A First Course in Stochastic Network Calculus

Dipl.-Math. Michael Alexander Beck

November 14, 2013

Contents

Foreword

This *First Course in Stochastic Network Calculus* is a “First” course in two different perspectives. One can see it as a introductory course to Stochastic Network Calculus. On the other side it was my first course I have given about SNC to a few students in June 2012. It builds on the lecture *Performance Modelling in Distributed Systems [6]* at the University of Kaiserslautern. The concepts of stochastic network calclulus parallels those of deterministic network calculus. This is why I reference on the lecture of 2011 at several points to stress these connections. This document however is thought of a stand-alone course and hence a deep study of [6] is not necessary (but recommended).

This course contains a rather large probability primer to ensure the students can really grasp the expressions, which appear in stochastic network calculus. A student familiar with probability theory might skip this first chapter and delve directly into the stochastic network calculus. For each topic exercises are given, which can (and should) be used to strengthen the understanding of the presented definitions and theory.

This document is still in process and hopefully will evolve at some day into a fully grown course about Stochastic Network Calculus, providing a good overview over this exciting theory. Hence, please provide feedback to `beck@cs.uni-kl.de`.

- Michael Beck

1 Basics of Probability

2 Stochastic Arrivals and Service

In this chapter we carefully motivate the need for stochastic arrivals and service and see two approaches how such stochastic processes can be described and bounded. Each of them leads to its own stochastic extension of the deterministic network calculus (see [3, 4, 2, 5]). So we cannot talk about *the* stochastic network calculus, but rather have to distinguish between several approaches, each with its own characteristics and strengths.

Overall we use a discrete time-scale, since it keeps formulas somewhat easier.

2.1 Stochastic Arrivals

As in deterministic network calculus we abstract data arriving at some service element as flows. The service element is abstracted as a node, which takes the arriving flow as input and relays it to a departing flow under a certain rate and scheduling policy. A node can have many different arriving flows and produces the same number of departing flows, each of them corresponding to one arrival flow. To describe such a flow we use cumulative functions satisfying:

$$\begin{aligned} A(n) &\geq 0 \qquad \forall\, n \in \mathbb{N}_0 \\ A(n) - A(m) &\geq 0 \qquad \forall\, n \geq m \in \mathbb{N}_0 \end{aligned} \tag{2.1}$$

Here the number of packets/data/bits/"whatever is flowing" up to time t is given by $A(t)$. In deterministic network calculus such a flow is given and no random decisions are involved. Now we want to generalise this idea, in the sense that chance plays a role. To do so we need the definition of stochastic processes:

Definition 2.1. Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, $(\Omega', \mathcal{A}')$ a measurable space and I an index set. A familiy of random variables $X = (X_i)_{i\in I}$ with $X_i : \Omega \to \Omega'$ is called *stochastic process*, with *state space* Ω' and *time space* I.

Remark 2.2. If I is totally ordered (e.g. when $I \subset \mathbb{R}$) we can fix one $\omega \in \Omega$ and interpret X_i as a function in the "time"-variable i:

$$\begin{aligned} X(\omega) :& I \to \Omega' \\ X(\omega) :& i \mapsto X_i(\omega) \end{aligned}$$

This mapping is called *trajectory* or *(sample) path* of X under ω.[1]

[1] This means, that ω contains "enough" information, to deduce complete trajectories from it. In fact, we often use more than one stochastic process simultaneously, i.e. have several families of random vari-

We give a few examples to illustrate what kind of structures the above definition can cover. However, we are not going into much detail here.

Example 2.3. Markov Chains are stochastic processes with a countable state space and discrete or continuous time space, i.e. $I \in \{\mathbb{N}_0, \mathbb{R}_0^+\}$. In the special case of $\Omega' = \mathbb{N}_0$ and the condition that $X_n \in [X_{n-1} - 1, X_{n-1} + 1]$ our Markov Chain becomes a Birth-Death-Process.

Example 2.4. The most famous continuous time stochastic process is the Brownian Motion or the Wiener Process (see for example §40 in [1]). The one-dimensional Wiener Process W_t has state space $\mathbb{R}$, time space $\mathbb{R}_0^+$ and satisfies:

- $W_0 = 0$
- The trajectories of W_t are continuous, a.s.
- $W_t - W_s$ has normal distribution with parameters $\mu = 0$ and $\sigma^2 = t - s$. $(t \geq s \geq 0)$
- If $t \geq s > u \geq v$, then the increments $W_t - W_s$ and $W_u - W_v$ are stochastically independent. (A similar condition holds for n different increments. See Definition **??**)

Showing the existence of such a process lies out of the scope of this course and is hence omitted.

Example 2.5. The time space is very often something "natural" like $\mathbb{N}_0$ or $\mathbb{R}_0^+$, but one should be aware, that the above definition, covers much more. Other appearing index sets are $[0, 1] \subset \mathbb{R}$ for the Brownian Bridge or $I = T$ for tree-indexed random processes, T being the set of nodes in a finite or infinite tree.

In our context of arrivals, we use a stochastic process with time space $\mathbb{N}_0$, while the state space is kept continuous ($\Omega' = \mathbb{R}_0^+$). Further, we want to preserve the cumulative nature of (2.1).

Definition 2.6. Let $(a(i))_{i \in \mathbb{N}}$ be a sequence of non-negative real random variables. We call the stochastic process A with time space $\mathbb{N}_0$ and state space $\mathbb{R}_0^+$ defined by

$$A(n) := \sum_{i=1}^{n} a(i)$$

a *flow*[2]. The $a(i)$ are called *increments* of the flow.

---

ables $X_i, Y_j, Z_k, \ldots$ each defining its own stochastic process. How this is possible, is best understood if one assumes $I = \mathbb{N}$. In this case, we might represent each event by $\omega = (\omega_{x_1}, \omega_{y_1}, \omega_{z_1}, \omega_{x_2}, \ldots)$ such that each X_i only depends on ω_{x_i}. Of course, the corresponding event space Ω is then quite large. In fact if one sets I equal to an uncountable set, the construction of the original probability space $(\Omega, \mathcal{A}, \mathbb{P})$ is a non-trivial task. In this course, however, we do not delve into this topic and are just happy that someone has done the work of constructing $(\Omega, \mathcal{A}, \mathbb{P})$ for us.

[2] As usual we define the empty sum as zero: $\sum_{i=1}^{0} a(i) = 0$

We see in the following two examples that it is quite hard to bound flows with *deterministic* arrival curves. (A flow has a deterministic arrival curve $\alpha : \mathbb{N}_0 \to \mathbb{R}_0^+$ if $A(n) - A(m) \leq \alpha(n-m)$ for all $n \geq m \geq 0$.)

Example 2.7. First consider the $a(i)$ to be i.i.d. Bernoulli distributed. This means for each $a(i)$ holds:

$$\mathbb{P}(a(i) = 0) = p$$
$$\mathbb{P}(a(i) = 1) = 1 - p$$

for some parameter $p \in (0,1)$. Is it possible that $A(n) - A(m) = n - m$? The answer is yes and we show this, by proving that the corresponding probability is larger than zero:

$$\begin{aligned}
\mathbb{P}(A(n) - A(m) = n - m) &= \mathbb{P}\left(\sum_{k=m+1}^{n} a(k) = n - m\right) \\
&= \mathbb{P}\left(\bigcap_{k=m+1}^{n} \{a(k) = 1\}\right) = \prod_{k=m+1}^{n} \mathbb{P}(a(k) = 1) \\
&= (1-p)^{n-m} > 0
\end{aligned}$$

Hence it is possible to encounter exactly $n - m$ arrivals in an interval of length $n - m$ (although the corresponding probability might be tiny). Hence the best deterministic arrival curve, we can give for such an arrival is $\alpha(n) = n$. With our deterministic "worst-case-glasses" we see A sending 1 in each time step. The possibility that the flow could send nothing in one time step is completely ignored! This buries any hopes to effectively bound a stochastic process by *deterministic* arrival curves. The next example shows, that the situation can become even worse.

Example 2.8. Next assume that the increments $a(i)$ are i.i.d. and exponentially distributed with parameter λ. In this case the best possible deterministic "arrival curve" is given by $\alpha(n) = \infty$ for all $n \in \mathbb{N}$, since we have for any $K \geq 0$:

$$\mathbb{P}(a(i) > K) = e^{-\lambda K} > 0$$

Hence no matter how large we choose K, there is always a small probability, that the arrivals in one time step exceeds K.

The previous two examples are bad news for someone trying to bound stochastic arrivals by deterministic arrival curves. The following corollary summarizes the possibilities one has to deterministically bound i.i.d. arrivals. The proof is generalised easily from example 2.7 and hence omitted.

Corollary 2.9. *Let the increments $a(i)$ of some flow be i.i.d. and define $x_+ := \inf\{x : F(x) = 1\} \in [0, \infty]$, where F is the cdf of $a(i)$. Then the best possible arrival curve for A is given by: $\alpha(n) = n \cdot x_+$.*

We see it is a somewhat pointless enterprise to describe stochastic arrivals by deterministic arrival curves. The solution to this problem is to exclude events, which are unlikely to happen. So instead of asking for some curve α such that $A(n) - A(m) \leq \alpha(n-m)$ we are rather interested in bounds, which hold with a high probability. Stated in a simple form, this looks like

$$\mathbb{P}(A(n) - A(m) \leq \alpha(n-m,\varepsilon)) \geq 1 - \varepsilon \qquad \forall n > m \in \mathbb{N}$$

which is equivalent to:

$$\mathbb{P}(A(n) - A(m) > \alpha(n-m,\varepsilon)) < \varepsilon \qquad \forall n > m \in \mathbb{N} \tag{2.2}$$

Here $\alpha(.,\varepsilon)$ is an arbitrary function, with some parameter $\varepsilon > 0$. [3] Starting from (2.2) two different kinds of bounds can be formulated. The first one is the so-called *tail-bound*, for which we need two more definitions in place:

Definition 2.10. An *envelope function,* or just *envelope,* is a function α with:

$$\alpha : \mathbb{N}_0 \times \mathbb{R}^+ \to \mathbb{R}_0^+$$

Definition 2.11. An *error function*, or just *error,* is a function with:

$$\eta : \mathbb{N}_0 \times \mathbb{R}^+ \to [0,1]$$

and

$$\eta(k,\varepsilon) \leq \eta(l,\varepsilon) \qquad \forall k \geq l,\, \varepsilon > 0.$$

Definition 2.12. A flow A is *tail-bounded* by envelope α with error η, if for all $\varepsilon > 0$ and $n \geq m \in \mathbb{N}_0$ it holds that:

$$\mathbb{P}(A(n) - A(m) > \alpha(n-m,\varepsilon)) \leq \eta(n-m,\varepsilon).$$

Example 2.13. If (2.2) holds for all $\varepsilon > 0$, we have indeed a tail-bound. In practice one often formulates the tail-bound in such a way, that α is of a simple form (e.g. linear). For example the tail-bound

$$\mathbb{P}\left(A(n) - A(m) > r\cdot(n-m) - \frac{1}{\theta}\log\left(\frac{\varepsilon}{M}\right)\right) \leq \varepsilon = \eta(k,\varepsilon)$$

can be reformulated: Choose $\epsilon > 0$ arbitrary and define $\varepsilon := Me^{-\theta\epsilon}$, then it holds that:

$$\mathbb{P}\left(A(n) - A(m) > r\cdot(n-m) + \epsilon\right) \leq Me^{-\theta\epsilon}$$

and since ϵ was choosen arbitrary, the above holds for all $\epsilon > 0$. We can define now $\alpha'(n-m,\epsilon) = r\cdot(n-m) + \epsilon$. This expression describes the probability, that flow A exceeds a maximal rate of r (in the interval $[m,n]$), by at least ϵ.

[3] A common, in literature often found, choice for α is $\alpha(n-m) = r\cdot(n-m) - \frac{1}{\theta}\log(\frac{\varepsilon}{M})$ for some positive θ.

Another kind of bound involving (2.2) is obtained by using Chernoff's inequality.

Theorem 2.14. *Let X be a non-negative random variable and $x > 0$, then holds* Markov's inequality*:*

$$\mathbb{P}(X > x) \leq \frac{\mathbb{E}(X)}{x}$$

Let X be a real random variable and $x \in \mathbb{R}$, then holds for all $\theta > 0$ Chernoff's inequality*:*

$$\mathbb{P}(X > x) \leq e^{-\theta x}\phi_X(\theta)$$

Proof. Let f be the density[4] of X, then it holds for all $x \geq 0$

$$\mathbb{P}(X > x) = \int_x^\infty f(y)dy \leq \int_x^\infty f(y)\frac{y}{x}dy = \frac{1}{x}\int_x^\infty yf(y)dy \leq \frac{\mathbb{E}(X)}{x}$$

which proves Markov's inequality. Now let X be a real random variable, then $e^{\theta X}$ is a non-negative real random variable. Hence, by the monotonicity of the function $h(x) = e^{\theta x}$

$$\mathbb{P}(X > x) = \mathbb{P}(e^{\theta X} > e^{\theta x}) \leq \frac{\mathbb{E}(e^{\theta X})}{e^{\theta x}} = e^{-\theta x}\phi_X(\theta)$$

where we have used Markov's inequality. □

If we apply Chernoff's inequality in (2.2) we get:

$$\mathbb{P}(A(n) - A(m) > \alpha(n-m,\varepsilon)) \leq e^{-\theta\alpha(n-m,\varepsilon)}\phi_{A(n)-A(m)}(\theta)$$

Returning to example (2.13) this reads:

$$\mathbb{P}(A(n) - A(m) > r \cdot (n-m) + \epsilon) \leq e^{-\theta r\cdot(n-m)+\epsilon}\phi_{A(n)-A(m)}(\theta)$$

For further computations the expression $\phi_{A(n)-A(m)}(\theta)$ needs to be taken care of, which triggers the following definitions:

Definition 2.15. A flow A is $(\sigma(\theta), \rho(\theta))$-*bounded* for some $\theta > 0$, if for all $n \geq m \geq 0$ the MGF $\phi_{A(n)-A(m)}(\theta)$ exists and

$$\phi_{A(n)-A(m)}(\theta) \leq e^{\theta\rho(\theta)(n-m)+\theta\sigma(\theta)}$$

holds.

A flow A is $f(\theta,\cdot)$-*bounded* for some $\theta > 0$ if for all $n \geq m \geq 0$ the MGF $\phi_{A(n)-A(m)}(\theta)$ exists and

$$\phi_{A(n)-A(m)}(\theta) \leq f(\theta, n-m)$$

holds.

[4]We give here the proof for distributions with densities. The general proof is almost the same, but needs knowledge about Lebesgue integration.

Obviously $f(\theta,\cdot)$-boundedness is a generalisation of $(\sigma(\theta),\rho(\theta))$-boundedness[5].

Before we study a few examples on tail- and MGF-bounds, we proof a theorem, which shows, that each of them can be converted to the other one, if mild requirements are met. This theorem connects the two branches of stochastic network calculus and we will always keep it in mind, when we derive results in one of the two branches.

Theorem 2.16. *Assume a flow A. If A is tail-bounded by envelope α with error $\eta(k,\varepsilon) = \varepsilon$, it is also $f(\theta,\cdot)$-bounded with:*

$$f(\theta,\cdot) := \int_0^1 e^{\theta\alpha(n-m,\varepsilon)} d\varepsilon$$

Conversely if A is $f(\theta,\cdot)$-bounded it is also tail-bounded with $\alpha(n-m,\varepsilon) = \varepsilon$ and $\eta(n-m,\varepsilon) = f(\theta,n-m)e^{-\theta\varepsilon}$.

Proof. We proof the first part under the assumption that for $A(n)-A(m)$ we have a density function $f_{m,n}$ (remember example **??**). The more general version needs the notion of Lebesgue's integral and - since it works out in the very same fashion - is left out. Denote by $F_{m,n} = \int_0^x f_{m,n}(t)dt$ the cumulative distribution function of $A(n)-A(m)$ and by G the inverse function of $1-F_{m,n}$. We have then for every $\varepsilon \in (0,1]$:

$$\mathbb{P}(A(n)-A(m) > G(\varepsilon)) = 1-\mathbb{P}(A(n)-A(m) \leq G(\varepsilon)) = 1-F(G(\varepsilon)) = \varepsilon$$

On the other hand, from the definition of the tail-bound we have:

$$\mathbb{P}(A(n)-A(m) > \alpha(n-m,\varepsilon)) < \varepsilon = \mathbb{P}(A(n)-A(m) > G(\varepsilon))$$

We can read this as follows: For some value ε the probability of $A(n)-A(m)$ being larger than $G(\varepsilon)$ is larger than the probability of $A(n)-A(m)$ being larger than $\alpha(n-m,\varepsilon)$. This implies:

$$G(\varepsilon) < \alpha(n-m,\varepsilon)$$

Writing the MGF of $A(n)-A(m)$ with the help of $f_{m,n}$ reads:

$$\phi_{A(n)-A(m)}(\theta) = \int_0^\infty e^{\theta x} f_{m,n}(x)dx$$

We substitute the variable x by $G(\varepsilon)$ and multiply by the formal expression $\frac{d\varepsilon}{d\varepsilon}$ (note that $\lim_{a\to 0} 1-F_{m,n}(a) = 0$ and $\lim_{a\to\infty} 1-F_{m,n}(a) = 0$):

$$\begin{aligned}
\phi_{A(n)-A(m)}(\theta) &= \int_1^0 e^{\theta G(\varepsilon)} f_{m,n}(G(\varepsilon)) \frac{dG(\varepsilon)}{d\varepsilon} d\varepsilon \\
&= -\int_0^1 e^{\theta G(\varepsilon)} f_{m,n}(G(\varepsilon)) \frac{1}{-f_{m,n}(G(\varepsilon))} d\varepsilon \\
&= \int_0^1 e^{\theta G(\varepsilon)} d\varepsilon \leq \int_0^1 e^{\theta\alpha(n-m,\varepsilon)} d\varepsilon
\end{aligned}$$

[5] The parameter $\rho(\theta)$ corresponds to the *effective bandwidth* of the flow A. This means if A is $(\sigma(\theta),\rho(\theta))$-bounded we have that $\rho(\theta) \geq \mathbb{E}(a(i))$ and if $a(i)$ is bounded by c, we further have $\rho(\theta) \leq c$ (or can improve to a $(\sigma(\theta),c)$-bound). In words: $\rho(\theta)$ lies between the average rate and the peak rate of the arrivals. (see also Lemma 7.2.3 and 7.2.6 in [3].)

We have used the rule of differentation for inverse functions in the transition from first to second line ($(F^{-1})' = \frac{1}{F'(F^{-1})}$).

For the second part of the theorem we use Chernoff's inequality:

$$\mathbb{P}(A(n) - A(m) < \varepsilon) \leq \phi_{A(n)-A(m)}(\theta)e^{-\theta\varepsilon} \leq f(\theta, n-m)e^{-\theta\varepsilon}$$

□

Next we show how MGF-bounds and tail-bounds can be derived for a given flow.

Example 2.17. Let the increments $a(n)$ of some flow A be i.i.d. exponentially distributed to the parameter λ. This means that our random variables $A(n) - A(m)$ are Erlang distributed with parameters $n - m$ and λ and we know from example **??**

$$\phi_{A(n)-A(m)}(\theta) = \left(\frac{\lambda}{\lambda - \theta}\right)^{n-m}$$

for all $\theta < \lambda$. Hence A is $(\sigma(\theta), \rho(\theta))$-bounded for all $\theta < \lambda$ with $\sigma(\theta) = 0$ and $\rho(\theta) = \frac{1}{\theta}\log\left(\frac{\lambda}{\lambda-\theta}\right)$.

Example 2.18. Let the increments of a flow A be i.i.d. Bernoulli with parameter p. We have then for all $\theta > 0$:

$$\begin{aligned}\phi_{A(n)-A(m)}(\theta) &= (\phi_{a(i)}(\theta))^{n-m} = (\mathbb{E}(e^{\theta a(i)}))^{n-m} \\ &= \left(p \cdot e^{\theta \cdot 0} + (1-p) \cdot e^{\theta \cdot 1}\right)^{n-m} = (p + (1-p)e^{\theta})^{n-m}\end{aligned}$$

Hence A is $f(\theta, n)$ bounded with $f(\theta, n) = (p + (1-p)e^{\theta})^n$ for all $\theta > 0$.

Example 2.19. Let the increments of a flow A be i.i.d. Pareto distributed with parameters x_m and α. We know already that the MGF of the Pareto distribution does not exist for any $\theta > 0$, hence no MGF-bound can be found for this flow.

Example 2.20. Assume the increments of a flow to be i.i.d. with bounded variance $\mathrm{Var}(a(n)) \leq \sigma^2 < \infty$. Note that $\mathbb{E}(A(n) - A(m)) = (n-m)\mathbb{E}(a(1))$ and hence $\mathbb{E}(a(1)) = \frac{\mathbb{E}(A(n)-A(m))}{n-m}$. Using the Chebyshev inequality[6] a tail-bound is constructed by:

$$\mathbb{P}(A(n) - A(m) > (n-m)(\mathbb{E}(a(n)) + \varepsilon)) \leq \mathbb{P}\left(\left|\frac{A(n) - A(m)}{n-m} - \mathbb{E}(a(n))\right| > \varepsilon\right) \leq \frac{1}{\varepsilon^2(n-m)}\sigma^2$$

In the following example, we see how a multiplexed flow can be bounded, if we already have bounds for the single flows.

[6] Chebyshev's inequality states that for a random variable X with finite variance it holds that: $\mathbb{P}(X - \mathbb{E}(X) > \varepsilon) \leq \varepsilon^{-2}\mathrm{Var}(X)$.

Example 2.21. Let A and B be two flows, which are $(\sigma_A(\theta), \rho_A(\theta))$- and $(\sigma_B(\theta), \rho_B(\theta))$-bounded, respectively (for the same $\theta > 0$). Further let A and B be stochastically independent. We call $A \oplus B(n) := A(n) + B(n)$ the multiplexed flow[7]. We have then:

$$\begin{aligned}\phi_{A\oplus B(n)-A\oplus B(m)}(\theta) &= \mathbb{E}(e^{\theta(A(n)-A(m)+B(n)-B(m))}) = \mathbb{E}(e^{\theta(A(n)-A(m))})\mathbb{E}(e^{\theta(B(n)-B(m))}) \\ &\leq e^{\theta\rho_A(\theta)(n-m)+\theta\sigma_A(\theta)}e^{\theta\rho_B(\theta)(n-m)+\theta\sigma_B(\theta)} \\ &= e^{\theta(\rho_A(\theta)+\rho_B(\theta))(n-m)+\theta(\sigma_A(\theta)+\sigma_B(\theta))}\end{aligned}$$

Hence $A \oplus B$ is bounded by $(\sigma_A(\theta) + \sigma_B(\theta), \rho_A(\theta) + \rho_B(\theta))$.

The last example raises the question what to do, when the flows A and B are not independent. The key for this problem is to take care of expressions of the form $\mathbb{E}(XY)$. The next lemma gives us a way to deal with that.

Lemma 2.22. *Let X and Y be two non-negative real random variables with finite second moment (i.e. $\mathbb{E}(X^2), \mathbb{E}(Y^2) < \infty$). Then holds the* Cauchy-Schwartz inequality*:*

$$\mathbb{E}(XY) = (\mathbb{E}(X^2))^{1/2}(\mathbb{E}(Y^2))^{1/2}$$

Let $p, q \in \mathbb{R}^+$ such that $\frac{1}{p} + \frac{1}{q} = 1$ and assume the p-th moment of X and the q-th moment of Y to be finite. Then holds Hölder's inequality*:*

$$\mathbb{E}(XY) = (\mathbb{E}(X^p))^{1/p}(\mathbb{E}(Y^q))^{1/q}$$

For the proof we need Lebesgue-integration, which would lead us too far away from our course. Hence a proof is omitted (see for example [3]). Let us resume now to the situation of example 2.21 with the difference, that the two flows are not stochastically independent. Using the above lemma we can achieve that $A \oplus B$ is $(\sigma_A(p\theta) + \sigma_B(q\theta), \rho_A(p\theta) + \rho_B(q\theta))$-bounded. But note that we need the conditional assumptions of A being $(\sigma_A(p\theta), \rho_A(p\theta))$-bounded and B being $(\sigma_B(q\theta), \rho_B(q\theta))$-bounded.

Exercises

Exercise 2.23. Prove corollary 2.9.

Exercise 2.24. Let A be some flow, which has a deterministic arrival curve α. What tail-bound and MGF-bound can be given? Let now A be some flow, which has a token bucket arrival curve $\alpha(n) = r \cdot n + B$, show that A is $(\sigma(\theta), \rho(\theta))$-bounded and determine $\sigma(\theta)$ and $\rho(\theta)$.

Exercise 2.25. Assume a flow A has the following properties: All increments $a(n)$ with $n \geq 2$ are i.i.d. uniformly distributed on the interval $[0, b]$ (this means, it has the density $f(x) = \frac{1}{b}$ for all $x \in [0, b]$ and $f(x) = 0$ elsewhere). The first increment $a(1)$ however is exponentially distributed with parameter λ. Give a $(\sigma(\theta), \rho(\theta))$-bound for A.

(Hint: Calculate first the MGF of $A(n) - A(m)$ for $m \neq 0$ and then for $m = 0$. Find then a bound which covers both cases)

[7] You can convince yourself easily that $A \otimes B$ is a flow.

Exercise 2.26. Assume a flow A is $(\sigma(\theta), \rho(\theta))$-bounded for all $\theta \in [0, b)$ with $\sigma(\theta) = \sigma$ and $\rho(\theta) = \rho$ for all θ and some $b > 0$. Show that the expected arrivals in some time interval $[m+1, n]$ is upper bounded by $\rho \cdot (n-m) + \sigma$, i.e.:

$$\mathbb{E}(A(n) - A(m)) \leq \rho \cdot (n-m) + \sigma \qquad \forall\, n \geq m \geq 0$$

(Hint: Show first that $\phi_{A(n)-A(m)}(0) = \lim_{\theta \searrow 0} e^{\theta\rho(\theta)(n-m)+\theta\sigma(\theta)} = 1$, then conclude from the $(\sigma(\theta), \rho(\theta))$-boundedness that the first derivative of $\phi_{A(n)-A(m)}$ at point 0 must be smaller or equal to $\frac{d}{d\theta} e^{\theta\rho(\theta)(n-m)+\theta\sigma(\theta)}\big|_{\theta=0}$ (why?). Then use exercise **??**)

Exercise 2.27. Assume a flow A is $(\sigma(\theta), \rho(\theta))$-bounded for all $\theta \in [0, b)$ and some $b > 0$. Further assume the following conditions on the bound:

$$\theta\rho(\theta) \xrightarrow{\theta\to\infty} 0$$
$$\theta\sigma(\theta) \xrightarrow{\theta\to\infty} 0$$
$$\theta\frac{d}{d\theta}\rho(\theta) \xrightarrow{\theta\to\infty} 0$$
$$\theta\frac{d}{d\theta}\sigma(\theta) \xrightarrow{\theta\to\infty} 0$$

Show that $\mathbb{E}(A(n) - A(m)) \leq \rho(0)(n-m) + \sigma(0)$ (Hint: Proceed as in the previous exercise).

Prove that the conditions are met if the increments are i.i.d. exponentially distributed with parameter λ. Apply this knowledge to recover $\mathbb{E}(A(n) - A(m)) \leq \left(\frac{1}{\lambda}\right)^{n-m}$.

Exercise 2.28. Let there be $J \in \mathbb{N}$ stochastically independent flows A_j $(j = 1, \ldots, J)$, all of them having i.i.d. Bernoulli distributed increments $a_j(i)$ with the same parameter p. One can easily see, that the number of flows sending a paket (i.e. $a_j(n) = 1$) at a certain time n follows a binomial distribution, with parameters p and J:

$$\mathbb{P}\left(\sum_{j=1}^{J} a_j(n) = k\right) = \binom{J}{k} p^k (1-p)^{J-k} \qquad \forall\, n \in \mathbb{N}$$

Hence we can think of the multiplexed flow $A(n) := \sum_{j=1}^{J} a_j(n)$ to be binomially distributed. Give a $(\sigma(\theta), \rho(\theta))$-bound for the multiplexed flow using that it is binomially distributed.

Now use instead the multiplexing property of $(\sigma(\theta), \rho(\theta))$-bounds, as seen in example 2.21.

Exercise 2.29. Next assume two flows A and B. We have that the increments of the first flow $a(i)$ are Bernoulli distributed with parameter p, further we assume that for the increments of B holds $b(i) = 1 - a(i)$. Convince yourself, that the increments of B are again Bernoulli distributed, but with the parameter $1-p$. Of what form is the flow $A \oplus B$? Give a $(\sigma(\theta), \rho(\theta))$-bound for $A \oplus B$ (Hint: Use exercise 2.24)! Next use the multiplexing property of $(\sigma(\theta), \rho(\theta))$-bounds (with Hölder) to bound the flow $A \oplus B$. Which bound is better? (Do not try to solve this analytically!)

Exercise 2.30. Remember the Cauchy-Schwarz inequality, as you know it from inner product spaces. An inner product space consists of a vector space V (usually $\mathbb{R}^d$ or $\mathbb{C}^d$) a field of scalars F (usually $\mathbb{R}$ or $\mathbb{C}$) and an inner product (dt.: Skalarprodukt)

$$\langle \cdot, \cdot \rangle : V \times V \to F$$

fulfilling:

- (Conjugate) symmetry: $\langle x, y \rangle = \overline{\langle x, y \rangle}$
- Linearity in the first argument: $\langle a \cdot x, y \rangle = a \cdot \langle x, y \rangle$ and $\langle x + y, z \rangle = \langle x, z \rangle + \langle y, z \rangle$
- Positive-definiteness: $\langle x, x \rangle \geq 0$ with equality only for $x = 0$.

Check that $\langle X, Y \rangle = \mathbb{E}(X \cdot Y)$ defines an inner product "on the space of one dimensional real random variables"

For inner product spaces one can define a norm by:

$$\|x\| := \sqrt{\langle x, x \rangle}$$

Convince yourself that the Cauchy-Schwarz inequality as given in lemma 2.22 is in fact the Cauchy-Schwarz inequality for inner product spaces:

$$|\langle x, y \rangle|^2 \leq \|x\| \cdot \|y\|$$

(with equality if x and y are linearly(!) independent)

Exercise 2.31. Prove the bound for multiplexed flows, being stochastically dependent (as given after lemma 2.22)!

Exercise 2.32. We now investigate the parameter θ in the $(\sigma(\theta), \rho(\theta))$-bounds, which we will call *acuity* later on. Generate some (in the order of thousands) exponentially distributed random variables with parameter $\lambda = 10$ (you might reuse your code from exercise **??**). Sort your realisations by magnitude and plot them. Now apply the function $f(x) = e^{\theta x}$ to your realisations with varying parameter $\theta \in [0, 10)$. How does the plot change for different values of θ? You should be able to see, that for large θ one has a few very large results, compared to the majority of smaller results. For smaller θ, however, one gets a more balanced picture.

We ask now how many of the realisations are larger than the expected value of one of the realisations $\mathbb{E}(e^{\theta X})$. For this we already know from example **??** that $\mathbb{E}(e^{\theta X}) = \frac{\lambda}{\lambda - \theta}$. Write a procedure, which counts for a fixed θ the number of realisations being larger $\frac{\lambda}{\lambda - \theta}$. What percentage of your realisations are larger for a small θ? What is it for a large θ? Explain the difference with the help of your previously produced plots!

Bibliography

[1] *H. Bauer.* Probability Theory. *De Gruyter, 1995.*

[2] *R.R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms.* Selected Areas in Communications, IEEE Journal on*, 18(12):2651–2664, 2000.*

[3] *C.-S. Chang.* Performance Guarantees in Communication Networks. *Springer, 2000.*

[4] *M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In* Quality of Service, 2006. IWQoS 2006. 14th IEEE International Workshop on*, pages 261 –270, 2006.*

[5] *Y. Jiang and Y. Liu.* Stochastic Network Calculus. *Springer, 2008.*

[6] *J. B. Schmitt. Lecture notes: Performance modeling of distributed systems. http://disco.informatik.uni-kl.de/content/PDS_ WS1112, October 2011.*