.1109/49.105173.
- A. Bouillard, M. Boyer, and E. Le Corronc, Deterministic Network Calculus: From Theory to Practical Implementation. Wiley-ISTE.
- Jörg Liebeherr (2017), "Duality of the Max-Plus and Min-Plus Network Calculus", Foundations and Trends® in Networking: Vol. 11: No. 3-4, pp 139-282. http://dx.doi.org/10.1561/1300000059.
- S. M. Tabatabaee, J.-Y. L. Boudec, and M. Boyer, "*Interleaved weighted round-robin: A network calculus analysis*," https://arxiv.org/pdf/2003.08372.pdf