Distributed Computer Systems Lab
https://disco.cs.uni-kl.de

TECHNISCHE UNIVERSITÄT
KAISERSLAUTERN

Dealing with Dependence in Stochastic Network Calculus

Using Independence as a Bound

Paul Nikolaus, Jens Schmitt, Florin Ciucu

WoNeCa 2020
Saarbrücken

Outline

1. Dependence in SNC with Moment-Generating Functions
2. Assumptions
3. Negative Dependence
4. Using Independence as a Bound in the SNC Analysis
5. Case Studies

SNC with MGFs Leads to Tighter Bounds

- Two Branches
  - Tail bounds / envelope functions
  - Moment generating functions (MGF)

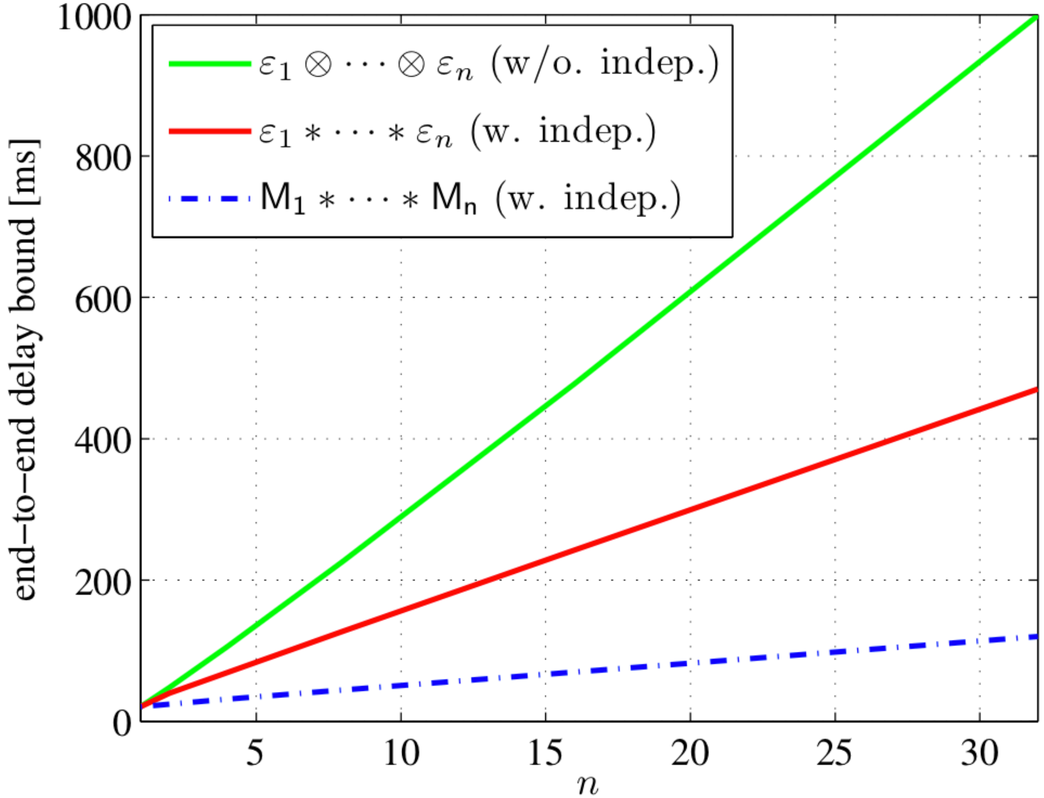

Approach with MGFs lead to tighter delay bounds
under independence [Rizk and Fidler, 2011]

Problem Statement

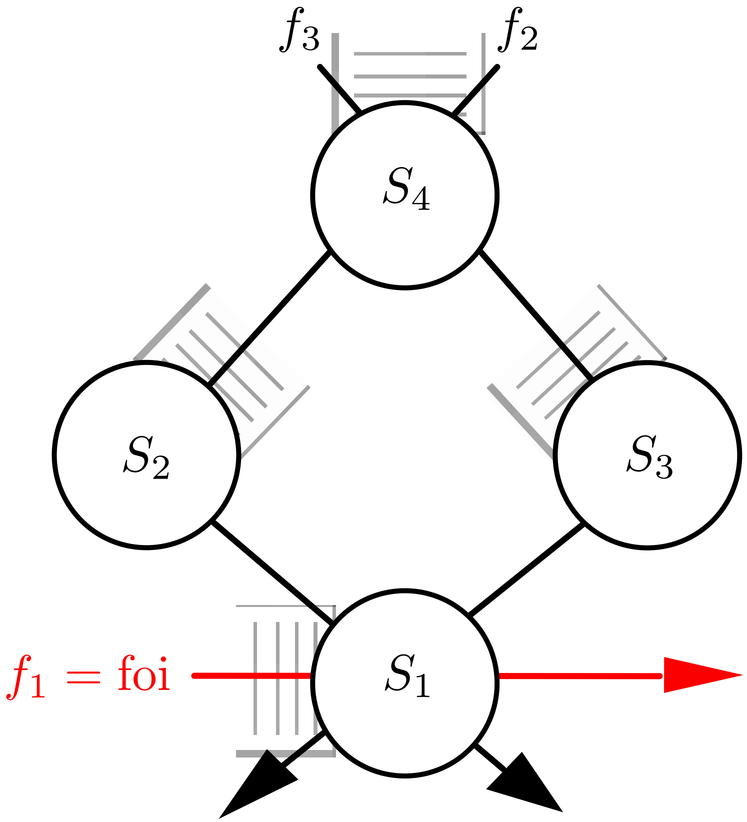

- Assume priority scheduling at server S_4
- Departure processes of flows f_2 and f_3 are dependent even if their arrival were assumed to be independent

Problem Statement

- If processes are dependent:

  Standard approach: Use **Hölder's inequality**:

$$\mathrm{E}\left[e^{\theta(A_1(s,t)+A_2(s,t))}\right] \leq \mathrm{E}\left[e^{p\theta A_1(s,t)}\right]^{1/p} \cdot \mathrm{E}\left[e^{q\theta A_2(s,t)}\right]^{1/q}$$

  for $\frac{1}{p} + \frac{1}{q} = 1$ and $p, q \in [1, \infty]$

- Problems:
  - Possible significant loss of accuracy
  - Additional parameter to optimize
  - *Ignores the knowledge about scheduling and dependence structure*

Can we avoid Hölder's inequality?

Arrival Process and Dynamic S-Server

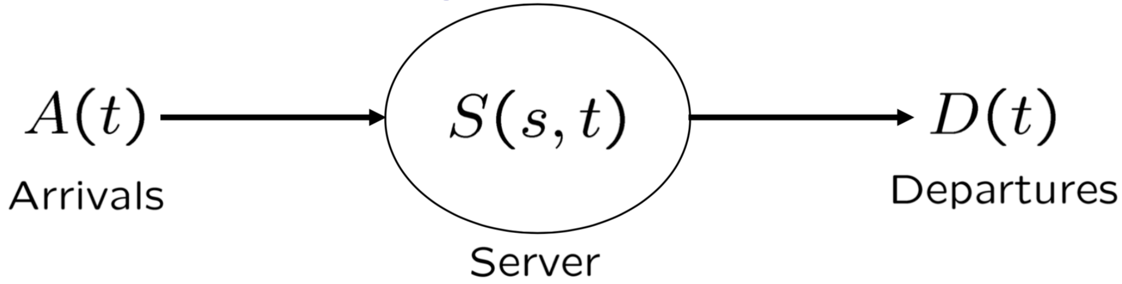

- Let A be an *discrete* arrival process, i.e., $A(s,t)$ is a stochastic process increasing in t such that

$$A(s,t) = \sum_{i=s+1}^{t} a_i$$

$a_i \geq 0$ and with existing moment generating function (MGF)

- Assume a dynamic S-server, i.e., $S(s,t)$ is a stochastic process increasing in t such that the departure process $D(0,t)$ is lower bounded:

$$D(0,t) \geq \inf_{0 \leq s \leq t} \{A(0,s) + S(s,t)\}$$

Bound on Delay Violation Probability

- Bound on the probability, that the delay at time t exceeds T [Chang, 2000]:

$$\mathrm{P}(d(t) > T) \leq \mathrm{E}\left[e^{\theta \sup_{0 \leq \tau \leq t}\{A(\tau,t)-S(\tau,t+T)\}}\right]$$

Idea of the Paper (in a Nutshell)

- Let two arrivals, A_1 and A_2, be multiplexed at one server
- Each sends a packet with a probability $p \in (0,1)$
- Server can serve exactly one packet, but A_2 is strictly prioritized

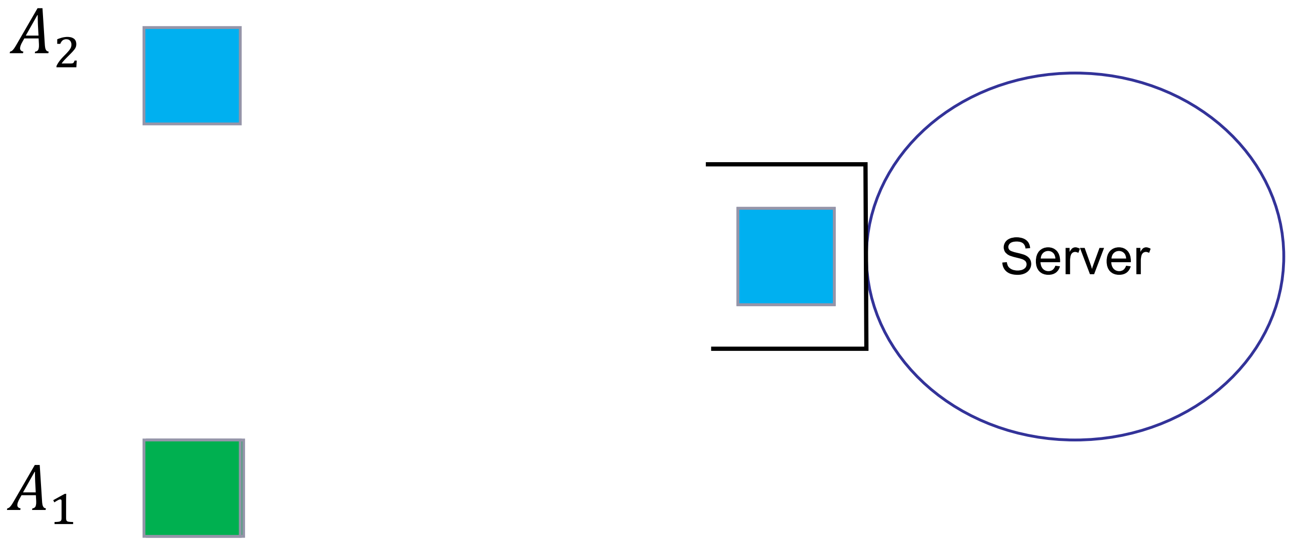

Idea of the Paper (in a Nutshell)

- For the output, we observe

$$\mathrm{P}$$

This dependence caused by scheduling (strict priority) is good:

… but

$$\mathrm{P}(D_1 = 1) \cdot \mathrm{P}(D_2 = 1)$$

$$\ldots (\text{Conditioning})$$

$$= (\quad {}_2 = 0)) \cdot \mathrm{P}(D_2 = 1)$$

$$= p$$

$$> 0$$

Independent scenario is an upper bound!

and thus

$$\mathrm{P}(D_1 = 1, D_2 = 1) < \mathrm{P}(D_1 = 1) \cdot \mathrm{P}(D_2 = 1)$$

Negative Dependence

Definition (Negative Dependence [Lehmann, 1966])

A finite family of random variables $\{X_1, \ldots, X_n\}$ is said to be *negatively (orthant) dependent (ND)* if the two following inequalities hold:

$$P\left(X_1 \leq x_1, \ldots, X_n \leq x_n\right) \leq \prod_{i=1}^{n} P\left(X_i \leq x_i\right),$$

$$P\left(X_1 > x_1, \ldots, X_n > x_n\right) \leq \prod_{i=1}^{n} P\left(X_i > x_i\right)$$

Negative Dependence and MGFs

Lemma ([Joag-Dev and Proschan, 1983])

If $\{X_1, \ldots, X_n\}$ is a set of ND random variables, then for any $\theta > 0$,

$$\mathrm{E}\left[e^{\theta \sum_{i=1}^{n} X_i}\right] \leq \prod_{i=1}^{n} \mathrm{E}\left[e^{\theta X_i}\right]$$

Challenge of Proving RVs to be ND

- Problem: Proving that random variables are negatively dependent is a challenging task!

- Some results exist:
  - “Zero-One Lemma” [Dubhashi and Ranjan, 1998]: If $X_1, \ldots, X_n \in \{0,1\}$ such that $\sum_i X_i = 1$, then they are ND (proves that output process in previous example are ND)
  - Permutation distribution, therefore random sampling without replacement, is ND [Joag-Dev and Proschan, 1983]

- The latter result is used to prove near-perfect load balancing for switches called “Sprinklers” [Ding et al., 2014]

Conjecture the Output to be ND

Conjecture

Let two independent flows with according arrival processes A_1 and A_2 traverse a work-conserving server with finite capacity. Assume both arrivals to have iid increments.

Then, we assume their respective output processes $D_1(s,t)$ and $D_2(s,t)$ to be ND for all $0 \leq s \leq t$.

- We do not have a proof…
  - ... but used 10^6 samples to estimate the joint and product (C)CDF of the output processes
  - over 5500 different combinations of intervals, x_1, x_2, (as in the CDF), utilizations (between 0.4 and 0.9) and
  - random packet sizes that were drawn from either exponential, Weibull, Gumpel, or log-normal distribution

The conjecture held in all experiments

Diamond Network

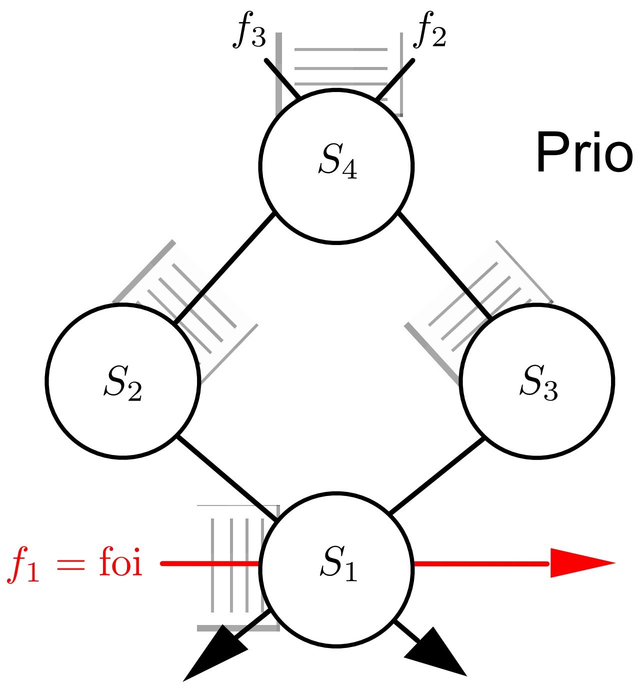

Priorities: foi $\leq f_2 \leq f_3$

- Assume work-conserving constant-rate servers
- S_2 and S_3 with equal rate

- We obtain for the standard bound

$$S_{\text{e2e}} = \left[S_1 - \left(\left(\left(A_2 \oslash [S_4 - A_3]^+\right) \oslash S_2\right) + ((A_3 \oslash S_4) \oslash S_3)\right)\right]^+$$

where $D_i^{(j)}$ is the output of flow i at server j

- For the approach using negative dependence, we use

$$S_{\text{e2e}} = \left[S_1 - \left(D_2^{(2)} + D_3^{(3)}\right)\right]^+$$

Diamond Network (Deriving a Delay Bound)

$$
\begin{aligned}
\mathrm{P}(d(t) > T) &\overset{\text{(Chernoff)}}{\leq} \mathrm{E}\left[e^{\theta A_1 \oslash S_{\mathrm{e2e}}(t+T,t)}\right] \\
&\overset{\text{(Boole)}}{\leq} \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta(A_1(\tau_1,t)-S_{\mathrm{e2e}}(\tau_1,t+T))}\right] \\
&\overset{\text{(Indep)}}{=} \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta A_1(\tau_1,t)}\right] \mathrm{E}\left[e^{-\theta\left[S_1-\left(D_2^{(2)}+D_3^{(3)}\right)\right]^+(\tau_1,t+T)}\right] \\
&\leq \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta A_1(\tau_1,t)}\right] e^{-\theta c_1(t+T-\tau_1)} \boxed{\mathrm{E}\left[e^{\theta\left(D_2^{(2)}+D_3^{(3)}\right)(\tau_1,t+T)}\right]}
\end{aligned}
$$

Diamond Network (Deriving a Delay Bound, cont.)

- Now, we can use the conjecture:

$$\begin{aligned}
& \mathrm{E}\left[e^{\theta\left(D_2^{(2)}+D_3^{(3)}\right)(\tau_1,t+T)}\right] \\
\overset{(\mathrm{ND})}{\leq}\; & \mathrm{E}\left[e^{\theta D_2^{(2)}(\tau_1,t+T)}\right]\mathrm{E}\left[e^{\theta D_3^{(3)}(\tau_1,t+T)}\right] \\
\overset{(\text{Output Bound})}{\leq}\; & \mathrm{E}\left[e^{\theta\left(\left(A_2\oslash[S_4-A_3]^+\right)\oslash S_2\right)(\tau_1,t+T)}\right]\mathrm{E}\left[e^{\theta((A_3\oslash S_4)\oslash S_3)(\tau_1,t+T)}\right]
\end{aligned}$$

Diamond Network (Deriving a Delay Bound, cont.)

- Assuming the arrivals to be (σ_A, ρ_A)-bounded, we obtain the time-independent bound

$$\mathrm{P}(d(t) > T) \leq \frac{e^{\theta\left(\left(\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-c_1\right)T+\sigma_1(\theta)+\sigma_{A_2}(\theta)+2\sigma_{A_3}(\theta)\right)}}{1-e^{\theta\left(\rho_{A_1}(\theta)+\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-c_1\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_2}(\theta)-c_2\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_3}(\theta)-c_3\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-c_4\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_3}(\theta)-c_4\right)}}$$

- No additional Hölder parameter to be optimized

Numerical Evaluation

- 10^4 Monte-Carlo simulations to sample parameter space for
  - Server rates
  - Exponentially distributed packet sizes

- Filtering such that utilization at server $S_1 \in [0.5,1)$

- Conducted parameter optimization of θ and Hölder parameter using a grid search followed by a downhill simplex algorithm (SciPy)

- Compute the improvement factor $\frac{\text{standard bound}}{\text{independence bound}}$

Diamond Network

Metric	Value
Number of analyzed scenarios	507
Number of improved delay bounds	389
Median (mean) of delay bound improvement factor	6.08 (2179.5)
Median (mean) of run time improvement factor	285.7 (290.5)

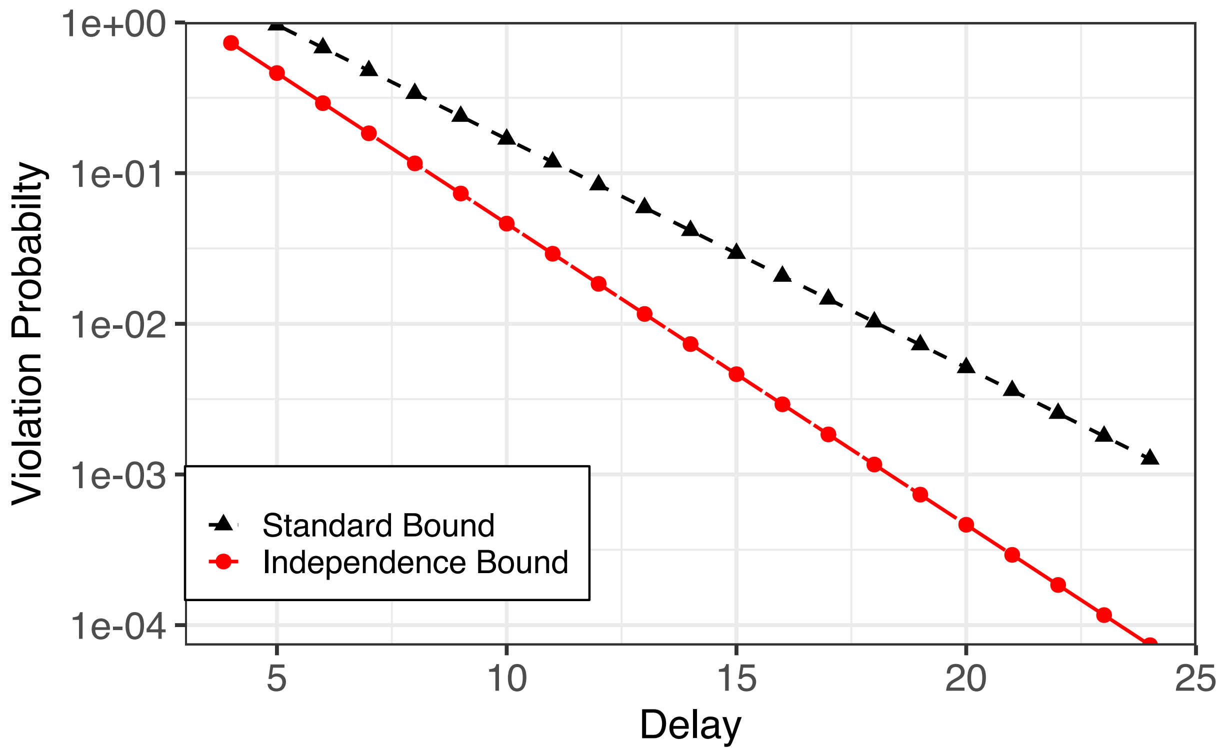

The $\mathbb{L}$

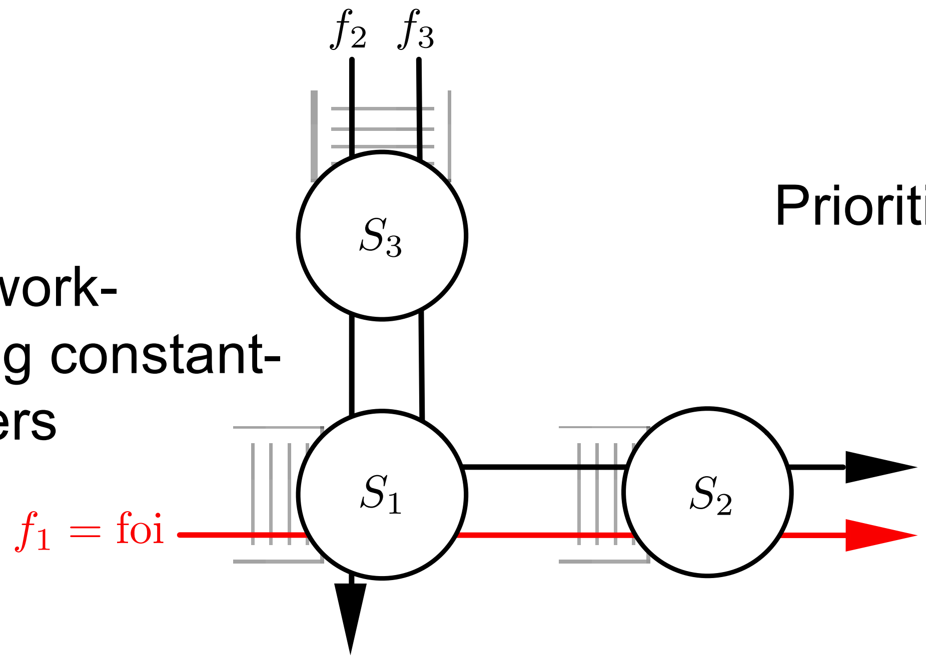

Priorities: $\mathrm{foi} \leq f_3 \leq f_2$

- Assume work-conserving constant-rate servers

Crucial difference: The convolution forces us to analyze the output processes at different intervals

- We receive

$$S_{\mathrm{e2e}} = \left[\left([S_1 - (A_2 \oslash S_3)]^+ \otimes S_2 \right) - \left(A_3 \oslash [S_3 - A_2]^+ \right) \right]^+$$

- and

$$S_{\mathrm{e2e}} = \left[\left(\left[S_1 - D_2^{(3)} \right]^+ \otimes S_2 \right) - D_3^{(3)} \right]^+$$

where, $D_i^{(j)}$ is the output of flow i at server j

The $\mathbb{L}$ (Deriving a Delay Bound)

$$
\begin{aligned}
&\mathrm{P}(d(t) > T) \\
&\leq \mathrm{E}\left[e^{\theta A_1 \oslash S_{\mathrm{net}}\,(t+T,t)}\right] \\
&\leq \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta(A_1(\tau_1,t)-S_{\mathrm{net}}(\tau_1,t+T))}\right] \\
&= \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta A_1(\tau_1,t)}\right] \mathrm{E}\left[e^{-\theta\left[\left(\left[S_1-D_2^{(3)}\right]^+ \otimes S_2\right)-D_3^{(3)}\right]^+(\tau_1,t+T)}\right] \\
&\vdots \\
&\leq \sum_{\tau_1=0}^{t} \mathrm{E}\left[e^{\theta A_1(\tau_1,t)}\right] \sum_{\tau_2=\tau_1}^{t+T} e^{-\theta c_1\cdot(\tau_2-\tau_1)} e^{-\theta c_2\cdot(t+T-\tau_2)}\, \mathrm{E}\left[e^{\theta D_3^{(3)}(\tau_1,t+T)} e^{\theta D_2^{(3)}(\tau_1,\tau_2)}\right]
\end{aligned}
$$

The $\mathbb{L}$ (Deriving a Delay Bound, cont.)

- Idea: leverage monotonicity and extend the interval

$$\mathrm{E}\left[e^{\theta D_3^{(3)}(\tau_1,t+T)} e^{\theta D_2^{(3)}\boxed{(\tau_1,\tau_2)}}\right] \leq \mathrm{E}\left[e^{\theta D_3^{(3)}(\tau_1,t+T)} e^{\theta D_2^{(3)}\boxed{(\tau_1,t+T)}}\right]$$

The $\mathbb{L}$ (Deriving a Delay Bound, cont.)

- Assuming the arrivals to be (σ_A, ρ_A)-bounded, we obtain the time-independent bound

$$\mathrm{P}(d(t) > T) \leq \frac{e^{\theta\left(\left(\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-\min\{c_1,c_2\}\right)\cdot T+\sigma_{A_1}(\theta)+2\sigma_{A_2}(\theta)+\sigma_{A_3}(\theta)\right)}}{1-e^{\theta\left(\rho_{A_1}(\theta)+\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-\min\{c_1,c_2\}\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_2}(\theta)-c_3\right)}} \cdot \frac{1}{1-e^{\theta\left(\rho_{A_2}(\theta)+\rho_{A_3}(\theta)-c_3\right)}} \cdot \frac{1}{1-e^{-\theta|c_1-c_2|}}$$

- Again, no additional Hölder parameter to be optimized

Numerical Evaluation

- 10^4 Monte-Carlo simulations to sample parameter space for
  - Server rates
  - Exponentially distributed packet sizes

- Filtering such that max utilization of S_1 and $S_2 \in [0.5,1)$

- Conducted parameter optimization of θ and Hölder parameter using a grid search followed by a downhill simplex algorithm (SciPy)

- Compute the improvement factor $\frac{\text{standard bound}}{\text{independence bound}}$

The $\mathbb{L}$

Metric	Value
Number of analyzed scenarios	729
Number of improved delay bounds	384
Median (mean) of delay bound improvement factor	1.27 (101.3)
Median (mean) of run time improvement factor	429.1 (474.5)

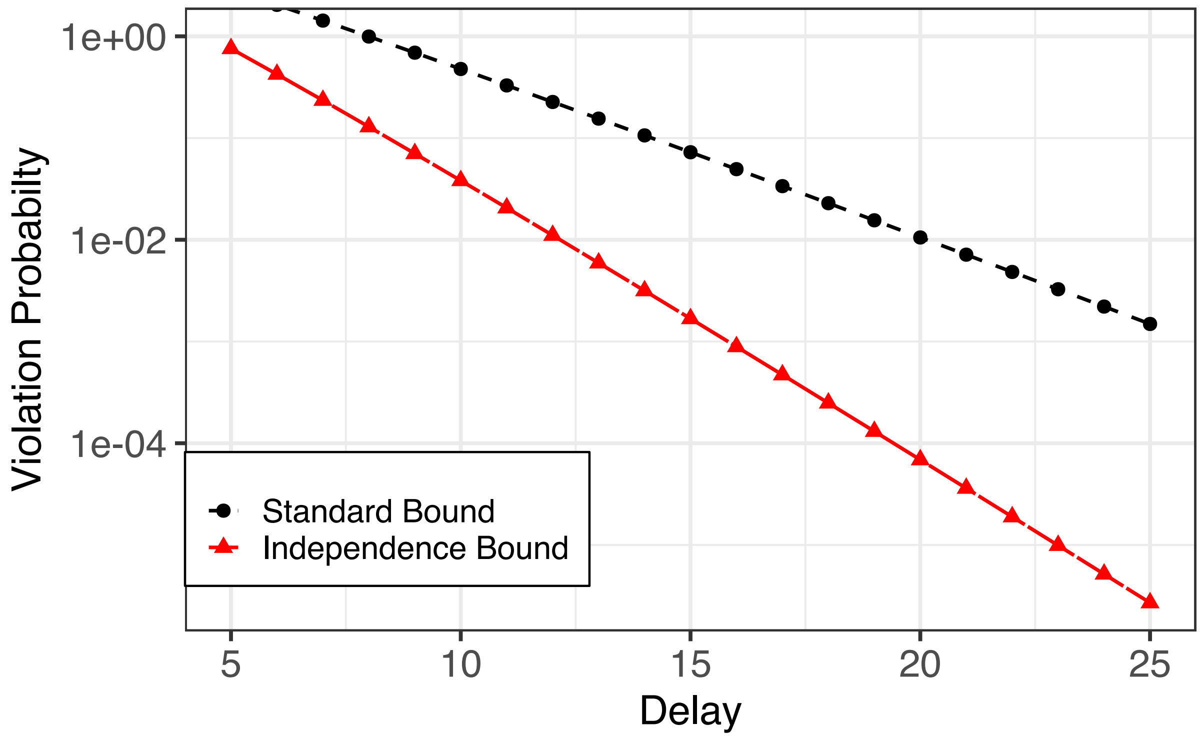

Discussion

- Make use of negative dependence in the stochastic network calculus
- Improve delay bounds in many cases depending on the topology
- Significant improvement of the run time, since less parameters need to be optimized
- Results are based on a conjecture (need scenario in which it can be proved rigorously)
- More scenarios in which the negative dependence can be exploited (large-scale experiments)

Thank you for your attention!

References

- Chang, C.-S. (2000).
  Performance guarantees in communication networks.
  Springer, London.
- Ding, W., Xu, J., Dai, J., Song, Y., and Lin, B. (2014).
  Sprinklers: A randomized variable-size striping approach to reordering-free load-balanced switching.
  In Proc. 10th ACM International on Conference on emerging Networking Experiments and Technologies, pages 89–100.
- Dubhashi, D. and Ranjan, D. (1998).
  Balls and bins: A study in negative dependence.
  Random Structures & Algorithms, 13(2):99–124.
- Joag-Dev, K. and Proschan, F. (1983).
  Negative association of random variables with applications.
  The Annals of Statistics, 11(1):286–295.
- Lehmann, E. L. (1966).
  Some concepts of dependence.
  The Annals of Mathematical Statistics, pages 1137–1153.

References

- Rizk, A. and Fidler, M. (2011).
  Leveraging statistical multiplexing gains in single-and multi-hop networks.
  In Proc. IEEE IWQoS ’11, pages 1–9.