A Unified Analysis of Markov Additive Processes

Florin Ciucu
University of Warwick

Joint work with Felix Poloczek (Warwick + TU Berlin)

The Problem. System Performance Analysis

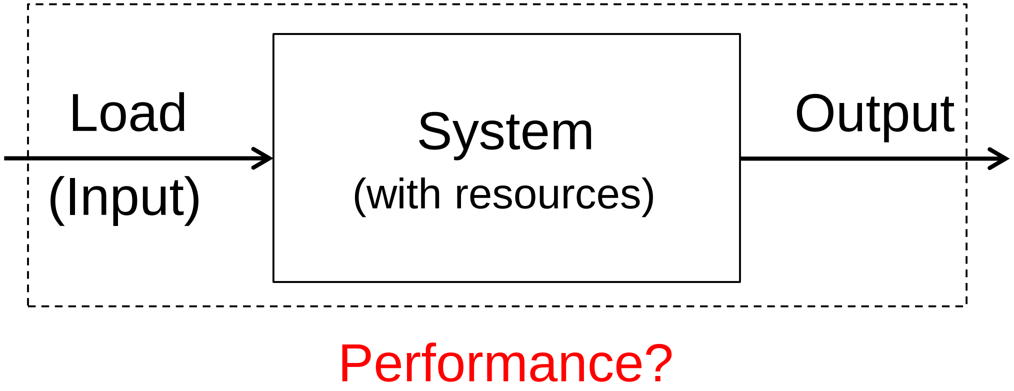

- Examples
  - The system: a network, a data center, the power grid, a cache
  - The resources: bandwidth, processors, batteries
  - The load: bits, jobs, energy demand/supply
  - The performance: reliable transmission, completion time, matching

- Problem formulations
  - Load + resources → performance
  - Load + performance → resources

Formalizing "the System": A Queueing Model

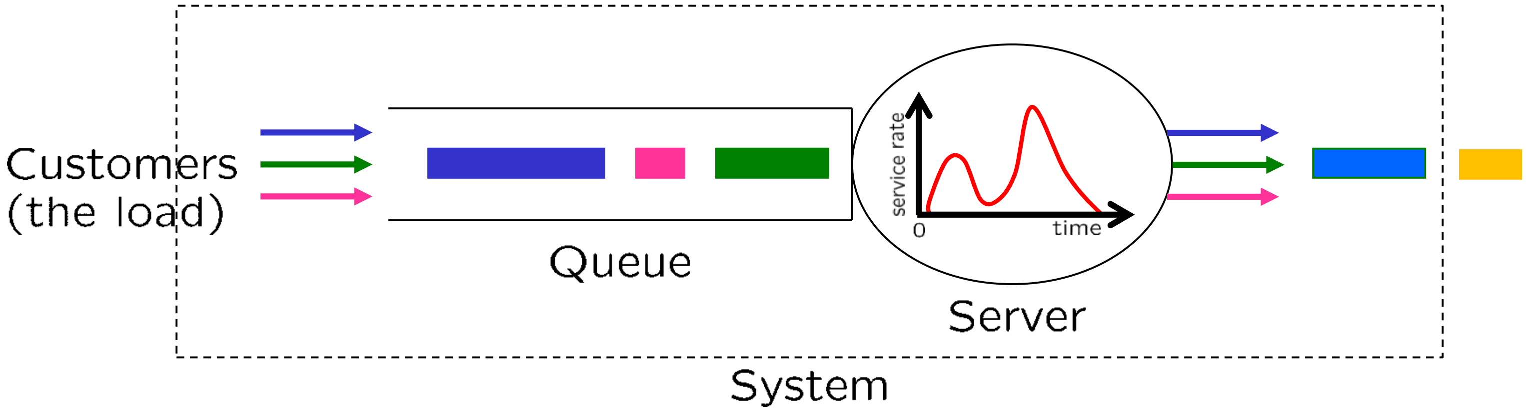

- Input
  - statistical descriptions on the load and server, e.g., How do customers arrive? How quickly are they served?
  - other factors, e.g., queue size, scheduling
- Output

$$\mathbb{P}_{\text{load, s\_rate}}\left(\text{delay} > x\right) = ?$$

$$\text{Battery Size s.t. } \mathbb{P}\left(\text{Load} \approx \text{Demand}\right) = 1 - \varepsilon$$

The Resource Allocation Problem. The Reality

- Constraints
  - Traffic is "complex" (i.e., non-Poisson, subject to short(long)-term correlations
  - Networks are complex
  - The underlying network/transport protocols add more complexity

- The goal: controllable delay/latency tails
  - Latency tail-tolerant systems
  - Internet at the speed of light, Tactile Internet

- Can it be done?
  - Yes: by overprovisioning, measurements, more overprovisioning, etc.
  - In some better way?

The Invention of Q. T. (A. K. Erlang, 1910's)

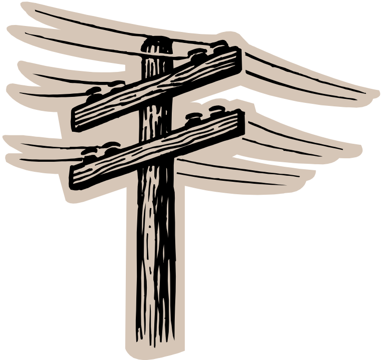

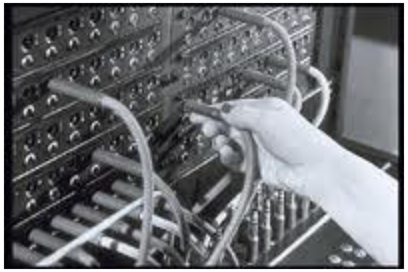

Remote Village

(Customers)

Telephone Lines

(Server)

Regional Office

Problem: given the number of phones and a target probability for getting a busy tone, determine the number of required telephone lines.

Q. T. for the Internet. The Rise (60's)

- Packet switching technology: all flows share the available bandwidth by interleaving packets

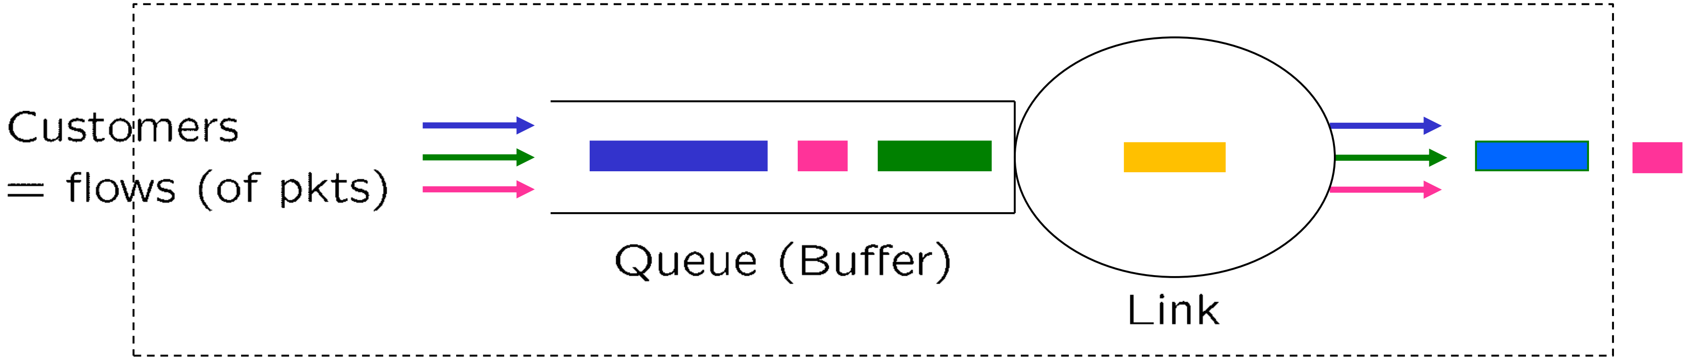

- Raison d'être: statistical multiplexing gain[1]

$$\begin{pmatrix}\text{Bandwidth needed to support} \\ \text{service for N flows}\end{pmatrix} << N \times \begin{pmatrix}\text{Bandwidth needed to support} \\ \text{service for 1 flow}\end{pmatrix}$$

[1]Liebeherr et al., 2001

Modeling Internet Traffic (60's)

- Alike the Telephone Network traffic
  - Packet arrivals: Poisson process
  - Packet sizes: exponential
- But … packets must change their size (?!) downstream

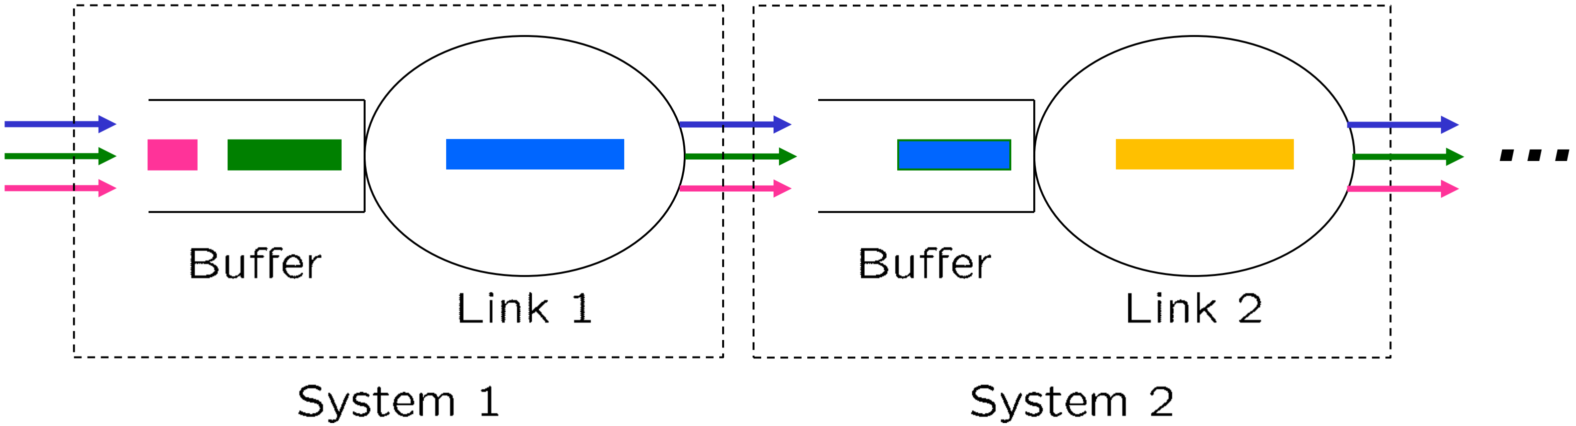

- This convenient assumption was numerically justified, but … it leads to incorrect scaling laws of, e.g., e2e delays[1]

$$\Theta(n) \text{ vs. } \Theta(n \log n)$$

[1]*Burchard/Liebeherr/Ciucu, ToN 2011*

Audio/Video Internet Traffic (80's)

- New models
  - Markov Fluid (MF)

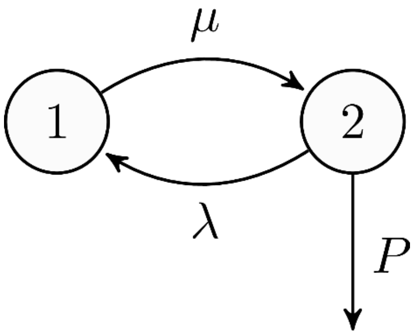

  - Markov Modulated Poisson (MMP)

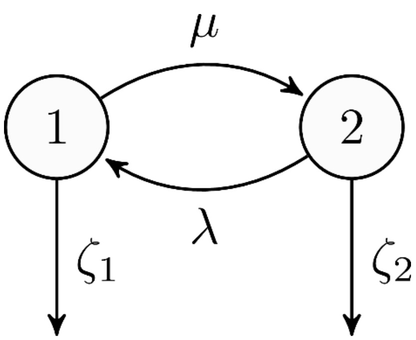

- And techniques (spectral decompositions, Wiener-Hopf factorization)
- Exact results but numerical complexity blows up

Bellcore Ethernet Traces (90's)

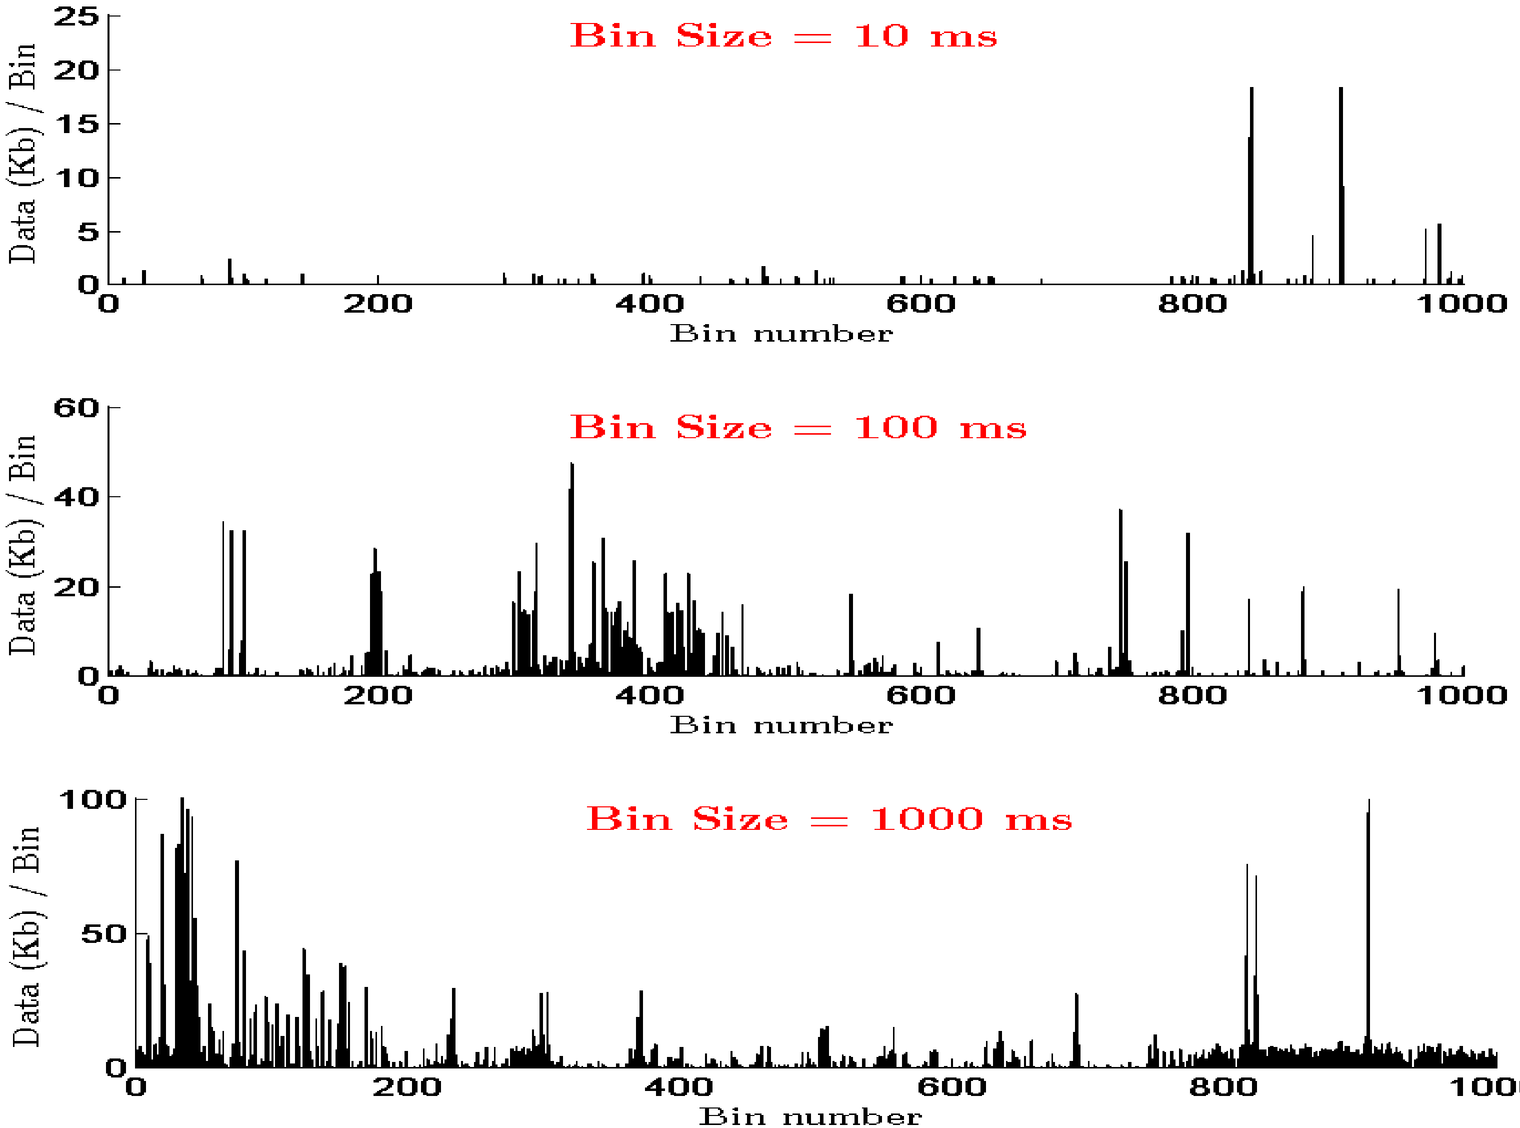

Q. T. for the Internet. The Decline

- A.k.a. the failure of Poisson modeling
- Applying classical results to modern Internet traffic can be very misleading
- New models (heavy-tailed, self-similar, alpha-stable) and techniques
  - capture *the exact* scaling behavior, e.g.,

$$\mathbb{P}_{\text{load, s\_rate}}\left(\text{delay} > x\right) = x^{1-\alpha},\ \ x >> 0$$

  - but inaccurate in finite regimes
  - … few scheduling, and overly-sophisticated (mathematically)
  - … the network case (?)

Alternatives. Effective Bandwidth (late 1980s-90s)

- Setting: multiplexing many flows (MF, MMP, FBM)

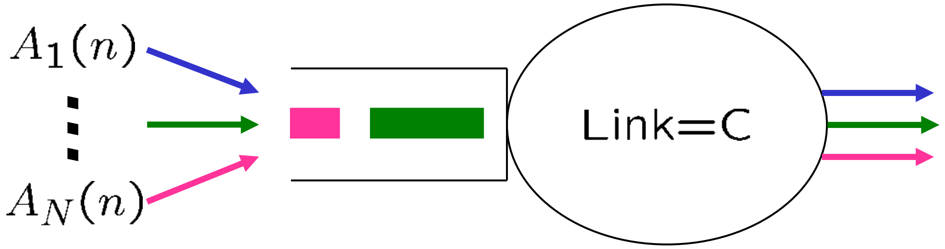

- Lindley's equation

$$Q = \sup_{t \geq 0} \{ \sum_i A_i(t) - Ct \}$$

- The tail (?)

$$\mathbb{P}\left(\sup_{t \geq 0} \{ \sum_i A_i(t) - Ct \} > x \right)$$

Computations and Main Result

- Union Bound

$$\mathbb{P}\left(\sup_{t\geq 0}\{\textstyle\sum_i A_i(t) - Ct\} > x\right)$$
$$\leq \textstyle\sum_t \mathbb{P}\left(\sum_i A_i(t) - Ct > x\right)$$

- Large deviations

$$\mathbb{P}\left(X_1 + X_2 + \cdots + X_N > NE[X] + z\right) \leq \ldots$$

- Effective bandwidth approximation (for loss)

$$\mathbb{P}\left(Q > x\right) \approx e^{-\eta x}$$

Numerical Accuracy? An Admission Control Problem

- Given a total capacity C how many flows should be admitted?

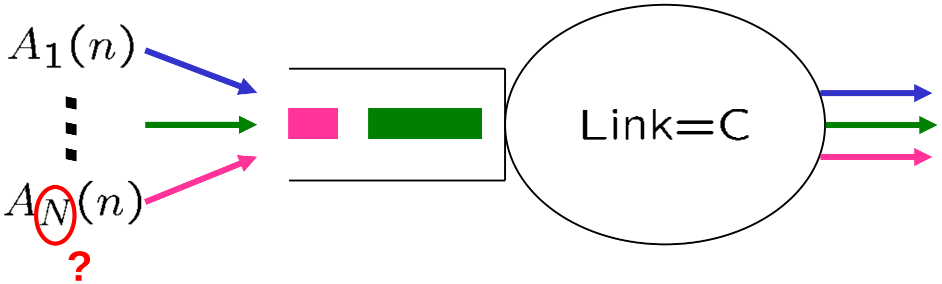

- … subject to the constraint

$$\mathbb{P}^i_{\text{loss}} \leq \varepsilon$$

- Answer: as many as long as

$$\sum_i C_i \leq C$$

- It (the effective bandwidth approximation) only “makes sense” for Poisson flows; too conservative otherwise

N. Shroff and M. Schwartz, Improved Loss Calculations at an ATM multiplexer, IEEE Transactions on Networking, 1998

The "Killing" Evidence

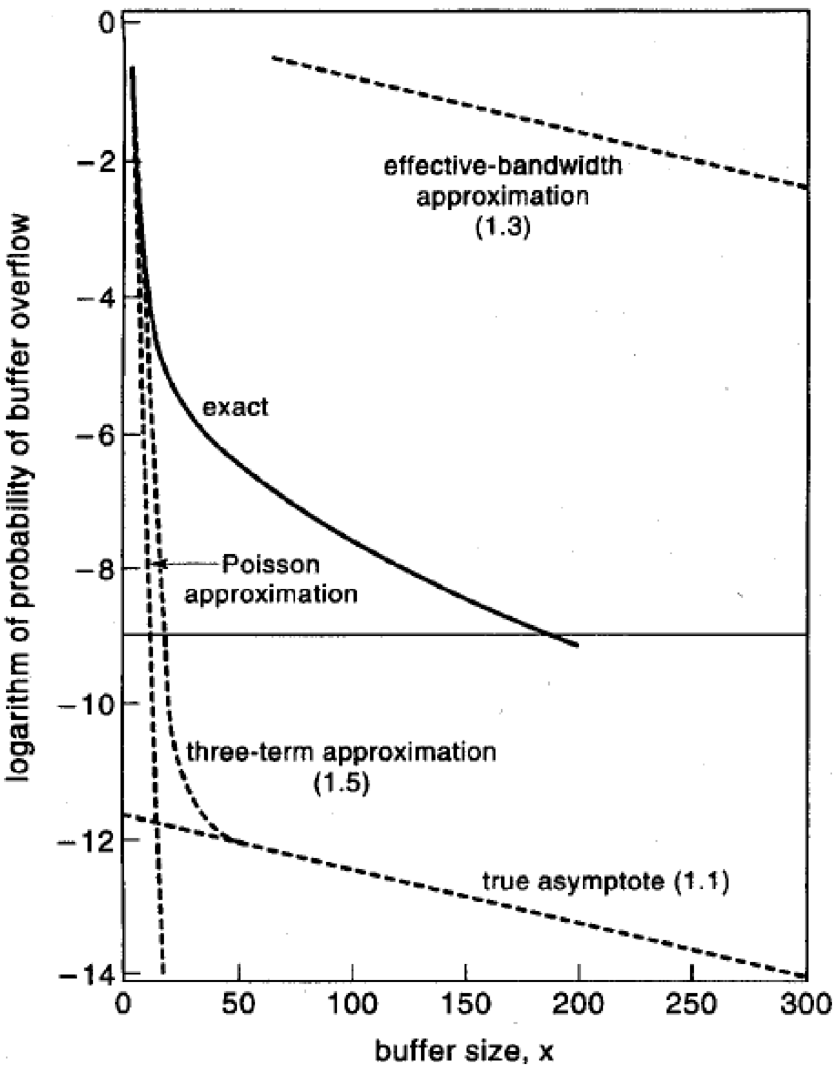

60 MMPP flows

$$(1.3)\ \mathbb{P}(Q > x) \approx e^{-\eta x}$$

$$(1.1)\ \approx \alpha e^{-\eta x}$$

$$(1.5)\ \approx \alpha_1 e^{-\eta_1 x} + \alpha_2 e^{-\eta_2 x} + \alpha_3 e^{-\eta_3 x}$$

$$(!)\ \mathbb{P}(Q > x) \approx \beta e^{-N\gamma} e^{-\eta x}$$

G. Choudhury, D. Lucantoni, and W. Whitt, Squeezing the Most out of ATM, IEEE Transactions on Communications, 1996

Another Alternative: Stochastic Network Calculus (1990s- ?)

- Extends the (deterministic) network calculus methodology in a probability space
- Essentially, it's a mix of

Effective Bandwidth $\mathbb{P}\left(Q > x\right) \approx e^{-\eta x}$

+

Scheduling

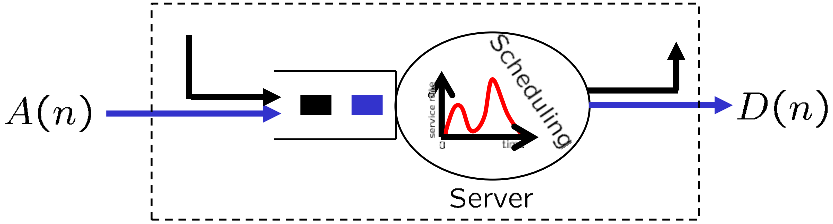

+

Multi-Node

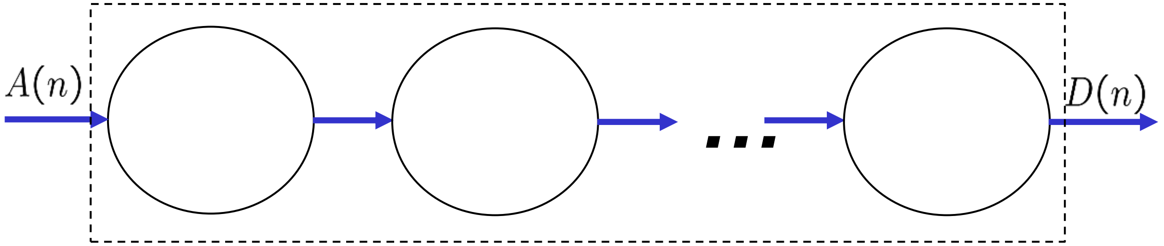

Key Property 1: Scheduling Abstraction

- Consider the following real system (from the perspective of $A(n)$) which is not linear

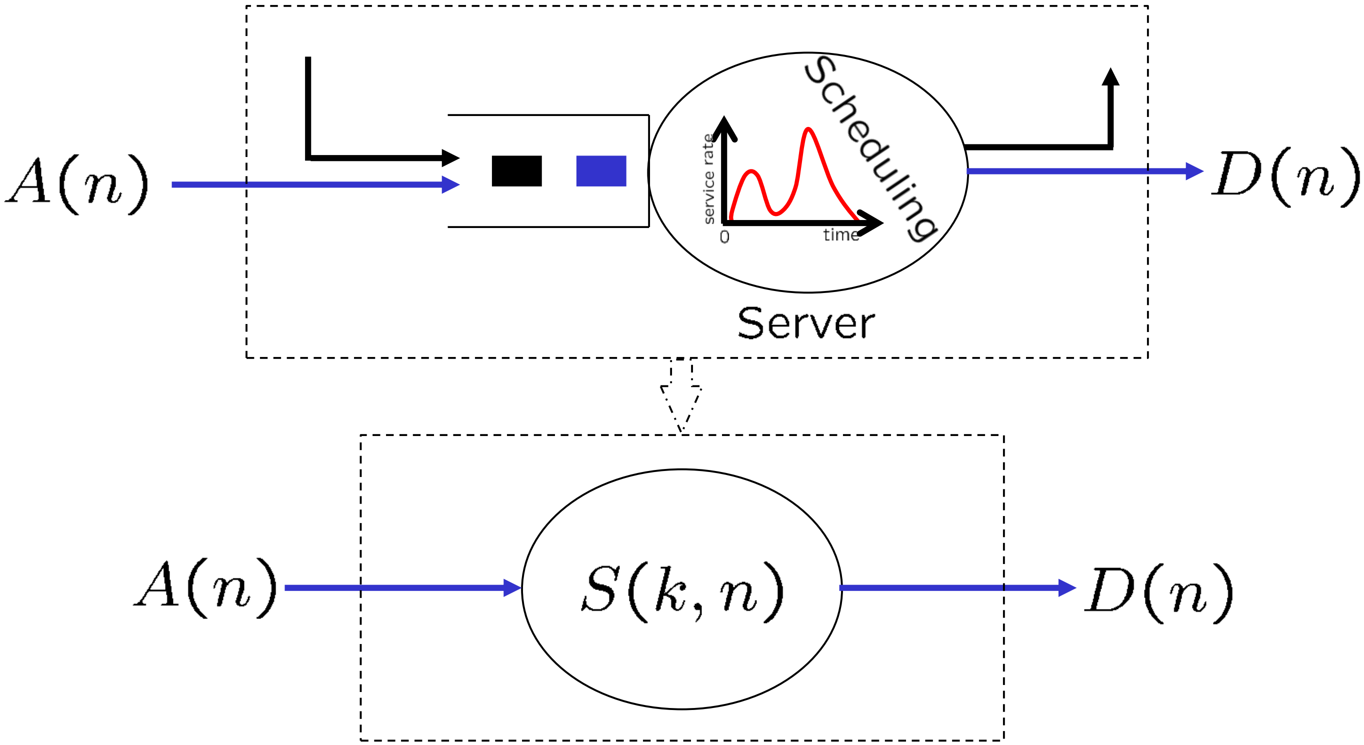

- The transformed virtual system is `somewhat' linear

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

Key Property 2: Convolution-Form Networks

- Consider a concatenation of systems with known service processes

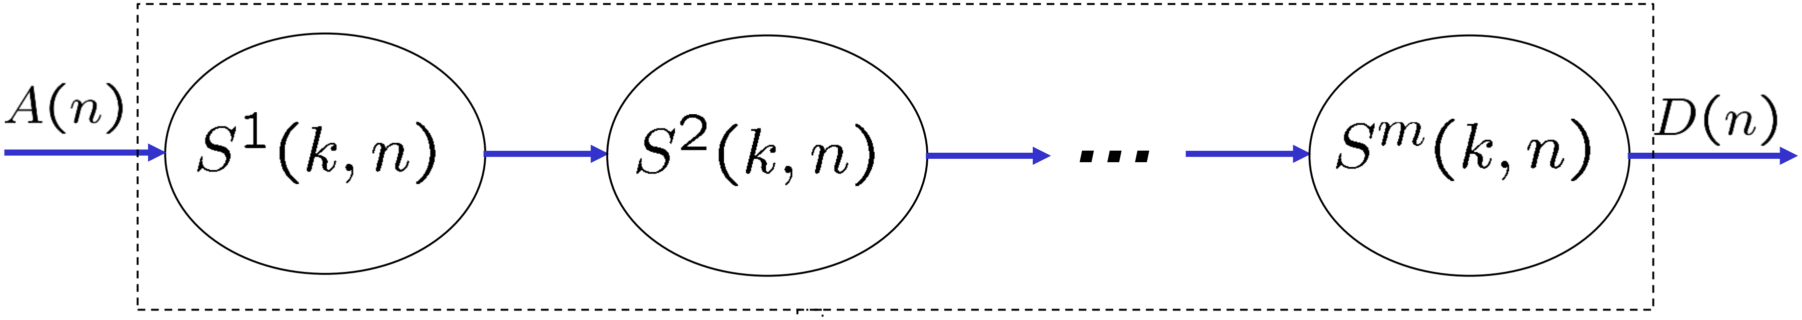

- NC transforms it to a single system

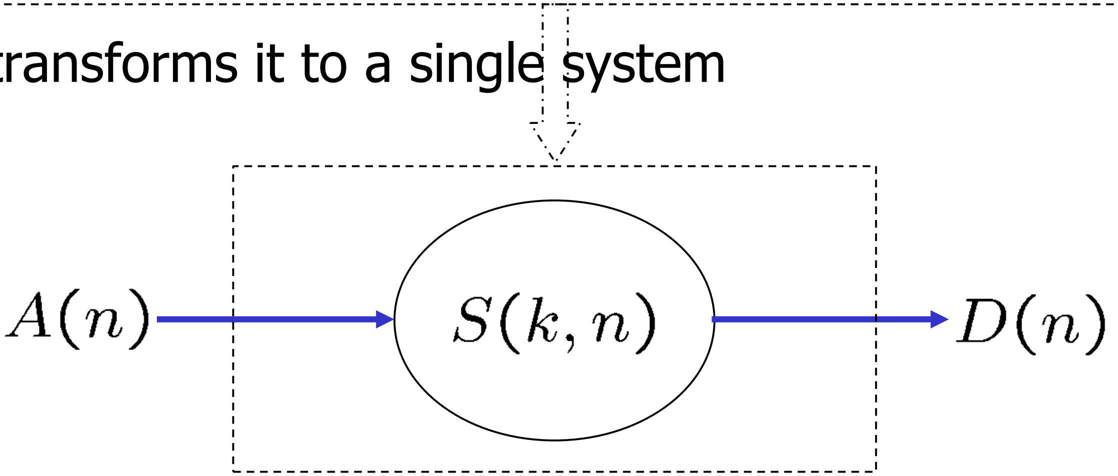

… where $S(k,n)$ is the (min,+) convolution of the others

Why Do Some Work on It?

- An admission control problem: How many type 1/2 flows can be admitted at some link subject to some latency constraints?

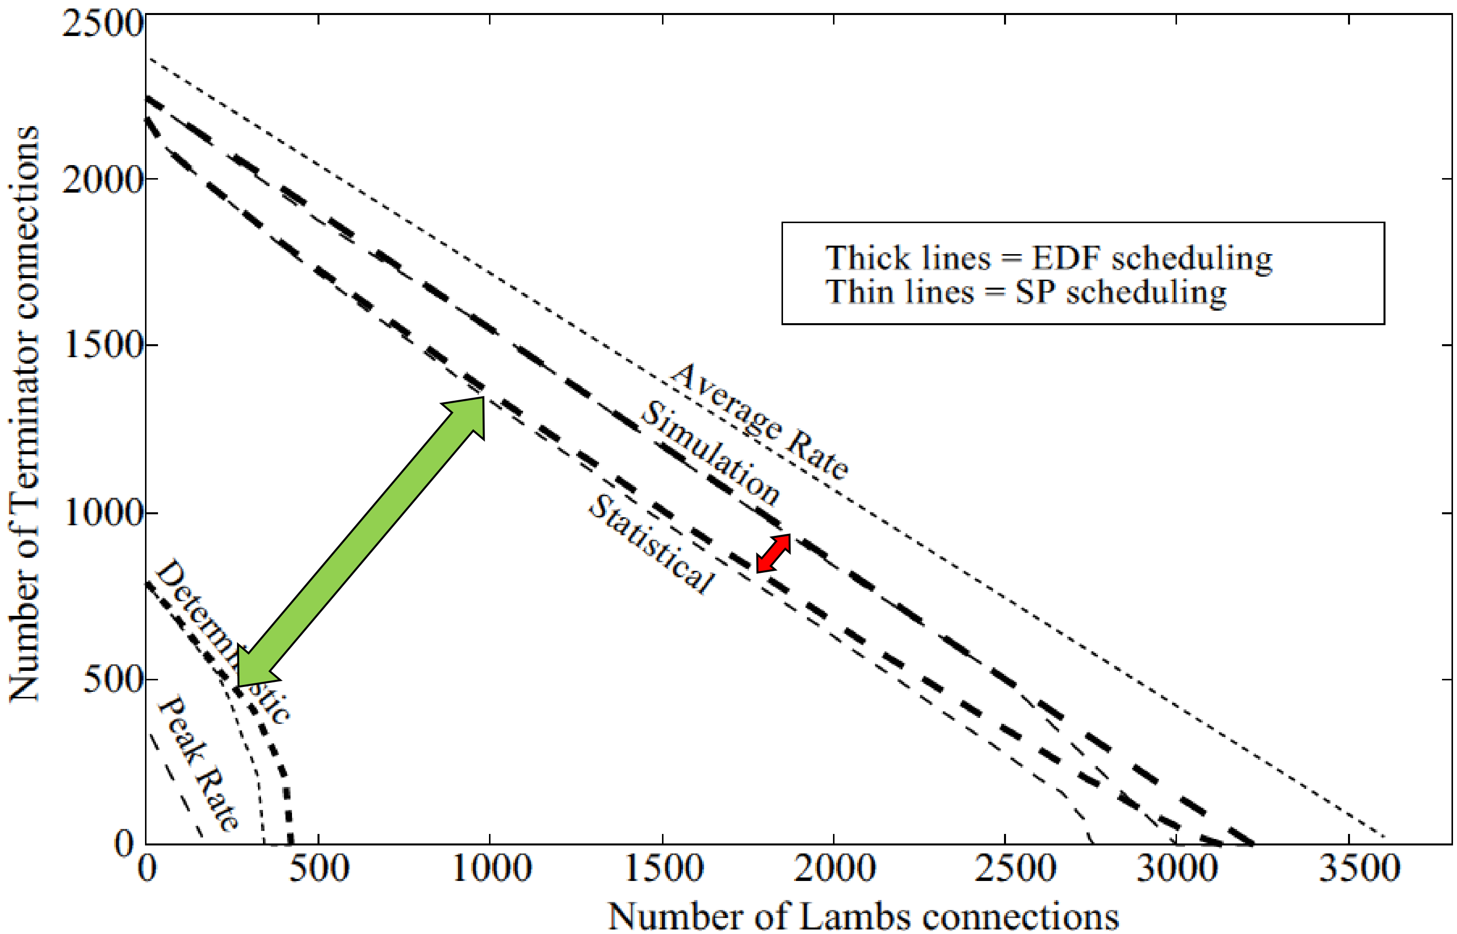

J. Liebeherr, A Note on Statistical Multiplexing and Scheduling in Video Networks at High Data Rates, University of Virginia, 2002

… but on a Closer Look

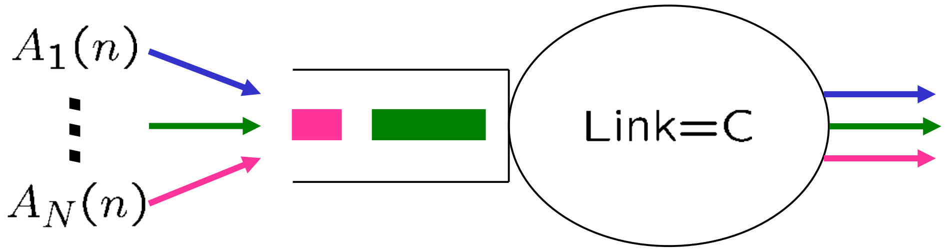

each flow:

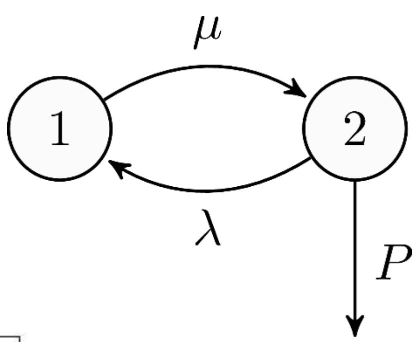

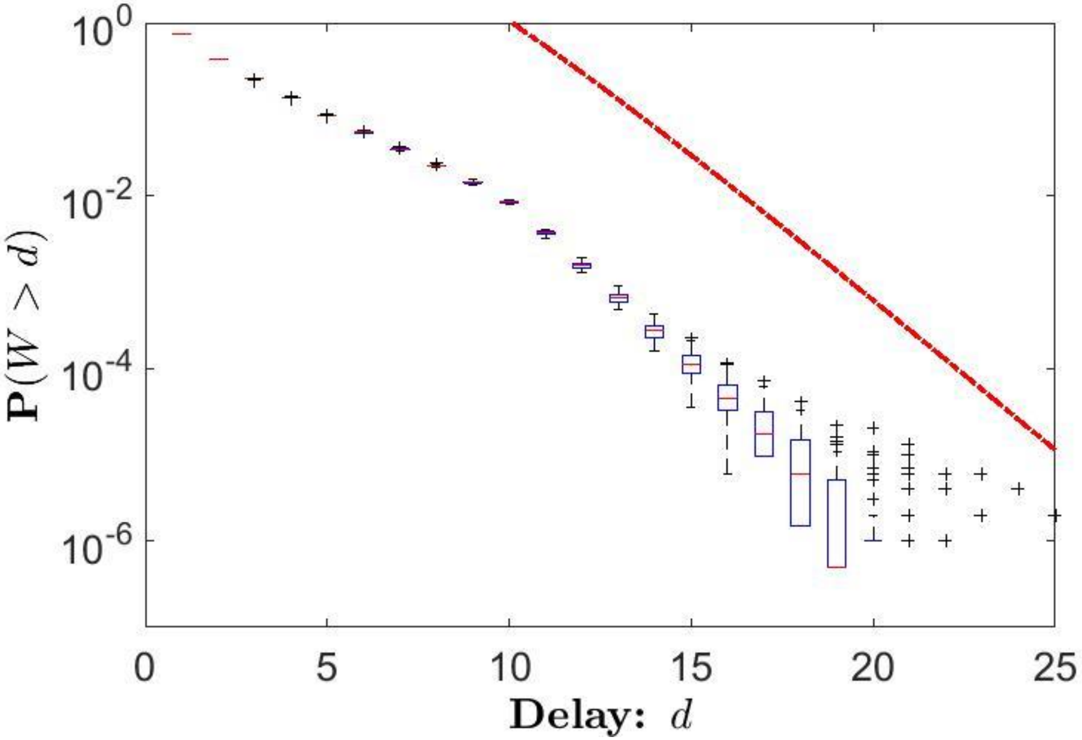

Moving Forward: How to Account for Correlations?

- … and avoid the "killing" step

$$\mathbb{P}\left(\sup_{t\geq 0}\{\textstyle\sum_i A_i(t) - Ct\} > x\right)$$
$$\leq \textstyle\sum_t \mathbb{P}\left(\sum_i A_i(t) - Ct > x\right)$$

- Insight #1: in queueing systems

  $E\,[\text{buffer change} \mid \text{history}] \leq 0$

  due to: average rate $\leq$ capacity (Loynes' condition)

- A *supermartingale* is a process X_n such that for each $n \in \mathbb{N}$

  $$E\left[X_{n+1} - X_n \mid X_1, \ldots, X_n\right] \leq 0$$

  (the expected increment is negative)

Dealing with the Actual Queueing Problem

$$\mathbb{P}\left(\sup_{t\geq 0}\{\sum_i A_i(t) - Ct\} > x\right)$$

- Defining the stopping time

$$\tau := \inf\left\{t : \sum_i A_i(t) - Ct > x\right\}$$

- … we need to compute

$$\mathbb{P}\big(\tau < \infty\big)$$

- Insight #2: (bounded) stopping times preserve martingale properties

$$E\left[X_1\right] = E\left[X_\tau\right]$$

Martingale-Envelope

- Idea: assign to a queueing system a suitable supermartingale M_n ("Martingale-Envelope")

- **Definition**:

  For $\theta > 0$ and h monotonically increasing, the flow A admits a (h, θ, C)-martingale-envelope if for $n \geq m$

  $$M_n := h(a_n)e^{\theta(A(m,n)-C(n-m))}$$

  is a (super-)martingale

  - C is the allocated capacity
  - θ and h capture the correlation structure of A

Operation 1: Multiplexing

- If two independent flows A_1 and A_2 admit martingale-envelopes (h_1, θ, C_1) and (h_2, θ, C_2) then the aggregate $A_1 + A_2$ admits the martingale-envelope

  $$(h_1 \otimes h_2, \theta, C_1 + C_2)$$

  where

  $$h_1 \otimes h_2(n) := \min_{0 \le m \le n} h_1(m) h_2(n-m) \quad ((\text{min}, \times) - \text{convolution})$$

- Why? (independent martingales are “closed” under the operation of multiplication)

Operation 2. Scheduling

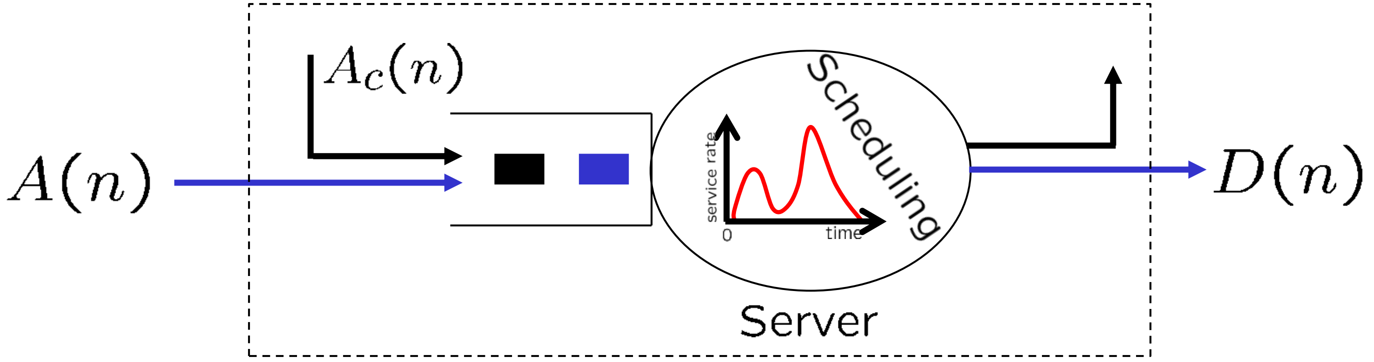

- If $M(n)$ captures $A(n)$ and $M_c(n)$ captures $A_c(n)$ then for a switching time k then

$$\tilde{M}(n) = \begin{cases} M_c(n) & n \leq k \\ M_c(n)M(n) & n \geq k \end{cases}$$

is a martingale.

- Construct a service martingale

$$M_{s,n} := h_s(s_n)e^{\theta(C(n-m)-S(m,n))}$$

Example 1: Markov On-On Processes

- two state Markov chain a_k
- stationary distribution $\pi = \left(\frac{q}{p+q}, \frac{p}{p+q}\right)$
- arrival process $A(n) = \sum a_k$
- Transform the transition matrix

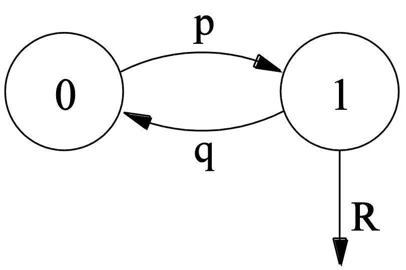

$$T := \begin{pmatrix} 1-p & p \\ q & 1-q \end{pmatrix} \rightsquigarrow T_\theta := \begin{pmatrix} 1-p & pe^{\theta R} \\ q & (1-q)e^{\theta R} \end{pmatrix}$$

- Let $\lambda(\theta)$ spectral radius and (v_0, v_1) eigenvector

- If C and θ satisfy $\lambda(\theta) = e^{\theta C}$ then

  $M_n := v_{a_n} e^{\theta(A(n)-Cn)}$ is a martingale.

  → A admits a (h, θ, C)-martingale-envelope (!), where

$h(0) := v_0$ and $h(R) := v_1$

Example 2: Autoregressive processes

- $Z_1, Z_2, \ldots \sim \mathcal{N}_{0,1}$" Gaussian White Noise" $\mu, \sigma > 0,\ 0 < \varphi < 1$
- Autoregressive process:

$$a_{n+1} := \varphi a_n + (1-\varphi)\mu + (1-\varphi)\sigma Z_n$$

- Let $\theta = 2\dfrac{C-\mu}{\sigma^2}$ and $h(t) := e^{\frac{\varphi}{1-\varphi}\theta t}$
- Then $M(n) := h(a_n)e^{\theta(A(n)-nC)}$ is a martingale
- → A admits a (h, θ, C)-martingale-envelope

Application #1. Per-Flow Delay

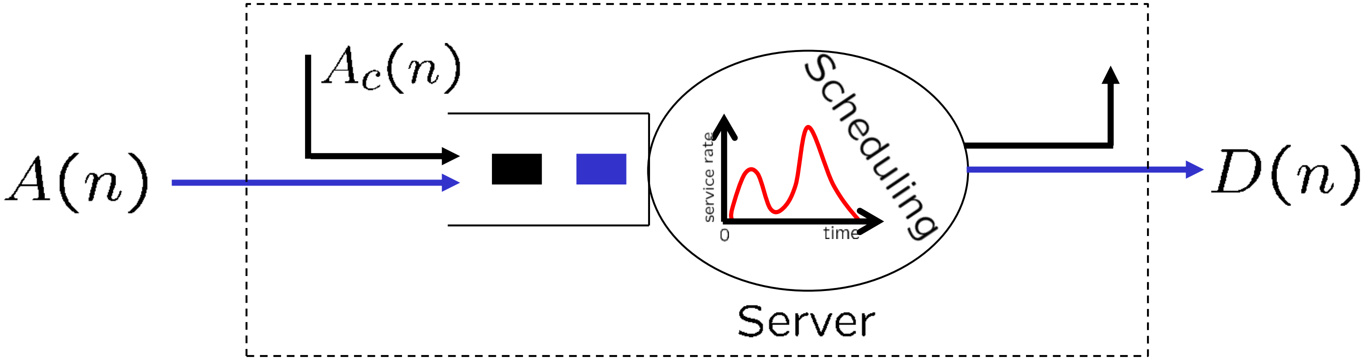

- For several scheduling algorithms

$$\text{FIFO}: \ \mathbb{P}(W > d) \le \gamma^n e^{-\theta \ldots}$$

$$\text{SP}: \ \mathbb{P}(W > d) \le \gamma^n e^{-\theta \ldots}$$

$$\text{EDF}: \ \mathbb{P}(W > d) \le \gamma^n e^{-\theta \ldots}$$

- Notes:
  - $0 < \gamma < 1$
  - previous exponential prefactors >1

Simulations. EDF

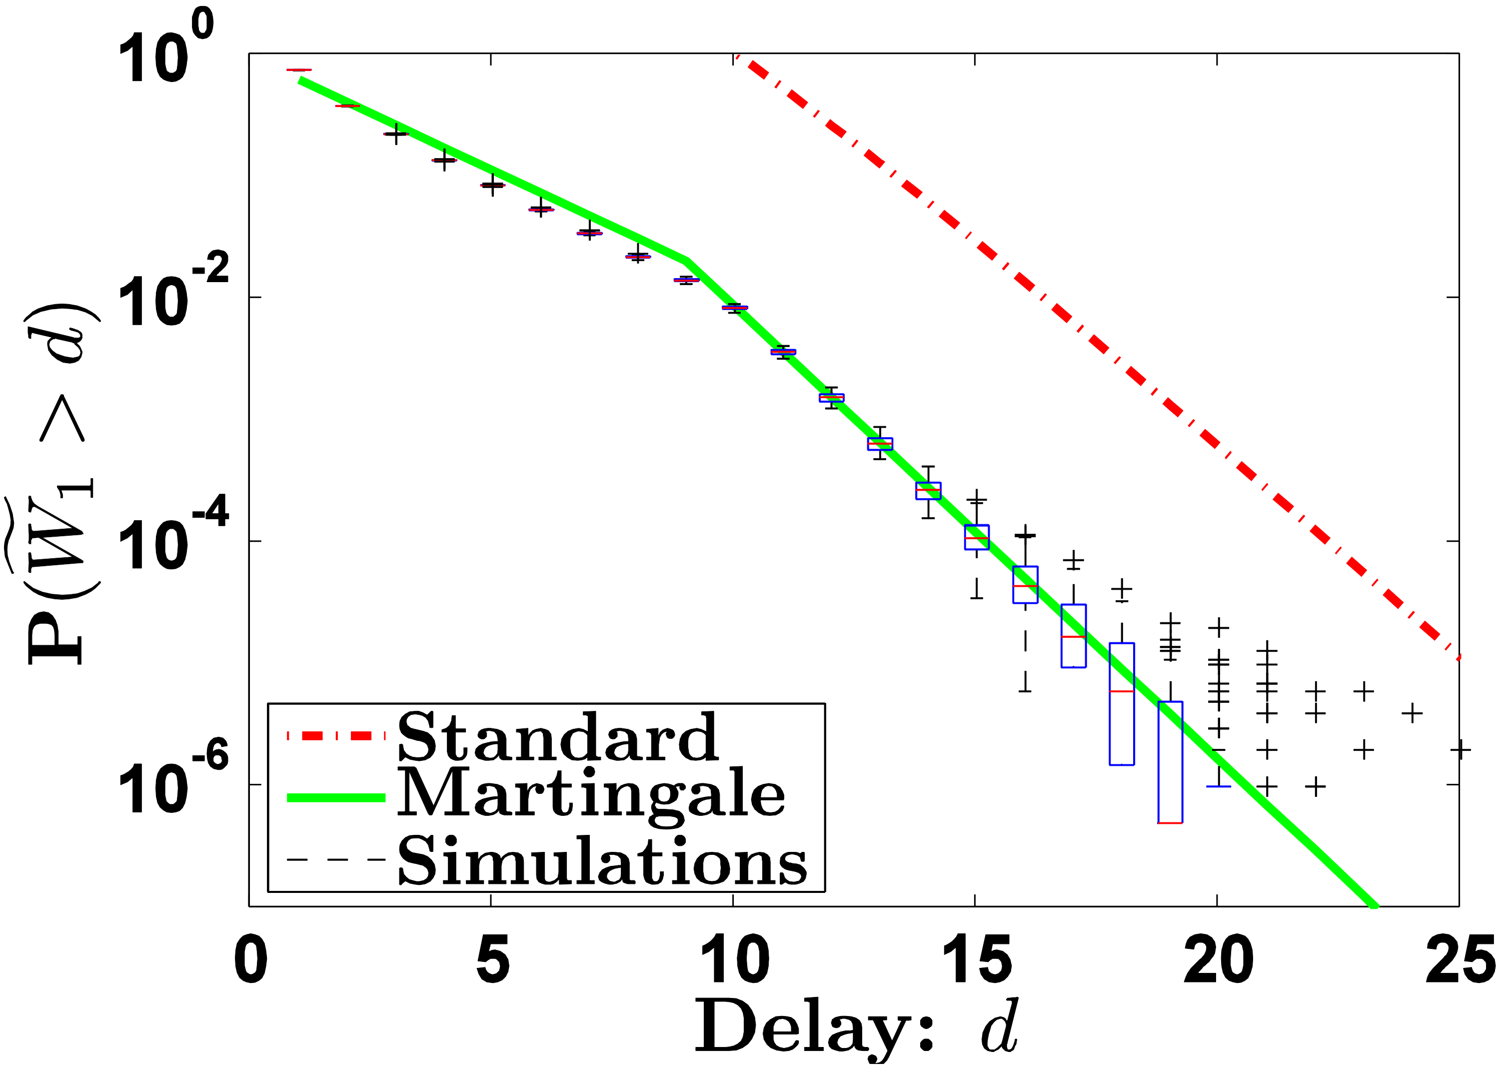

Application 2. Per-Flow Delay (access + queueing)

- A flow (Markov-Modulated, etc.) is competing at a wireless channel

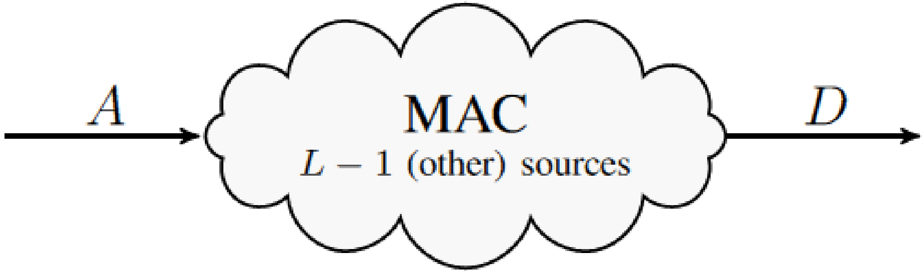

- MAC: Aloha or CSMA/CA

CSMA/CA

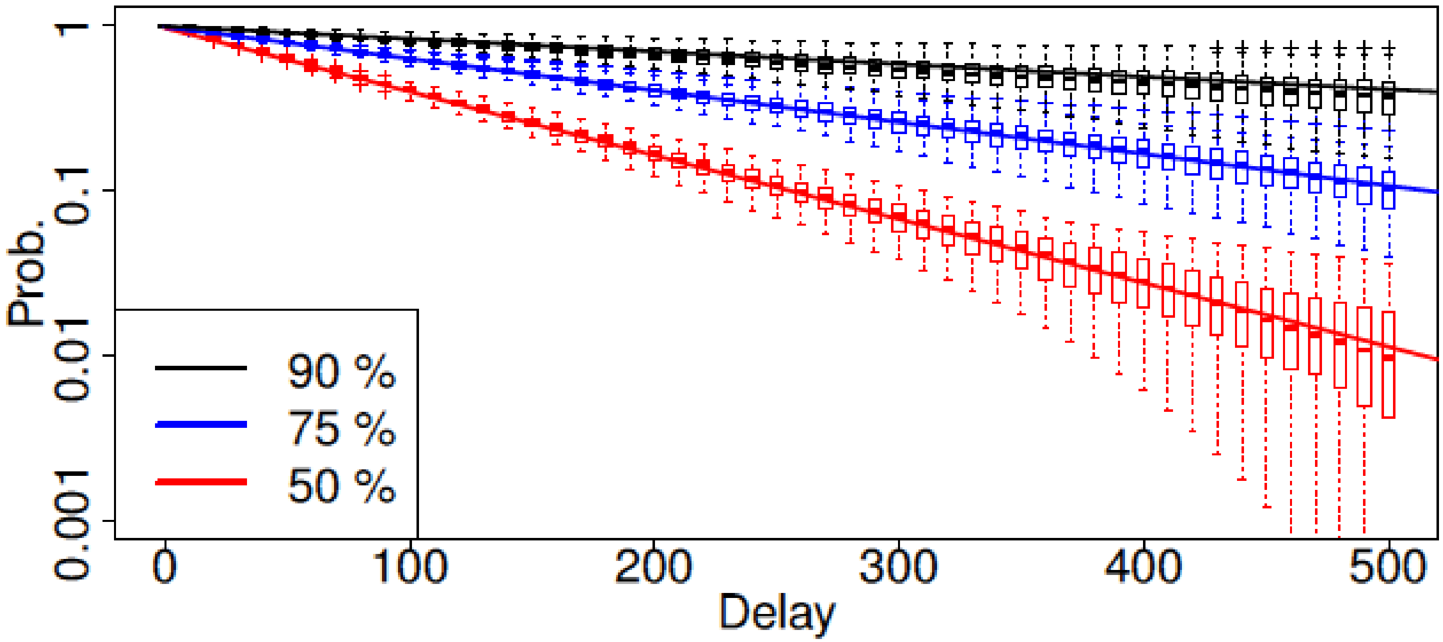

The Continuous Case. Markov Additive Processes

- Definition

A bivariate process $(A(t), M_t)_t$ is a *Markov Additive Process* if and only if

1. the pair $(A(t), M_t)$ is a Markov process in $\mathbb{R}^2$,
2. $A(0) = 0$ and $A(t)$ is nondecreasing,
3. the (joint and conditional) distribution of $(A(s,t), M_t \mid A(s), M_s)$ depends only on M_s.

- Examples
  - Markov Fluid
  - Markov Modulated Poisson Process (MMPP)
  - Markov Arrival Processes (MAP)
  - etc.

Example 1. Markov Fluid (MF)

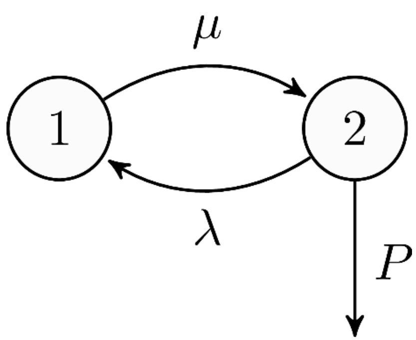

Let

$$\theta = \frac{\lambda}{P - C} - \frac{\mu}{C}\ , \quad h(P) = \frac{\theta C + \mu}{\mu}\ , \quad \text{and} \quad h(0) = 1\ .$$

Then the process

$$h(M_t)e^{\theta(A(t)-tC)}$$

is a martingale.

Example 2. Markov Modulated Poisson Process (MMPP)

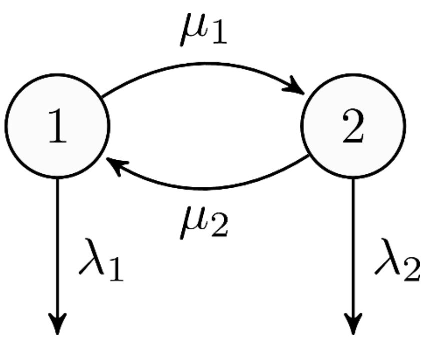

For $\theta > 0$, let T_θ denote the following 2×2-matrix:

$$T_\theta := \begin{pmatrix} \lambda_1 e^\theta - \mu_1 - \lambda_1 & \mu_1 \\ \mu_2 & \lambda_2 e^\theta - \mu_2 - \lambda_2 \end{pmatrix} .$$

Further, let $\lambda(\theta)$ denote its spectral radius. Pick $\theta > 0$ such that $\lambda(\theta) = \theta C$, and let $h = (h_1, h_2)$ denote an eigenvector corresponding to T_θ and $\lambda(\theta)$. Then the process

$$h(M_t)e^{\theta(A(t)-tC)}$$

is a martingale.

Example 3. Markov Arrival Process (MAP)

For $\theta > 0$, let $\lambda(\theta)$ denote the spectral radius of the matrix

$$D_0 + e^{\theta} D_1 \ ,$$

If $\lambda(\theta) = \theta C$, and h is a corresponding eigenvector, then the process

$$h(M_t) e^{\theta(A(t) - tC)}$$

is a martingale.

Multiplexing MAPs

In the situation with two MAPs, for $\theta > 0$, let $\lambda(\theta)$ and $\lambda'(\theta)$ denote the spectral radii of the matrices

$$D_0 + e^{\theta} D_1 \;, \text{and} \quad D_0' + e^{\theta} D_1' \;,$$

respectively. If $\lambda(\theta) + \lambda'(\theta) = \theta C$ and h a corresponding eigenvector, then the process

$$h(M_t) e^{\theta\left(A(t) + A'(t) - tC\right)}$$

is a martingale.
(!) No blow-up of numerical complexity.

The Take-Aways

- Modern (queueing) systems are complex

- Stochastic Network Calculus (SNC) started as a very promising approach (broad arrivals, scheduling, multi-node)

- Need for improvements/alternatives
  - Martingale approach (?)
  - Copula analysis (?)