distributed computer systems

TECHNISCHE UNIVERSITÄT KAISERSLAUTERN

Window Flow Controller and Subadditivity

Michael A. Beck
TU Kaiserslautern

WoNeCa 2016

Introduction

What is Window Flow Control?

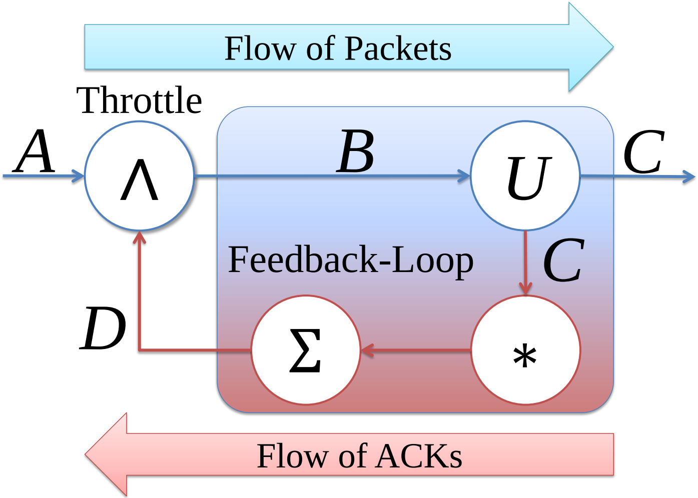

- Most important example: window-based transport protocols

End-to-end Description

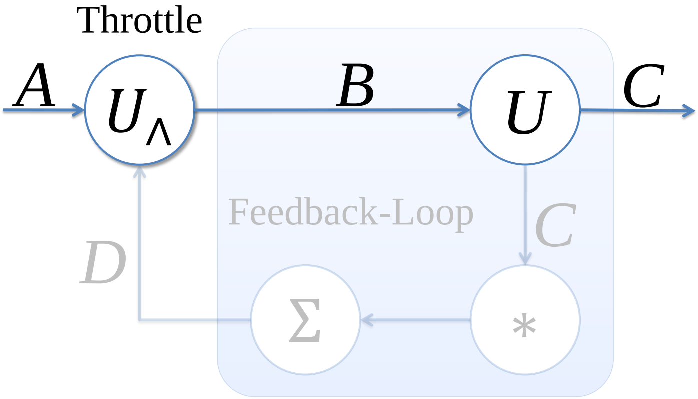

$$U_\wedge = \bar{U}_{fb} := \bigwedge_{n=0}^{\infty} U_{fb}^{(n)}(t)$$

Subadditivity

Definition

U is *subadditive* if

$$U(s) + U(t) \geq U(t + s)$$

$$\Big(U(s, r) + U(r, t) \geq U(s, t)\Big)$$

"The whole is **smaller** than the sum of its parts"

Subadditivity

Definition

U is *subadditive* if

$$U(s) + U(t) \geq U(t+s)$$

$$\Big(U(s,r) + U(r,t) \geq U(s,t)\Big)$$

"The whole is **smaller** than the sum of its parts"

- If U is subadditive its subadditive closure equals $U \wedge \mathbf{1}$. Where $\mathbf{1}$ is the neutral element of the min-plus convolution. ($\mathbf{1}(t) = \infty$ for all $t > 0$).

Subadditivity

Definition

U is *subadditive* if

$$U(s) + U(t) \geq U(t + s)$$

$$\Big(U(s, r) + U(r, t) \geq U(s, t) \Big)$$

"The whole is **smaller** than the sum of its parts"

- If U is subadditive its subadditive closure equals $U \wedge \mathbf{1}$. Where $\mathbf{1}$ is the neutral element of the min-plus convolution. ($\mathbf{1}(t) = \infty$ for all $t > 0$).
- Convolution **does not** preserve subadditivity.
- The calculation of leftover service curves **does not** preserve subadditivity.
- Rule of thumb: "Interesting things are not subadditive."

MGF Calculus

MGF-Bounds

- MGF-Bounds live on a bivariate formulation of network calculus.
- Throttle's service is still the subadditive closure $\bar{U}_{fb}$.

MGF-Bounds

- MGF-Bounds live on a bivariate formulation of network calculus.
- Throttle's service is still the subadditive closure $\bar{U}_{fb}$.
- In the calculation of performance bounds the MGF

$$\mathbb{E}\left(e^{\theta A \oslash (\bar{U}_{fb} \otimes U)(s,t)}\right) \leq \sum_{n=0}^{\infty} \mathbb{E}\left(e^{\theta A \oslash (U_{fb}^{(n)} \otimes U)(s,t)}\right)$$

must be bounded.

MGF-Bounds

- MGF-Bounds live on a bivariate formulation of network calculus.
- Throttle's service is still the subadditive closure $\bar{U}_{fb}$.
- In the calculation of performance bounds the MGF

$$\mathbb{E}\left(e^{\theta A \oslash (\bar{U}_{fb} \otimes U)(s,t)}\right) \leq \sum_{n=0}^{\infty} \mathbb{E}\left(e^{\theta A \oslash (U_{fb}^{(n)} \otimes U)(s,t)}\right)$$

  must be bounded.
- However: $\sum_{n=0}^{\infty} \mathbb{E}\left(e^{\theta A \oslash (U_{fb}^{(n)} \otimes U)(s,t)}\right) = \infty$.

First Successes

Method 0: Bivariate Calculus

- Subadditivity leads to a closed form of the subadditive closure: $\bar{U}_{fb} = U_{fb} \wedge \mathbf{1}$.
- First idea: Consider subadditive feedback loops.
- In bivariate formulations leftover service descriptions preserve subadditivity (no arrival curves involved).

Method 0: Bivariate Calculus

- Subadditivity leads to a closed form of the subadditive closure: $\bar{U}_{fb} = U_{fb} \wedge \mathbf{1}$.
- First idea: Consider subadditive feedback loops.
- In bivariate formulations leftover service descriptions preserve subadditivity (no arrival curves involved).

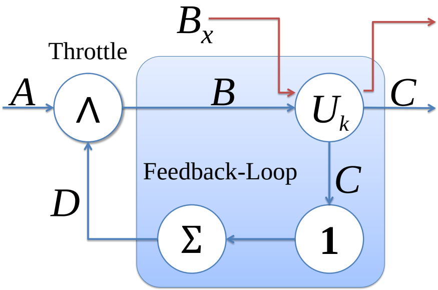

$$U_k(s,t) = k(t-s)$$

Method 1: Change Topologies

- Part of “Window Flow Control in Stochastic Network Calculus” (TR, 2015)
- By changing the topology subadditivity can be enforced.
- Can be costly.

Method 1: Change Topologies

- Part of "Window Flow Control in Stochastic Network Calculus" (TR, 2015)
- By changing the topology subadditivity can be enforced.
- Can be costly.

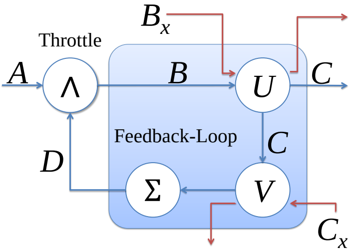

Method 1: Change Topologies

- Part of "Window Flow Control in Stochastic Network Calculus" (TR, 2015)
- By changing the topology subadditivity can be enforced.
- Can be costly.

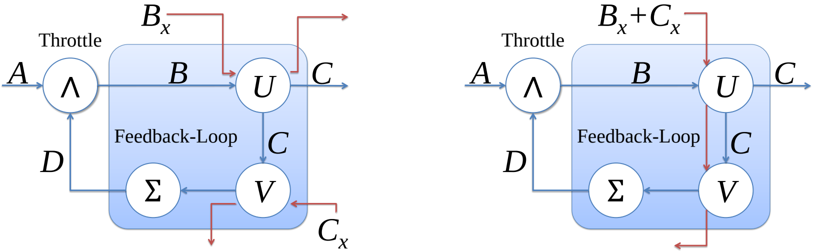

- Nevertheless: first non-trivial performance bounds on WFC-Systems!
- Works better if:
  - service elements have **similar** rates
  - each service element can work the entire aggregate of crossflows

Method 2: "Assume" Subadditivity

- "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- A condition for subadditivity:

$$(E): \; U \otimes V(s,t) - (U \otimes V)^{(2)}(s,t) \leq \Sigma \text{ for all } s \leq t$$

$$\Rightarrow \quad U_{fb} \text{ is subadditive.}$$

Method 2: "Assume" Subadditivity

- "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- A condition for subadditivity:

$$(E):\ U \otimes V(s,t) - (U \otimes V)^{(2)}(s,t) \leq \Sigma \text{ for all } s \leq t$$
$$\Rightarrow \quad U_{fb} \text{ is subadditive.}$$

- If the probability of $(\neg E)$ can be bounded the WFC system can be analyzed as well

Method 2: "Assume" Subadditivity

- "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- A condition for subadditivity:

$$
\begin{aligned}
&(E):\ U \otimes V(s,t) - (U \otimes V)^{(2)}(s,t) \leq \Sigma \text{ for all } s \leq t \\
\Rightarrow \quad & U_{fb} \text{ is subadditive.}
\end{aligned}
$$

- If the probability of $(\neg E)$ can be bounded the WFC system can be analyzed as well

$$
\mathbb{P}(\mathfrak{d}_{sys}(t) > T) \leq \mathbb{P}(\mathfrak{d}_{sys}(t) > T \mid U_{fb} \text{ is subadditive}) + \mathbb{P}(\neg E)
$$

Method 2: "Assume" Subadditivity

- "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- A condition for subadditivity:

$$
\begin{aligned}
&(E):\ U \otimes V(s,t) - (U \otimes V)^{(2)}(s,t) \leq \Sigma \text{ for all } s \leq t \\
\Rightarrow\quad &U_{fb} \text{ is subadditive.}
\end{aligned}
$$

- If the probability of $(\neg E)$ can be bounded the WFC system can be analyzed as well

$$\mathbb{P}(\mathfrak{d}_{sys}(t) > T) \leq \mathbb{P}(\mathfrak{d}_{sys}(t) > T \mid U_{fb} \text{ is subadditive}) + \mathbb{P}(\neg E)$$

- Works better for
  - **differing** service elements
- The catch: $\mathbb{P}(\neg E)$ diverges in t!

SNC at Work

General Case

- From "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- Throttled vs. Unthrottled:

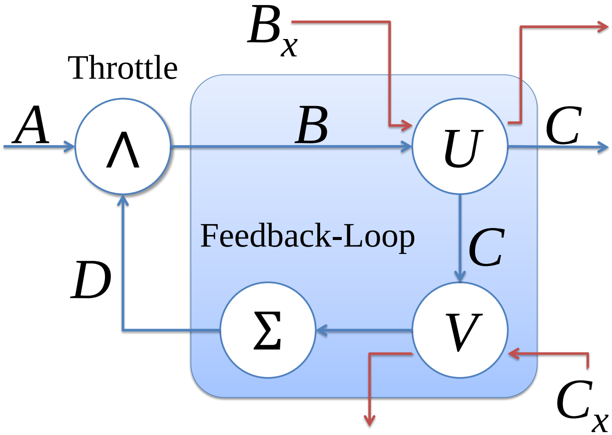

General Case (throttled vs. unthrottled)

- From "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- I.i.d. exp(λ) increments for all flows involved.

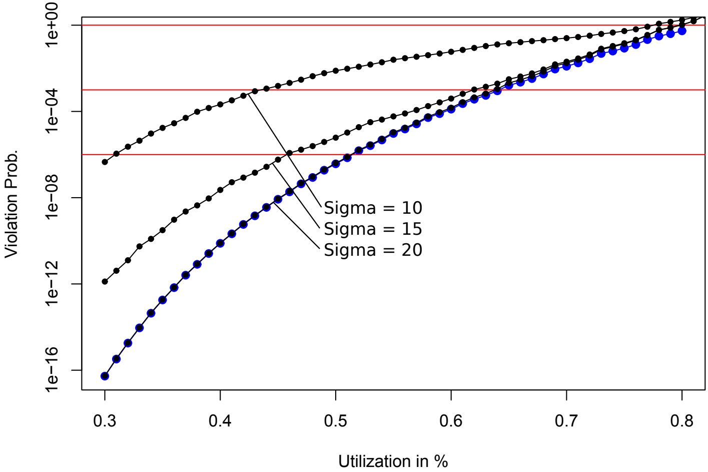

Convergence to Unthrottled Systems?

- From "Window Flow Control in Stochastic Network Calculus – The General Case" (Valuetools 2015)
- Bounds for improving window sizes Σ:

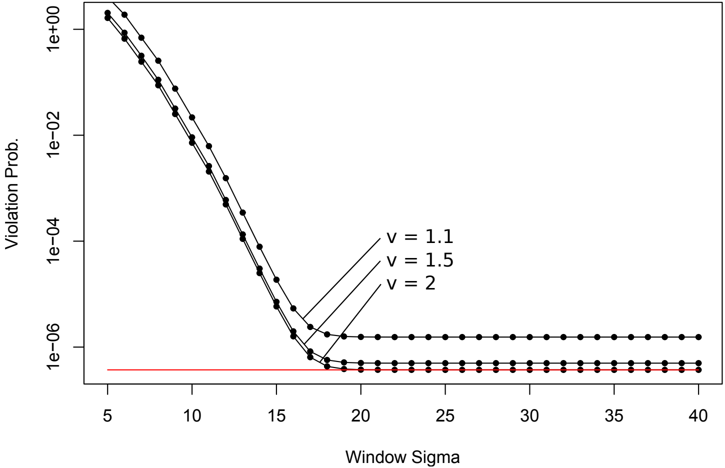

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Given the same scenario we can analyze it deterministically or by the two stochastic methods.
- Admission problem:

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Given the same scenario we can analyze it deterministically or by the two stochastic methods.
- Admission problem:
  - Admitting one subflow of A increases utilization at U by 1%.
  - U is utilized by a certain amount of cross-flows already.

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Given the same scenario we can analyze it deterministically or by the two stochastic methods.
- Admission problem:
  - Admitting one subflow of A increases utilization at U by 1%.
  - U is utilized by a certain amount of cross-flows already.
  - How many flows can we admit to the system without breaking a given (probabilistic) backlog bound at the throttle element?

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Reachable utilizations: DNC (up to 52%), by topology change (up to 65%), and by bounding $\mathbb{P}(\neg E)$ (up to 90%, $t = 1000$)

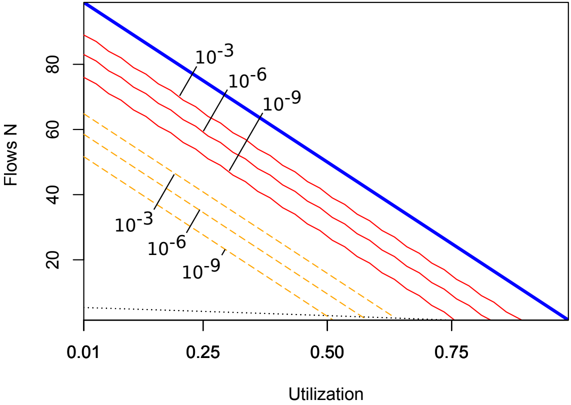

Asymmetric situation: V runs faster than U.

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Now V and U have the same rate. V handles 10 crossflows.

DNC vs. Topology Changes vs. $\mathbb{P}(\neg E)$

- From "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (under review)
- Now V and U have the same rate. V handles 10 crossflows.

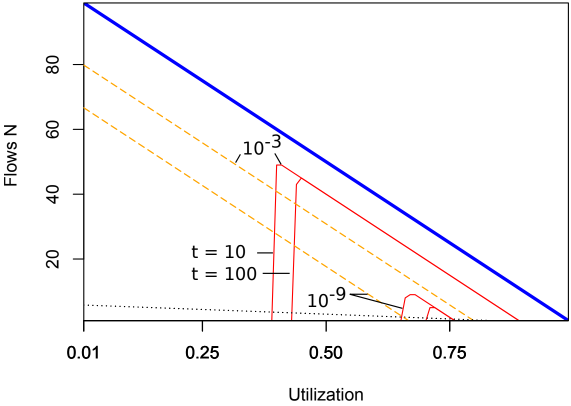

- For a low number of crossflows at U the service elements are very similar! →bounding $\mathbb{P}(\neg E)$ becomes harder.

Conclusion

Conclusion

- WFC is a hard problem in SNC.
- Recent work in *bivariate* formulations gives first results, though.
- Current approaches either:

Conclusion

- WFC is a hard problem in SNC.
- Recent work in *bivariate* formulations gives first results, though.
- Current approaches either:
  - consider specific feedback loops only (fixed delay elements)

Conclusion

- WFC is a hard problem in SNC.
- Recent work in *bivariate* formulations gives first results, though.
- Current approaches either:
  - consider specific feedback loops only (fixed delay elements)
  - underestimate available service rates (topological changes)

Conclusion

- WFC is a hard problem in SNC.
- Recent work in *bivariate* formulations gives first results, though.
- Current approaches either:
  - consider specific feedback loops only (fixed delay elements)
  - underestimate available service rates (topological changes)
  - depend on the evaluation time t (bounding $\mathbb{P}(\neg E)$)

Conclusion

- WFC is a hard problem in SNC.
- Recent work in *bivariate* formulations gives first results, though.
- Current approaches either:
  - consider specific feedback loops only (fixed delay elements)
  - underestimate available service rates (topological changes)
  - depend on the evaluation time t (bounding $\mathbb{P}(\neg E)$)
- Methods complement one another to some extent.

Thank you for your attention!

For details:

- "Window Flow Control in Stochastic Network Calculus" (Beck and Schmitt. TR, 2015)
- "Window Flow Control in Stochastic Network Calculus – The General Case" (Beck and Schmitt, Valuetools 2015)
- "Stochastic Worst Case Analysis of Window Flow Controlled Systems" (Beck. Under review)