RUHR-UNIVERSITÄT BOCHUM

SEARCHING FOR UPPER DELAY BOUNDS IN FIFO MULTIPLEXING FEEDFORWARD NETWORKS

Alexander Scheffler (RUB), Jens Schmitt (TUK) and Steffen Bondorf (RUB)

Overview

- **Deterministic Network Calculus (DNC) Motivation and Basics**
- **Objective and Approaches**
- **Related Work**
- **LB-FF and DS-FF**
- **Evaluation**

DNC Motivation and Basics

- **Theory of deterministic queueing systems [Cruz91]**
  - Worst-case bounds such as delay and backlog

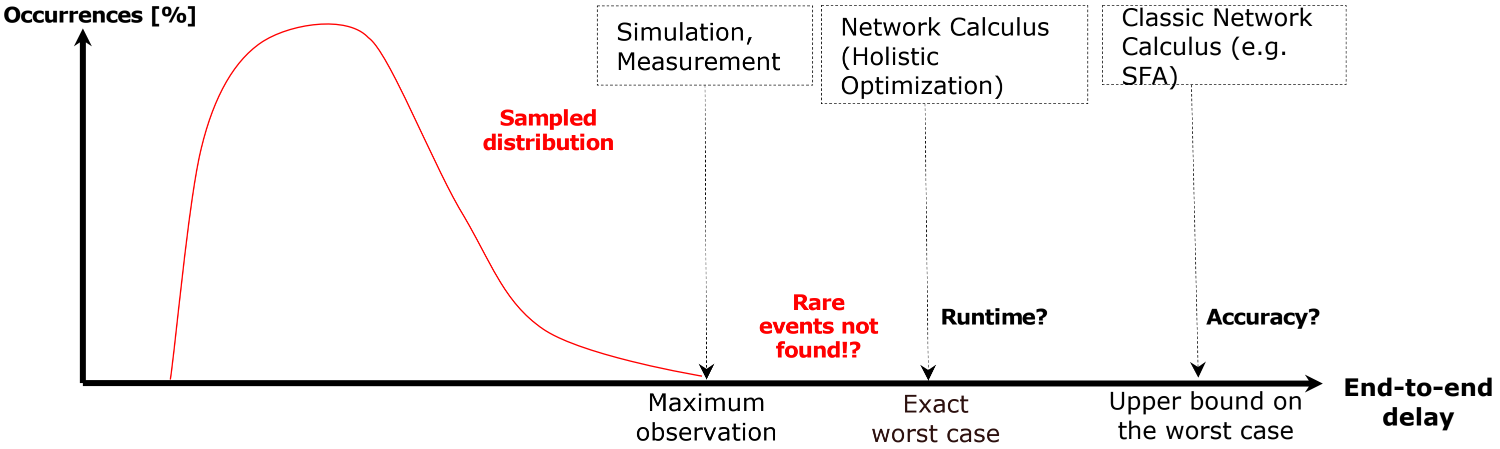

DNC Motivation and Basics (2)

- **Can be used for certifying performance guarantees of cyber-physical systems, e.g., airplanes [Boyer08]**
- **Can aid in ranking different network topologies and configurations**

DNC Motivation and Basics (3) [LeBoudec01]

- **Arrival curve** $\forall 0 \leq s \leq t : A(t) - A(t-s) \leq \alpha(s)$
- **Service curve** $\forall t \geq 0 : A'(t) \geq \inf_{0 \leq s \leq t} \{A(t-s) + \beta(s)\} := A \otimes \beta(t)$

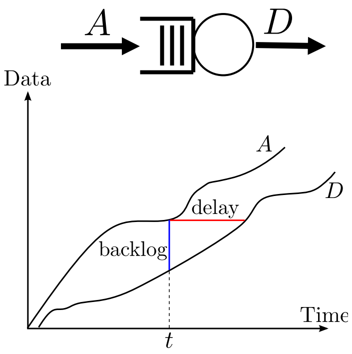

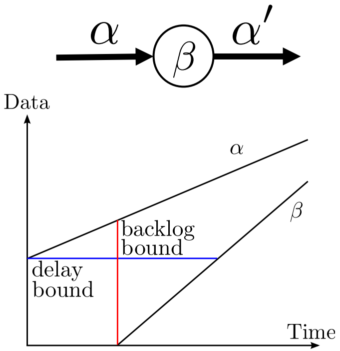

DNC Motivation and Basics (4) [LeBoudec01]

- **Concatenation of servers** $\beta_1 \otimes \beta_2 = \beta_{1,2}$
- **Output bound** $\alpha'(t) = \alpha \oslash \beta(t) := \sup_{u \geq 0}\{\alpha(t+u) - \beta(u)\}$
- **Delay bound** $hdev(\alpha, \beta) = \inf\{d \geq 0 : (\alpha \oslash \beta)(-d) \leq 0\}$

DNC Motivation and Basics (5) [LeBoudec01]

- **FIFO left-over service curve**

$$\beta_{f_1}^{\text{l.o.}}(t,\theta) = [\beta(t) - \alpha_2(t-\theta)]^{\uparrow} \cdot 1_{\{t>\theta\}} \forall \theta \geq 0$$

$$[g(x)]^{\uparrow} = \sup_{0 \leq z \leq x} g(z)$$

$$1_{\{t>\theta\}} := \begin{cases} 1, & t > \theta \\ 0, & t \leq \theta \end{cases}$$

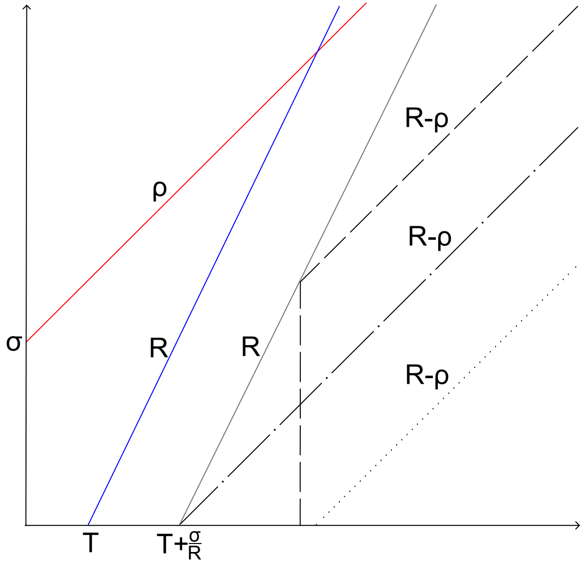

RUHR UNIVERSITÄT BOCHUM

Objective and Approaches

- **LUDB-FF [Scheffler21] optimizes these free parameters**
- **New Approach**
  - Do not use optimization
  - For a better tradeoff between accuracy and runtime, we provide a static setting of the free parameters (LB-FF) as well as one found by a directed search (DS-FF)
- **Find our code at**
  - https://github.com/NetCal/DNC

Related Work

- **SFA-FIFO**
  - Server-by-server analysis
  - Compute at each server the residual service curve, convolve them
  - Each occuring θ set statically, $T + \frac{\sigma}{R}$
  - Simple left-over curve but local view

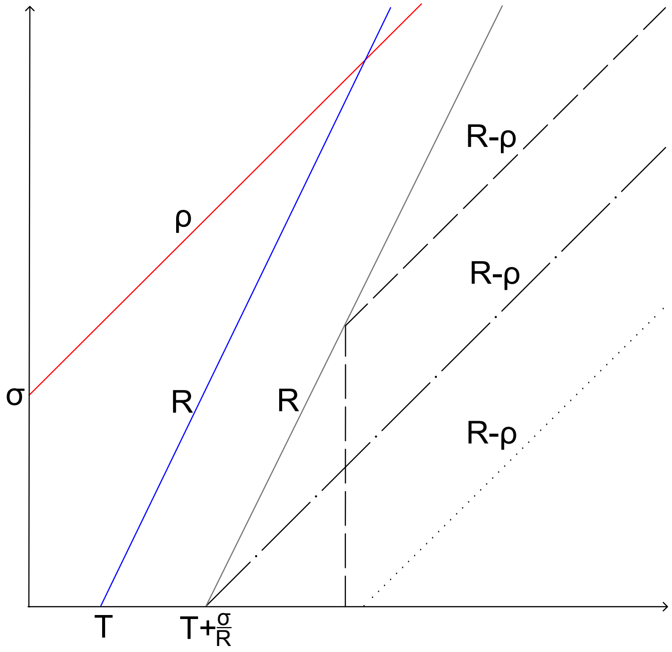

RUHR UNIVERSITÄT BOCHUM RUB

Related Work (2)

- **LUDB [Bisti08]**
  - Nested interference: A tandem has nested interference iff for every pair of flows either both flows do not have common servers or the path of one flow is completely included in the path of the other.

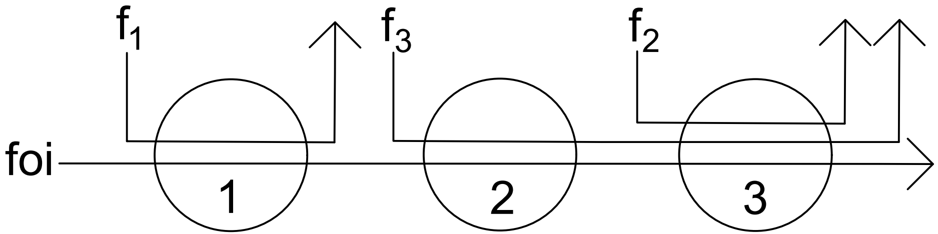

Related Work (3)

- **LUDB [Bisti08]**
  - "convolution before subtraction" for nested interference
  - Optimizes $(\theta_1, ..., \theta_{|F_x|})$ w.r.t. delay bound

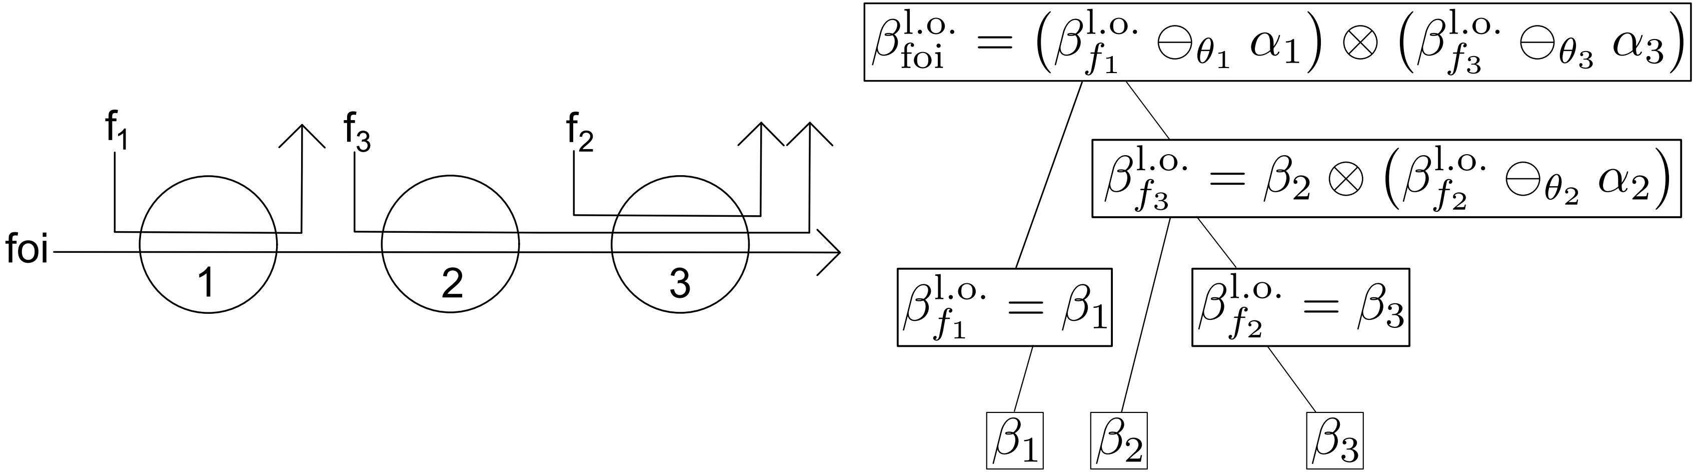

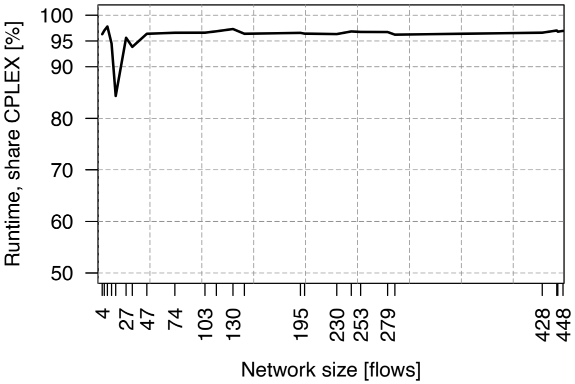

Related Work (4)

- **LUDB-FF [Scheffler21]**
  - Extension of LUDB [Bisti08] for feedforward networks
    - Step 1: Take the foi's path as first tandem to analyze.
    - Step 2: Check for non-nested interference patterns, cut into nested tandems if necessary (is always possible).

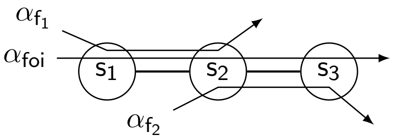

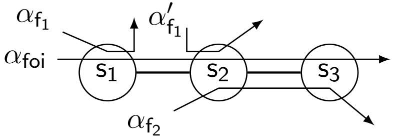

    - Step 3: Compute the bounds on crossflow arrivals to the nested tandems by starting Step 1 with them as F_{foi}
    - Step 4: For each nested tandem, compute the nesting tree that encodes the "convolution before subtraction" scheme according to the nesting of flows.
    - Step 5: Start the LUDB analysis for the nesting tree. The LUDB produces a Piecewise Linear Program (PLP).
    - Step 6: For each linear decomposition of the PLP, one LP will be formulated and solved (with CPLEX e.g.).

Related Work (5)

- **Holistic Network Optimization [Bouillard15, Bouillard22]**
  - Transforms entire feedforward network with its NC modelling constraints into a MILP formulation
  - Computes the WCD but exponential runtime, thus only feasible for small/medium-sized networks
  - The LP relaxation, FF-LPA, has a better runtime but is still exponential
  - Recent efforts try to improve the runtime of this approach by dropping some constraints and adding constraints from analyses such as SFA-FIFO

LB-FF and DS-FF

- **LB-FF and DS-FF share the ”-FF” suffix with LUDB-FF**
  - Embedded into the same feedforward analysis
- **Instead of optimizing $(\theta_1, ..., \theta_{|F_x|})$, we set the free FIFO variables in a static way (LB-FF) / employ a search on top of LB-FF (DS-FF)**
- **Def. (Lower Bound on θ)**
  - Given an arrival curve $\alpha_f := \gamma_{\rho,\sigma}$ and left-over service curve $\beta_f^{\text{l.o}}$; we define $\underline{\theta}_f(\beta_f^{\text{l.o.}}, \alpha_f) := \inf\{t \geq 0 : \beta_f^{\text{l.o.}}(t) \geq \sigma\}$
  - Note that $\underline{\theta}_f(\beta_f^{\text{l.o.}}, \alpha_f) = hdev(\alpha_f, \beta_f^{\text{l.o.}})$

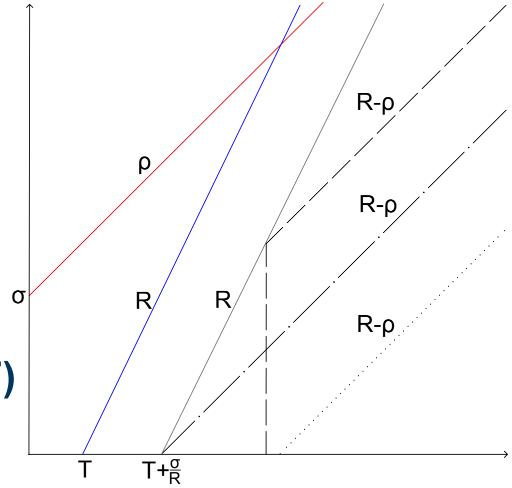

LB-FF and DS-FF (2)

- **LB-FF**
  - Procedure: Traverse nesting tree bottom-up, set each θ_i to the lower bound
  - $\theta_1 = T_1 + \frac{\sigma_1}{R_1}$ $\qquad$ $\theta_2 = T_3 + \frac{\sigma_2}{R_3}$ $\qquad$ $\theta_3 = T_2 + \theta_2 + \frac{\sigma_3}{\min\{R_2, R_3 - \rho_2\}}$

$$\beta^{\text{l.o.}}_{\text{foi}} = \beta_{\min\{R_1-\rho_1, R_2-\rho_3, R_3-\rho_2-\rho_3\}, T_1+T_2+T_3+\frac{\sigma_1}{R_1}+\frac{\sigma_2}{R_3}+\frac{\sigma_3}{\min\{R_2, R_3-\rho_2\}}}$$

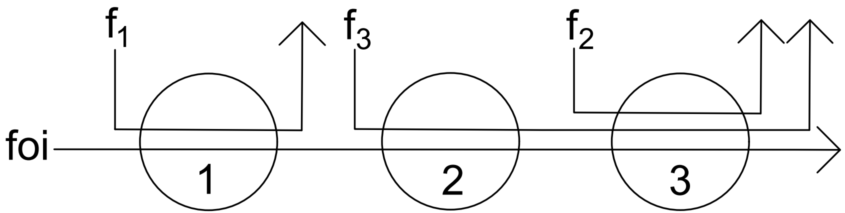

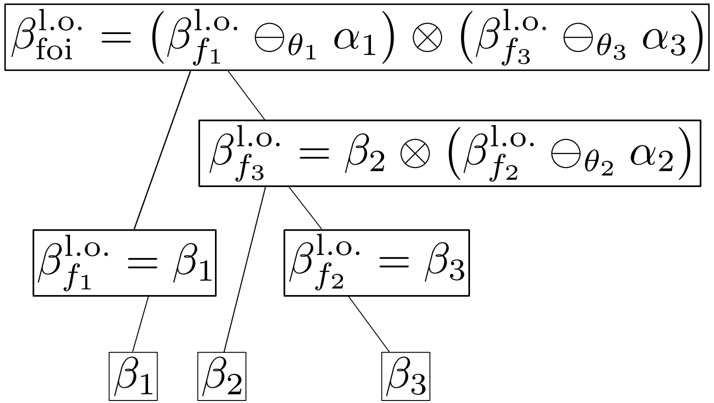

LB-FF and DS-FF (3)

- **DS-FF**
  - Uses a directed search (Hooke and Jeeves) to find a reasonable setting of $(\theta_1, ..., \theta_{|F_x|})$
  - Local Optima?

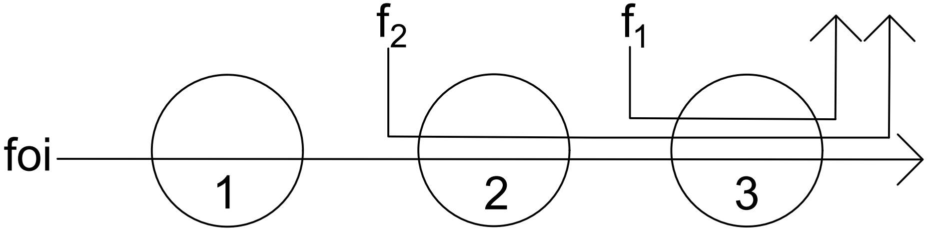

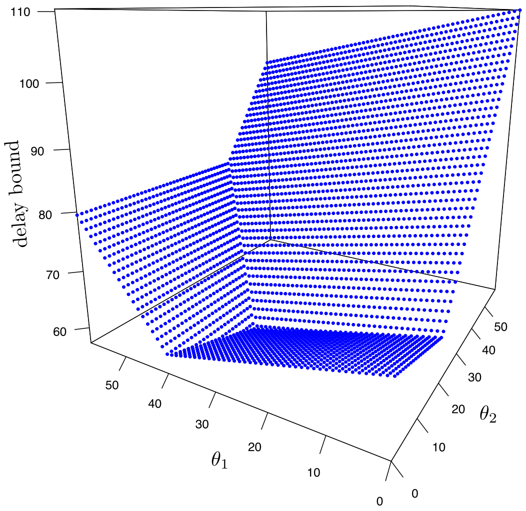

LB-FF and DS-FF (4)

- **DS-FF**
  - Start combination $\Theta^{\text{LB-FF}} := (\theta_1^{\text{LB-FF}}, \ldots, \theta_{|F_x|}^{\text{LB-FF}})$
  - Exploratory Phase
    - Search within the current environment for a better setting $(\theta_1^{\text{new}}, \ldots, \theta_{i-1}^{\text{new}}, \theta_i^{\text{new}} \pm sp_i, \theta_{i+1}^{\text{new}}, \ldots, \theta_{|F_x|}^{\text{new}})$
  - Pattern Move Phase
    - Compute direction of improvement $\Delta = \Theta^{\text{new}} - \Theta^{\text{old}}$
    - Walk along the line of improvement $\Theta^{\text{new}} + \Delta, \Theta^{\text{new}} + 2 \cdot \Delta, \Theta^{\text{new}} + 3 \cdot \Delta, \ldots$
    - Go back to the Exploratory Phase
  - Reduce stepsize $sp_i \leftarrow sp_i \cdot \xi$ with $0 < \xi < 1$
  - Terminate if $\min_{i=1,..,|F_x|} \{sp_i\} < \epsilon$ with $\epsilon > 0$

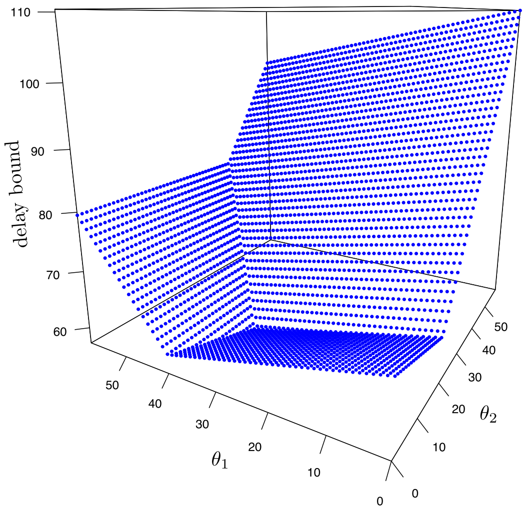

LB-FF and DS-FF (5)

- **DS-FF: Runtime improvements**
  - Selection of initial stepsize
    - (i) $f_i \in \text{Children}(\text{foi})$

      $$\overline{\theta}_i(\Theta^{\text{LB-FF}}) := d^{\text{LB-FF}} - \sum_{k \in S(\text{foi}) \setminus S(\text{Children}(\text{foi}))} T_k - \sum_{f_i \neq f_k \in \text{Children}(\text{foi})} \theta_k^{\text{LB-FF}}$$

    - (ii) $f_i \notin \text{Children}(\text{foi})$. Let $f_p \in F_x \text{s.t.} f_i \in \text{Children}(f_p)$. Then,

      $$\overline{\theta}_i(\Theta^{\text{LB-FF}}) := \overline{\theta}_p(\Theta^{\text{LB-FF}}) - \sum_{k \in S(f_p) \setminus S(\text{Children}(f_p))} T_k - \sum_{f_i \neq f_k \in \text{Children}(f_p)} \theta_k^{\text{LB-FF}}$$

    - Reasoning: Consider $\Theta^{\text{LB-FF}}$,then $(\theta_1^{\text{LB-FF}}, \ldots, \theta_{i-1}^{\text{LB-FF}}, \overline{\theta}_i(\Theta^{\text{LB-FF}}), \theta_{i+1}^{\text{LB-FF}}, \ldots, \theta_{|F_x|}^{\text{LB-FF}})$ results in a worse delay bound for any i
    - Hence, we set the stepsize $sp_i = \dfrac{\overline{\theta}_i(\Theta^{\text{LB-FF}}) - \theta_i^{\text{LB-FF}}}{c-1}$

LB-FF and DS-FF (6)

- **DS-FF: Runtime improvements**
  - Before computing the left-over service curve w.r.t. a specific combination (in any phase), check that each $\theta_i \geq 0$ and $\theta_i < d$ with d being the lowest delay bound so far found by the search

Evaluation

- **Setup**
  - 31 random feedforward networks following Erdös-Rényi model with a total of 4479 flows
  - Arrival curves set to token bucket $\gamma_{\rho,\sigma} = \gamma_{1,1}$
  - Service curve set to rate latency $\beta_{R,T}$ with $T = 0$ and R set to achieve a desired utilization of the server between 50% and 99%
- **Goal**
  - Find a suitable termination criterion ϵ for DS-FF giving a good tradeoff between effort and quality of bounds

Evaluation (2)

- **Comparison of Delay Bounds** $delay_{\text{analysisA,analysisB}} = \frac{delay_{\text{analysisA}} - delay_{\text{analysisB}}}{delay_{\text{analysisB}}}$

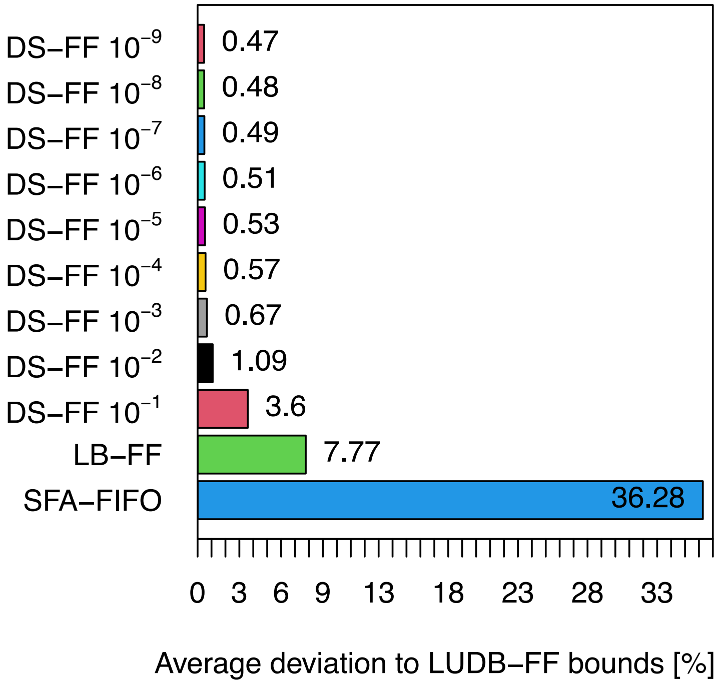

Evaluation (3)

- **Comparison of Delay Bounds** $delay_{\text{analysisA,analysisB}} = \frac{delay_{\text{analysisA}} - delay_{\text{analysisB}}}{delay_{\text{analysisB}}}$

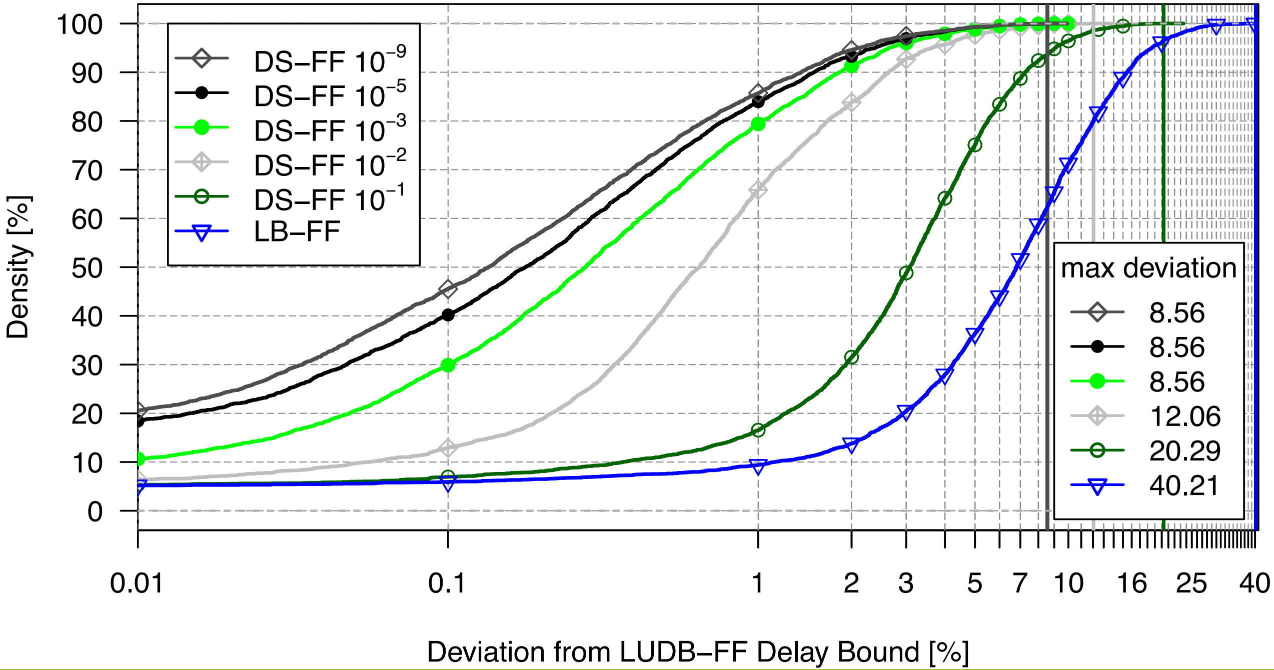

Evaluation (4)

- **Comparison of Runtimes**

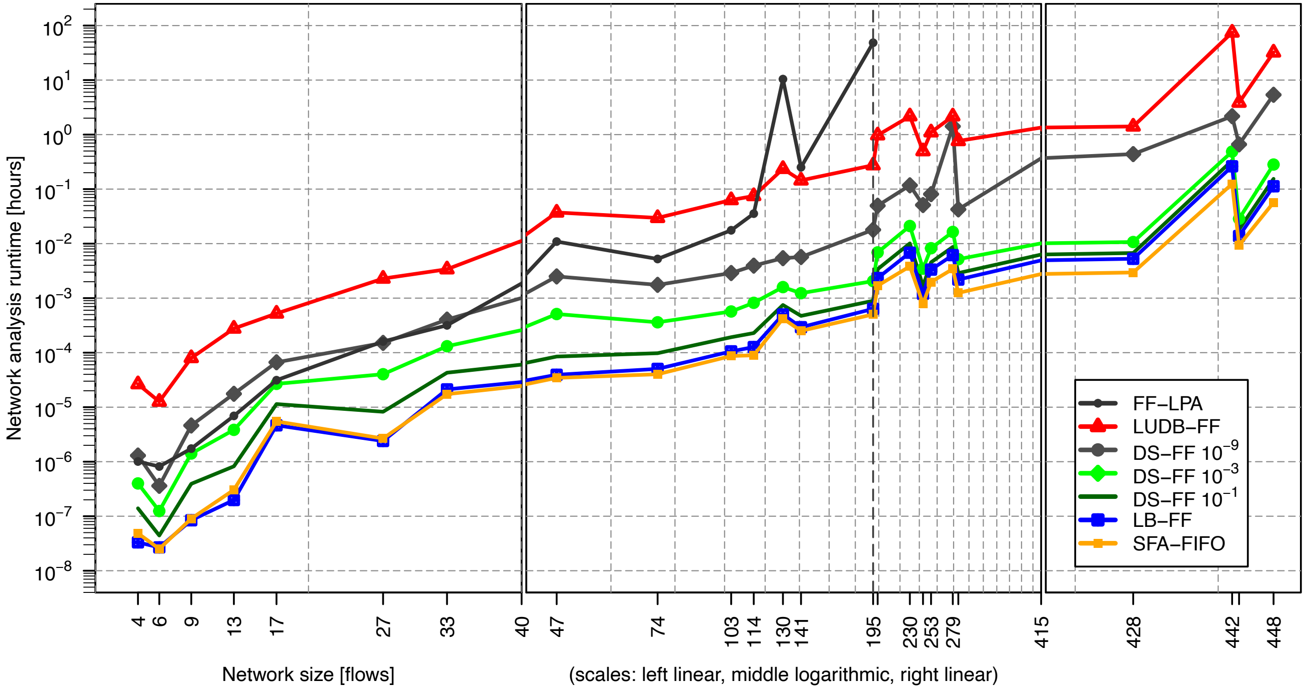

Conclusion

- **Novel and scalable FIFO-aware analysis techniques for feedforward networks LB-FF and DS-FF**
- **LB-FF sets all FIFO parameters statically**
- **DS-FF employs a search on-top of LB-FF**
- **DS-FF ($\epsilon = 10^{-3}$) is on average 120 times faster than LUDB-FF**
- **DS-FF ($\epsilon = 10^{-3}$) delay bounds are worse by only 0.57% on average and maximally 8.56% compared to LUDB-FF**
- **LUDB-FF restricted to token-bucket rate-latency curves while LB-FF and DS-FF can be applied to more complex curves, e.g., staircase functions**

Thanks for your attention!

Questions?

References

[Bisti08] Bisti, L., Lenzini, L., Mingozzi, E., Stea, G.: Estimating the worst-case delay in FIFO tandems using network calculus. In: Proceedings of the ICST ValueTools (2008)

[Bouillard15] Anne Bouillard and Giovanni Stea. 2015. Exact Worst-Case Delay in FIFO- Multiplexing Feed-Forward Networks. IEEE/ACM Transactions on Networking 23, 5 (2015), 1387–1400.

[Bouillard22] Anne Bouillard. 2022. Trade-off between accuracy and tractability of network calculus in FIFO networks. Elsevier Performance Evaluation 153 (2022), 102250.

[Boyer08] Boyer, Marc, and Christian Fraboul. "Tightening end to end delay upper bound for AFDX network calculus with rate latency FIFO servers using network calculus." 2008 IEEE International Workshop on Factory Communication Systems. IEEE, 2008.

[Cruz91] R. L. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. In IEEE Transactions on Information Theory, 1991, and R. L. Cruz. A Calculus for Network Delay, Part II: Network Analysis. In IEEE Transactions on Information Theory, 1991.

[LeBoudec01] Le Boudec, J.Y., Thiran, P.: Network Calculus: A Theory of Deterministic Queuing Systems for the Internet. Springer-Verlag, Berlin (2001)

[Scheffler21] Alexander Scheffler and Steffen Bondorf. 2021. Network Calculus for Bounding Delays in Feedforward Networks of FIFO Queueing Systems. In Proc. of the International Conference on Quantitative Evaluation of Systems (QEST).