Towards a Statistical Network Calculus - Dealing with Uncertainties in Arrivals

Michael A. Beck, Sebastian A. Henningsen, Simon B. Birnbach, Jens B. Schmitt

WoNeCa 2014/MMB 2014

Motivation

Motivation

Why StatNC?

- Apply Stochastic Network Calculus in a given (new) setup.
- To do so, we need to make assumptions about the processes involved, especially the arrivals.

Motivation

Why StatNC?

- Apply Stochastic Network Calculus in a given (new) setup.
- To do so, we need to make assumptions about the processes involved, especially the arrivals.

- How accurate or representative are our measurements?
- Can we be certain about the deduced arrival model?
- What if our assumption or measurements are wrong? With what probability are they wrong?

Motivation

Why StatNC?

- Apply Stochastic Network Calculus in a given (new) setup.
- To do so, we need to make assumptions about the processes involved, especially the arrivals.

- How accurate or representative are our measurements?
- Can we be certain about the deduced arrival model?
- What if our assumption or measurements are wrong? With what probability are they wrong?

→ Statistical Network Calculus (StatNC) takes uncertainties about the traffic behaviour into account.

Motivation

Why StatNC?

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
    - How can we weave statistics into the framework of SNC ($\rightarrow$ sufficient condition)?

Motivation

Why StatNC?

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
    - How can we weave statistics into the framework of SNC ($\rightarrow$ sufficient condition)?
    - How much do we lose?

Motivation

Why StatNC?

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
  - How can we weave statistics into the framework of SNC ($\rightarrow$ sufficient condition)?
  - How much do we lose?
  - Are dynamic adaptations possible (detrend seasonal effects)?

Motivation

Why StatNC?

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
  - How can we weave statistics into the framework of SNC ($\rightarrow$ sufficient condition)?
  - How much do we lose?
  - Are dynamic adaptations possible (detrend seasonal effects)?
  - Can we achieve a higher flexibility?

Notations

Notations

- We use discrete time and allow data to be continuous: Negative indices denote past timeslots, positive indices upcoming timeslots.

Notations

- We use discrete time and allow data to be continuous: Negative indices denote past timeslots, positive indices upcoming timeslots.
- A flow is characterized by its non-negative increments $(a_k)_{k\in\mathbb{Z}}$:

$$A(m,n) := \sum_{k=m+1}^{n} a_k$$

$$\sup_{m\in\mathbb{Z}}\{\mathbb{E}(e^{\theta A(m,m+k)})\} \leq e^{k\theta\rho_A(\theta)+\theta\sigma_A(\theta)} \qquad k \in \mathbb{N}$$

Notations

- We use discrete time and allow data to be continuous: Negative indices denote past timeslots, positive indices upcoming timeslots.

- A flow is characterized by its non-negative increments $(a_k)_{k\in\mathbb{Z}}$:

$$A(m,n) := \sum_{k=m+1}^{n} a_k$$

$$\sup_{m\in\mathbb{Z}}\{\mathbb{E}(e^{\theta A(m,m+k)})\} \leq e^{k\theta\rho_A(\theta)+\theta\sigma_A(\theta)} \qquad k\in\mathbb{N}$$

- A dynamic S-Server fulfills for all $n\in\mathbb{N}_0$ and flows A:

$$D(0,n) \geq \min_{0\leq k\leq n}\{A(0,k)+S(k,n)\}$$

$$\sup_{m\geq 0}\{\mathbb{E}(e^{-\theta S(m,m+k)})\} \leq e^{k\theta\rho_S(\theta)+\sigma_S(\theta)} \qquad k\in\mathbb{N}$$

Notations

For A and S being stochastically independent the backlog $q(n) = A(0,n) - D(0,n)$ is bounded by:

$$\mathbb{P}(q(n) > x) \leq e^{-\theta x} \sum_{k=0}^{n} \mathbb{E}(e^{\theta A(k,n)})\mathbb{E}(e^{-\theta S(k,n)})$$
$$\leq e^{-\theta x} e^{\theta(\sigma_A(\theta)+\sigma_S(\theta))} \sum_{k=0}^{n} e^{k\theta(\rho_A(\theta)+\rho_S(\theta))}$$

Notations

For A and S being stochastically independent the backlog $q(n) = A(0, n) - D(0, n)$ is bounded by:

$$\mathbb{P}(q(n) > x) \leq e^{-\theta x} \sum_{k=0}^{n} \mathbb{E}(e^{\theta A(k,n)}) \mathbb{E}(e^{-\theta S(k,n)})$$

$$\leq e^{-\theta x} e^{\theta(\sigma_A(\theta)+\sigma_S(\theta))} \sum_{k=0}^{n} e^{k\theta(\rho_A(\theta)+\rho_S(\theta))}$$

How to deal with $\mathbb{E}(e^{\theta A(k,n)})$, if we face uncertainties about A?

Statistical Framework

A Sufficient Condition

...for merging SNC and a Statistic

Define $f \in \mathcal{F}$ if $f : \mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{R}^+ \to \mathbb{R}_0^+$ (example: The MGF of $A(\cdot, \cdot)$ evaluated at θ).

Theorem

Let $\Phi : \mathbb{R}^{|n_0|} \to \mathcal{F}$ *be a statistic on* $a := (a_{n_0}, \dots, a_{-1})$ *such that:*

$$\sup_{\theta \in (0, \theta^*)} \mathbb{P}\left(\bigcup_{m \leq n} \Phi(a)(m, n, \theta) < \mathbb{E}(e^{\theta A(m,n)}) \right) \leq \alpha$$

Then for all $n \in \mathbb{N}_0$, $\theta < \theta^*$

$$\mathbb{P}(q(n) > x) \leq \alpha + e^{-\theta x} \sum_{k=0}^{n} \Phi(a)(k, n, \theta) \mathbb{E}(e^{-\theta S(k,n)})$$

A Sufficient Condition

...for merging SNC and a Statistic

Theorem

Let $\Phi : \mathbb{R}^{|n_0|} \to \mathcal{F}$ *be a statistic on* $a := (a_{n_0}, \ldots, a_{-1})$ *such that:*

$$\sup_{\theta \in (0,\theta^*)} \mathbb{P}\left(\bigcup_{m \leq n} \Phi(a)(m,n,\theta) < \mathbb{E}(e^{\theta A(m,n)}) \right) \leq \alpha$$

Then for all $n \in \mathbb{N}_0$, $\theta < \theta^*$

$$\mathbb{P}(q(n) > x) \leq \alpha + e^{-\theta x} \sum_{k=0}^{n} \Phi(a)(k,n,\theta)\mathbb{E}(e^{-\theta S(k,n)})$$

$$\mathbb{P}(q(n) > x) \leq e^{-\theta x} \sum_{k=0}^{n} \mathbb{E}(e^{\theta A(k,n)})\mathbb{E}(e^{-\theta S(k,n)})$$

Examples

Examples

- Parametric Estimator
  - As example i.i.d. exponentially distributed increments

$$\Phi(a)(m,n,\theta) = \left(\frac{\bar{\lambda}}{\bar{\lambda}-\theta}\right)^{n-m} \geq \mathbb{E}(e^{\theta A(m,n)})$$

with $\bar{\lambda}$ being a parametric estimator on the rate parameter λ.

Examples

- Parametric Estimator
  - As example i.i.d. exponentially distributed increments

$$\Phi(a)(m,n,\theta) = \left(\frac{\bar{\lambda}}{\bar{\lambda}-\theta}\right)^{n-m} \geq \mathbb{E}(e^{\theta A(m,n)})$$

with $\bar{\lambda}$ being a parametric estimator on the rate parameter λ.

- Bandwidth-Limited i.i.d. Traffic
  - No knowledge on the increments distribution needed (very flexible)

$$\Phi(a)(m,n,\theta) = \left(\varepsilon(e^{\theta M}-1) + \frac{1}{|n_0|}\sum_{k=n_0}^{-1} e^{\theta a_k}\right)^{n-m}$$

Examples

- Parametric Estimator
  - As example i.i.d. exponentially distributed increments

$$\Phi(a)(m,n,\theta) = \left(\frac{\bar{\lambda}}{\bar{\lambda}-\theta}\right)^{n-m} \geq \mathbb{E}(e^{\theta A(m,n)})$$

  with $\bar{\lambda}$ being a parametric estimator on the rate parameter λ.

- Bandwidth-Limited i.i.d. Traffic
  - No knowledge on the increments distribution needed (very flexible)

$$\Phi(a)(m,n,\theta) = \left(\varepsilon(e^{\theta M}-1) + \frac{1}{|n_0|}\sum_{k=n_0}^{-1} e^{\theta a_k}\right)^{n-m}$$

- Markov-Modulated Arrivals
  - *On*- and *Off*-State, with unknown i.i.d. increments in *On*-state.
  - Φ is a combination of the above statistic and an estimation on the transition probabilities of the Markov chain.

StatNC at Work

StatNC at Work

Questions and Opportunities

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
  - How can we weave statistics into the framework of SNC ($\rightarrow$ necessary conditions)? Done!
  - **How much do we lose?**
  - Are dynamic adaptations possible (detrend seasonal effects)?
  - Can we achieve a higher flexibility?

StatNC at Work

The Price of StatNC

- For this we consider a Markov-modulated arrival with a peak-utilization of 98% and an average utilization of 49%.

StatNC at Work

The Price of StatNC

- For this we consider a Markov-modulated arrival with a peak-utilization of 98% and an average utilization of 49%.
- SNC uses full knowledge on the arrivals.
- StatNC does *neither* know the transition probabilities *nor* the distribution of the increments in the *On*-state.

StatNC at Work

The Price of StatNC

How much do we lose?

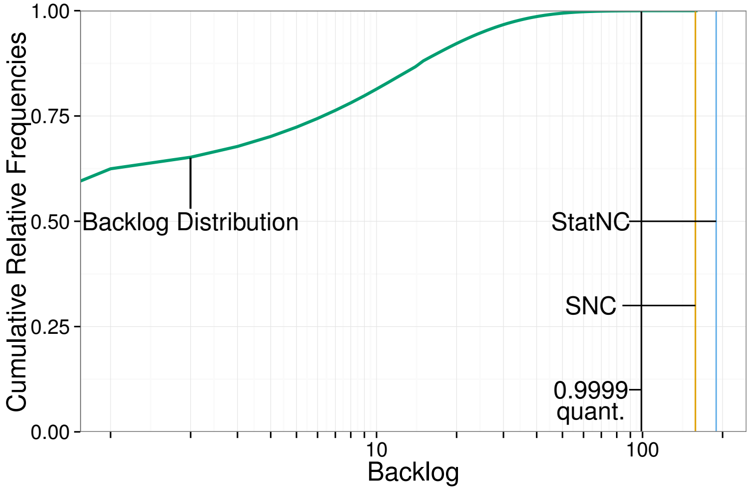

StatNC at Work

Questions and Opportunities

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
  - How can we weave statistics into the framework of SNC ($\rightarrow$ necessary conditions)? Done!
  - How much do we lose? Done!
  - **Are dynamic adaptations possible (detrend seasonal effects)?**
  - Can we achieve a higher flexibility?

StatNC at Work

Dynamic adaptation

- SNC encounters problems for flows with seasonal changes.
- E.g: A diminishing flow: Large initial increments result in loose bounds later.

StatNC at Work

Dynamic adaptation

- SNC encounters problems for flows with seasonal changes.
- E.g: A diminishing flow: Large initial increments result in loose bounds later.

- We consider a Markov-modulated arrival, but with *High*- and *Low*-states.
- Transition probabilities are very small $\rightarrow$ seasonal effects.
- StatNC uses an observation window, while SNC full knowledge on the arrivals.

StatNC at Work

Dynamic adaptation

Are dynamic adaptations possible?

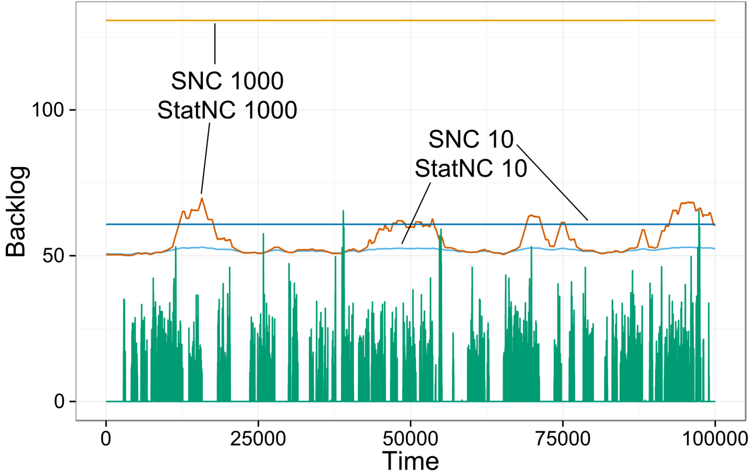

StatNC at Work

Questions and Opportunities

- Deriving a (statistical) model from observed arrivals rises new questions and opportunities:
  - How can we weave statistics into the framework of SNC ($\rightarrow$ necessary conditions)? Done!
  - How much do we lose? Done!
  - Are dynamic adaptations possible (detrend seasonal effects)? Done!
  - **Can we achieve a higher flexibility?**

StatNC at Work

Robustness

- When making assumptions about the arrivals there is always the possibility we might err!
- What effects can this have on SNC and StatNC?

StatNC at Work

Robustness

- When making assumptions about the arrivals there is always the possibility we might err!
- What effects can this have on SNC and StatNC?
- We consider a scenario, where SNC makes a false assumption (assumes exponential increments) on the arrivals (having Pareto increments).
- StatNC however makes no assumption about the actual distribution of the increments.

StatNC at Work

Robustness

What about the robustness of statistical methods?

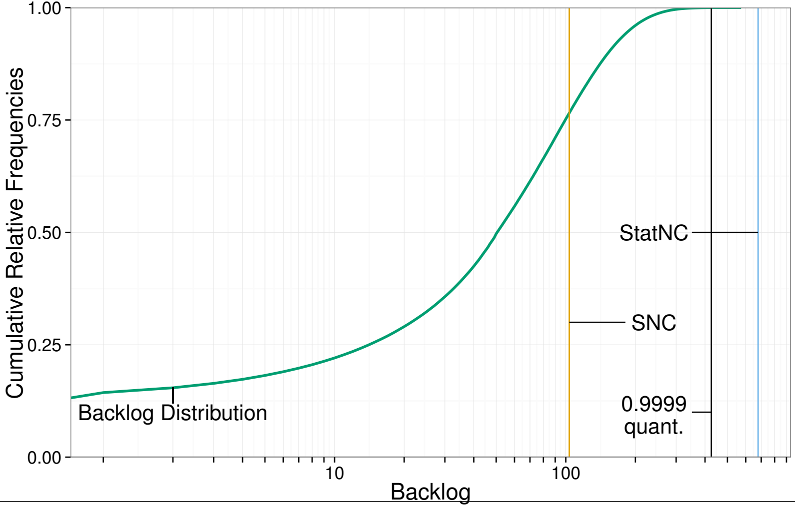

Conclusion

Conclusion

- We gave a framework, which takes uncertainties about arrivals directly into account.
- This merges statistical methods with SNC.
- StatNC can detrend seasonal effects much better, even if SNC has complete knowledge about the arrivals.
- The risk of making wrong assumptions on arrivals can be decreased by using robust statistics.
- Using statistical methods is an important step towards the applicability of SNC.

Thank you for your attention!