Bandwidth Allocation in Large Stochastic Networks

Mathieu Feuillet

Soutenance de thèse

12/07/2012

Introduction

Modeling

Objectives:

- Modeling
- Design
- Dimensioning

What Are We Talking About?

- In a distributed storage system with failures, what is the life expectancy of a file?
- Does the Internet collapse if users are selfish and don’t use congestion control?
- Does CSMA/CA, as used in WiFi, ensure efficient use of bandwidth?

Contents

Modeling

Modeling

Objectives:

- Modeling
- Design
- Dimensioning

Tools:

- Markov processes
- Queueing models
- Scaling methods

Stochastic Models

State: $(X(t))$ a Markov jump process in $\mathbb{N}^d$:

- Number of files,
- Number of active flows in the Internet,
- Number of messages to be transmitted.

Stochastic Models

State: $(X(t))$ a Markov jump process in $\mathbb{N}^d$:

- Number of files,
- Number of active flows in the Internet,
- Number of messages to be transmitted.

Markov assumptions:

- Poisson arrivals
- Exponentially distributed sizes/durations.

Stochastic Models

State: $(X(t))$ a Markov jump process in $\mathbb{N}^d$:

- generally, non-reversible,
- when ergodic, invariant distribution not known,
- results on transient properties are rare (for $d \geq 2$).

Scaling Methods

Scaling Methods

Principle: N a scaling parameter

Analyze the evolution of the sample path of

$$\left(\frac{X^N(\Psi_N(t))}{\Phi_N}\right)$$

as $N \to \infty$, for some convenient $(\Psi_N(t))$ and (Φ_N).

Time scale $t \to \Psi_N(t)$ is used as a tool to focus on some specific part of sample paths.

Scaling Methods

Principle: N a scaling parameter

Analyze the evolution of the sample path of

$$\left(\frac{X^N(\Psi_N(t))}{\Phi_N}\right)$$

as $N \to \infty$, for some convenient $(\Psi_N(t))$ and (Φ_N).

Time scale $t \to \Psi_N(t)$ is used as a tool to focus on some specific part of sample paths.

There may be more than one time scale of interest!

Scaling Methods: Goals

Give a First order description of $(X^N(t))$:

$$X^N(\Psi_N(t)) \approx \Phi_N . x(t)$$

where,

$(x(t))$ is a simpler stochastic process or even a deterministic dynamical system:

$$\dot{x}(t) = F(x(t))$$

Classical Example: Fluid Limit

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = \|X(0)\|.$$

Scaling parameter: initial state

Time scale: $t \mapsto Nt$

Classical Example: Fluid Limit

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = \|X(0)\|.$$

Scaling parameter: initial state

Time scale: $t \mapsto Nt$

Fluid limit reaches 0

$\Downarrow$

Process is stable

Example: Fluid Limit of M/M/1 Queue

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = X(0).$$

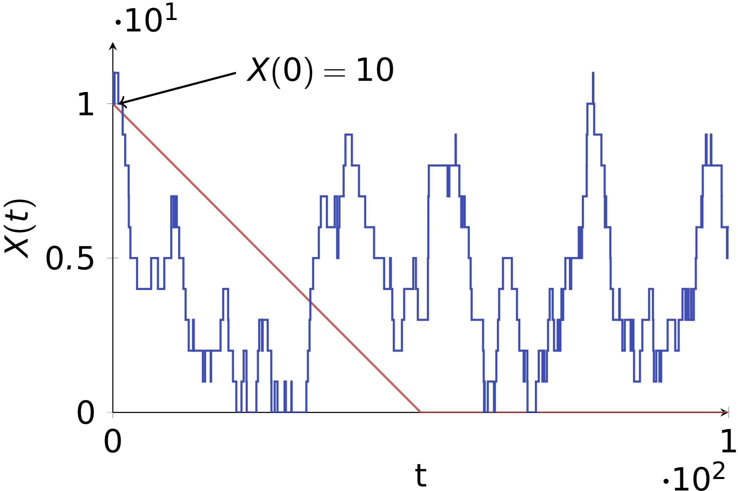

Example: Fluid Limit of M/M/1 Queue

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = X(0).$$

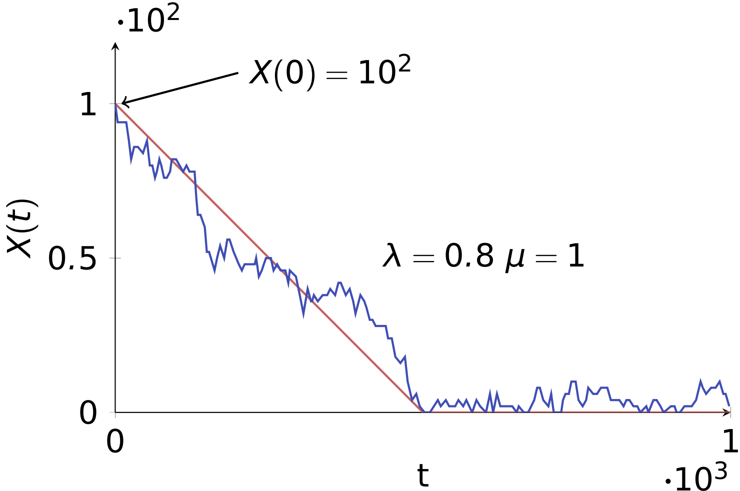

Example: Fluid Limit of M/M/1 Queue

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = X(0).$$

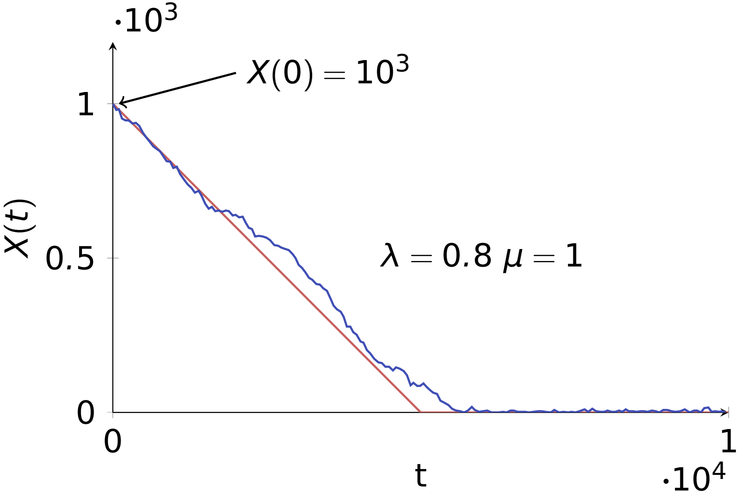

Example: Fluid Limit of M/M/1 Queue

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = X(0).$$

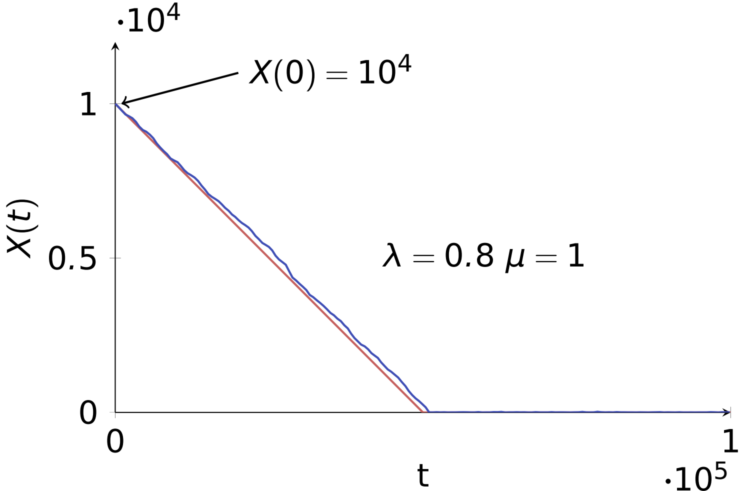

Example: Fluid Limit of M/M/1 Queue

$$\left(\bar{X}(t)\right) = \left(\frac{X(Nt)}{N}\right), \quad \text{with } N = X(0).$$

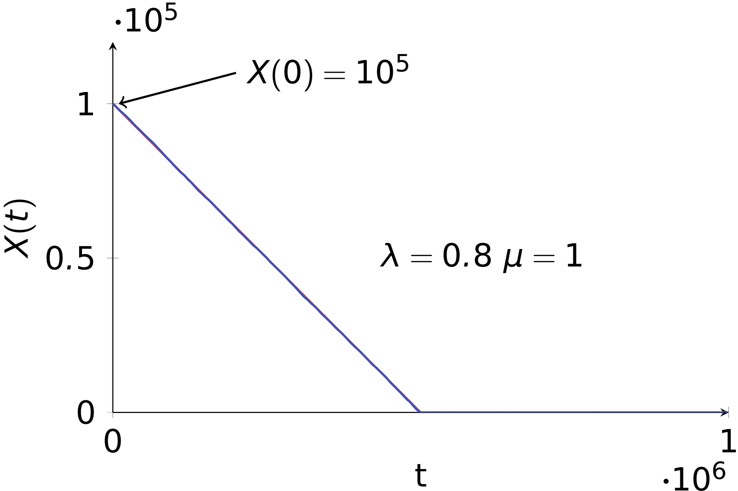

References

Fluid limits for queueing systems:

[Malyshev 93]
[Rybko-Stolyar 92]
[Dai 95]

Scaling methods:

[Khasminskii 56]
[Freidlin-Wentzell 79]
[Ethier-Kurtz 86]
[Robert 03]

Technical Corner

Proof of the tightness of the scaled process

$$\left(\frac{X^N\left(\Psi_N(t)\right)}{\Phi_N}\right)$$

- Stochastic Differential Equation representation of $(X^N(t))$ with martingales
- Standard tightness criteria

Technical Corner

Proof of the tightness of the scaled process

$$\left(\frac{X^N(\Psi_N(t))}{\Phi_N}\right)$$

- Stochastic Differential Equation representation of $(X^N(t))$ with martingales
- Standard tightness criteria

Difficulties:

- Discontinuities: Skorokhod Problem Techniques
- Stochastic averaging

Technical Corner

Proof of the tightness of the scaled process

$$\left(\frac{X^N(\Psi_N(t))}{\Phi_N}\right)$$

- Stochastic Differential Equation representation of $(X^N(t))$ with martingales
- Standard tightness criteria

Difficulties:

- Discontinuities: Skorokhod Problem Techniques
- Stochastic averaging

Technical Corner

Proof of the tightness of the scaled process

$$\left(\frac{X^N(\Psi_N(t))}{\Phi_N}\right)$$

- Stochastic Differential Equation representation of $(X^N(t))$ with martingales
- Standard tightness criteria

Difficulties:

- Discontinuities: Skorokhod Problem Techniques
- Stochastic averaging

Each example has its specific difficulties

Stochastic Averaging

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t)),$$
$$\dot{y}_N(t) = G(x_N(t), y_N(t))$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t)), \qquad \text{Fast time-scale}$$

$$\dot{y}_N(t) = G(x_N(t), y_N(t)) \qquad \text{Slow time-scale}$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t)), \qquad \text{Fast time-scale}$$
$$\dot{y}_N(t) = G(x_N(t), y_N(t)) \qquad \text{Slow time-scale}$$

Fast time-scale:

$$\dot{x}_N(t/N) = F(x_N(t/N)).$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t)),$$ Fast time-scale

$$\dot{y}_N(t) = G(x_N(t), y_N(t))$$ Slow time-scale

Fast time-scale:

$$\dot{x}_N(t/N) = F(x_N(t/N)).$$

Slow time-scale: If $x(t)$ tends to a fixed point $x^\star$: $(y_N(t))$ converges to $(y(t))$ with

$$\dot{y}(t) = G(x^\star, y(t))$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t), y_N(t)), \qquad \text{Fast time-scale}$$
$$\dot{y}_N(t) = G(x_N(t), y_N(t)) \qquad \text{Slow time-scale}$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t), y_N(t)), \qquad \text{Fast time-scale}$$
$$\dot{y}_N(t) = G(x_N(t), y_N(t)) \qquad \text{Slow time-scale}$$

Fast time-scale: When $N \to \infty$, $y_N(t/N) \approx z$

$$\dot{x}_N(t/N) \approx F(x(t/N), z)$$

A Deterministic Example

Deterministic sequences $(x_N(t))$ and $(y_N(t))$ with:

$$\dot{x}_N(t) = NF(x_N(t), y_N(t)), \qquad \text{Fast time-scale}$$
$$\dot{y}_N(t) = G(x_N(t), y_N(t)) \qquad \text{Slow time-scale}$$

Fast time-scale: When $N \to \infty$, $y_N(t/N) \approx z$

$$\dot{x_N}(t/N) \approx F(x(t/N), z)$$

Slow time-scale: If $(x_N(t))$ tends to a fixed point $x^\star_z$:
$(y_N(t))$ converges to $(y(t))$ with

$$\dot{y}(t) = G\left(x^\star_{y(t)}, y(t)\right)$$

Stochastic vs Deterministic

	Deterministic	**Stochastic**
Fast process	ODE $(x(t))$ $\dot{x} = F(x(t), y)$	Markov process $(X(t))$ $\Omega(y)$
Slow process	ODE $(y(t))$	Markov process $(Y(t))$
Equilibrium	Fixed point $x^{\star}_{y}$	Stationary distribution π_y
Convergence	Regularity of $y \mapsto x^{\star}_{y}$	Regularity of $y \mapsto \pi_y$

Stochastic vs Deterministic

	Deterministic	**Stochastic**
Fast process	ODE $(x(t))$ $\dot{x} = F(x(t), y)$	Markov process $(X(t))$ $\Omega(y)$
Slow process	ODE $(y(t))$	Markov process $(Y(t))$
Equilibrium	Fixed point $x_y^\star$	Stationary distribution π_y
Convergence	Regularity of $y \mapsto x_y^\star$	Regularity of $y \mapsto \pi_y$

References

Statistical mechanics:

[Bogolyubov 62]

Stochastic calculus:

[Khasminskii 68],
[Papanicolaou et al. 77],
[Freidlin-Wenzell 79].

Loss networks:

[Kurtz 92],
[Hunt-Kurtz 94]

Contributions

The Law of the Jungle:

- Stochastic averaging
- Scaling over the stationary distributions

Flow-Aware CSMA:

- Suboptimality of CSMA (mono/multi-channel)
- Optimality of Flow-Aware CSMA (mono/multi)
- Time-scale separation

An unreliable file system:

- Three time-scales
- Stochastic averaging (simpler proof)

Transient properties of Engset and Ehrenfest:

- Positive martingales
- Asymptotics on hitting times

Contributions

The Law of the Jungle:

- Stochastic averaging
- Scaling over the stationary distributions

Flow-Aware CSMA:

- Suboptimality of CSMA (mono/multi-channel)
- Optimality of Flow-Aware CSMA (mono/multi)
- Time-scale separation

An unreliable file system:

- Three time-scales
- Stochastic averaging (simpler proof)

Transient properties of Engset and Ehrenfest:

- Positive martingales
- Asymptotics on hitting times

Example 1:

An Unreliable File System

Model

βN files
2 copies/file

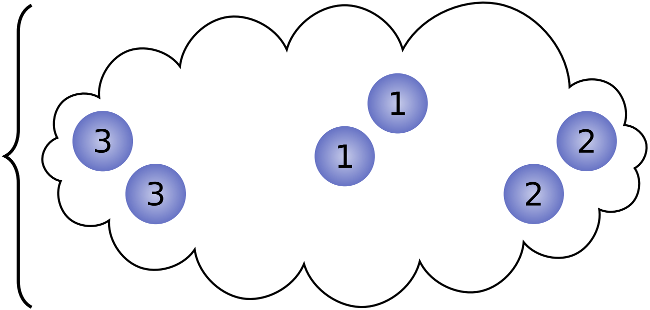

Back-up: λN

Model

βN files
2 copies/file

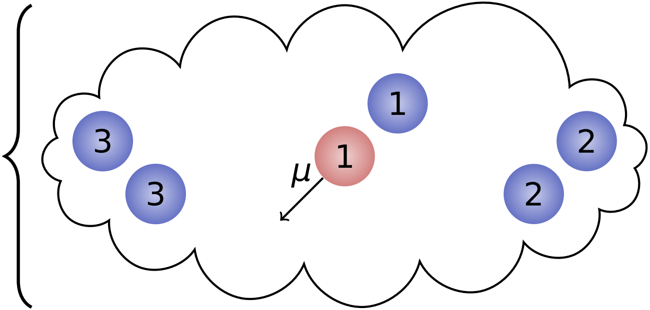

Back-up: λN

Each copy is lost at rate μ

Model

βN files
2 copies/file

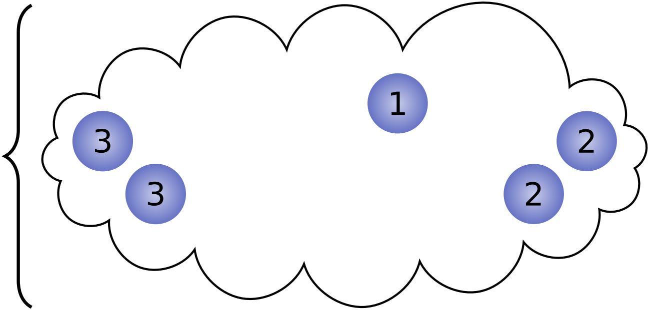

Back-up: λN

Model

βN files
2 copies/file

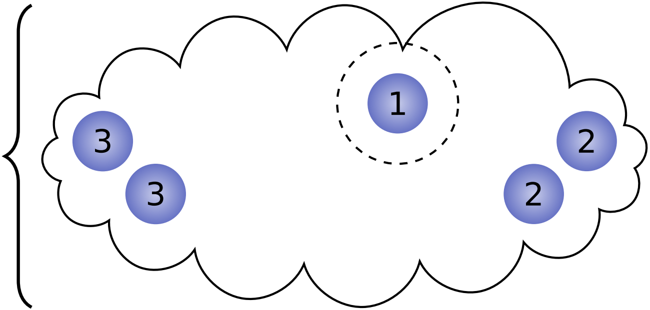

Back-up: λN

A file with 1 copy can be backed up

Model

βN files
2 copies/file

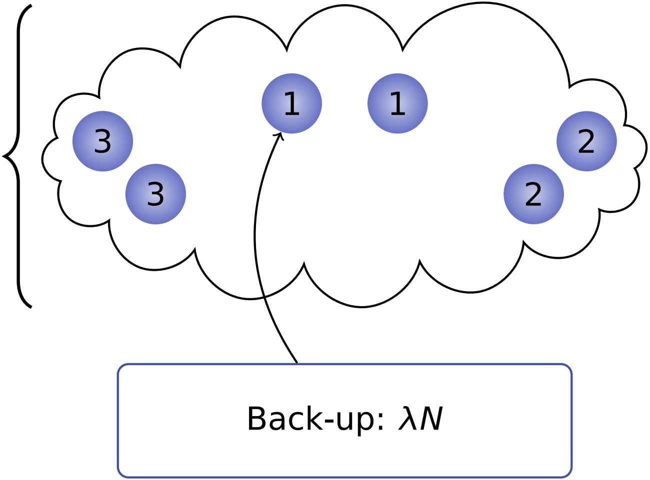

A file with 1 copy can be backed up

Model

βN files
2 copies/file

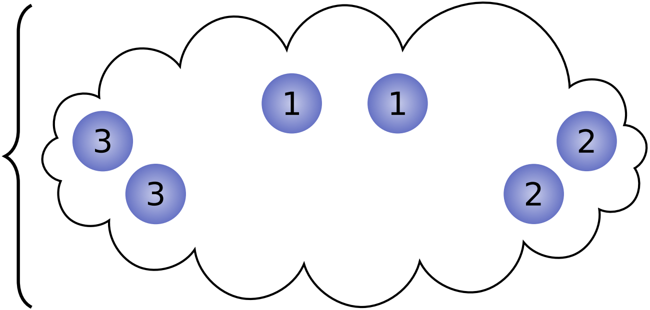

Back-up: λN

Model

βN files
2 copies/file

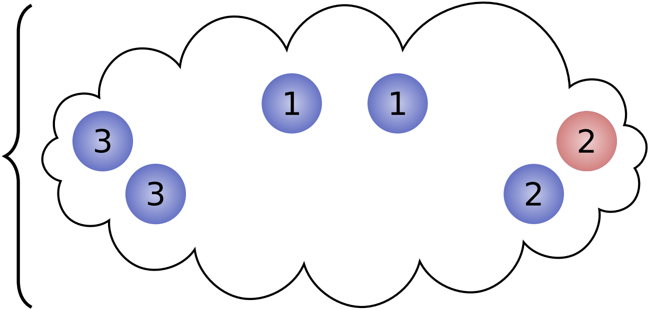

Back-up: λN

A file with 0 copies is lost

Model

βN files
2 copies/file

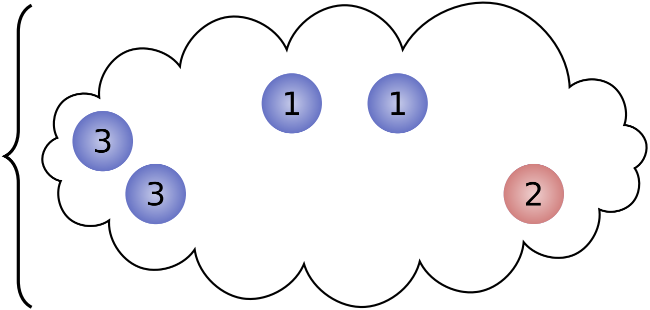

Back-up: λN

A file with 0 copies is lost

Model

βN files
2 copies/file

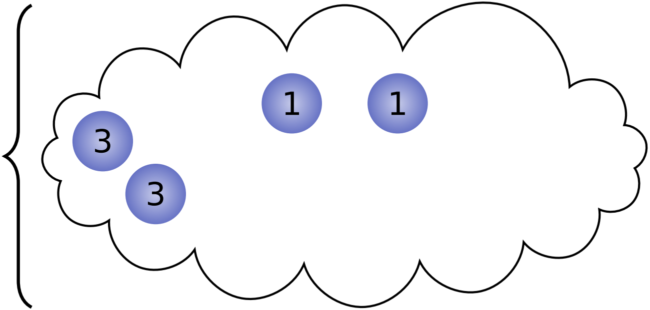

Back-up: λN

A file with 0 copies is lost

Model

βN files
2 copies/file

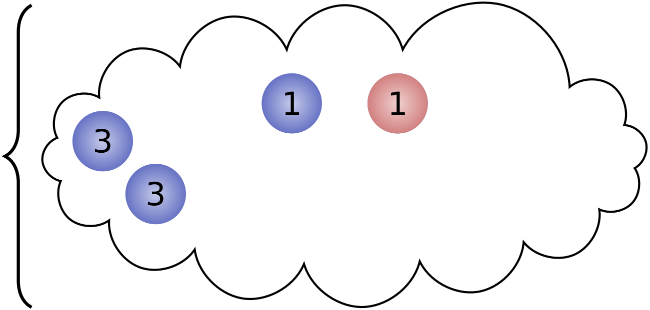

Back-up: λN

Model

βN files
2 copies/file

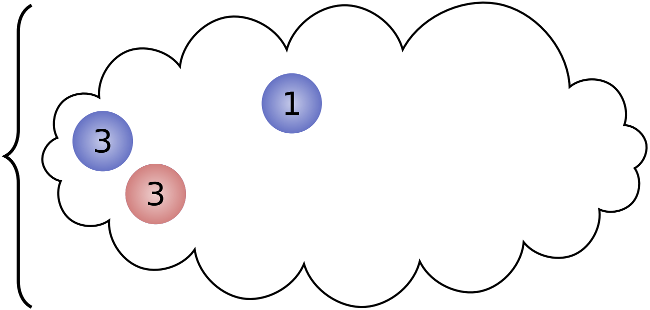

Back-up: λN

Model

βN files
2 copies/file

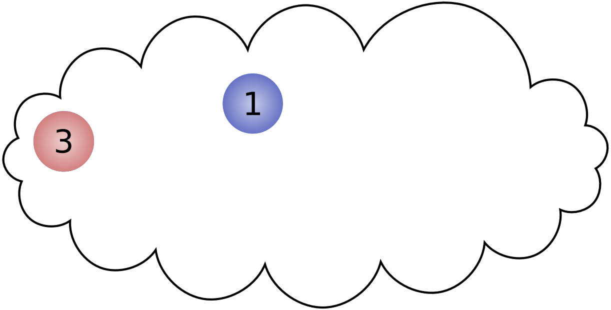

Back-up: λN

Model

βN files
2 copies/file

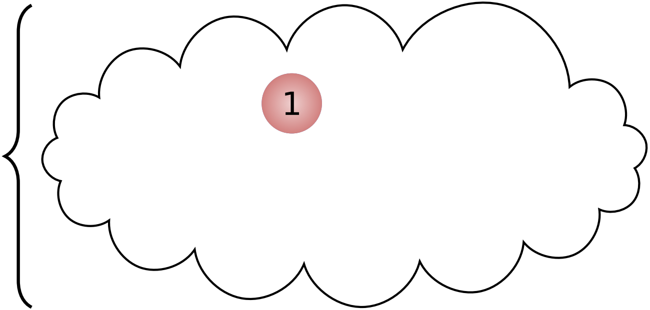

Back-up: λN

Model

βN files
2 copies/file

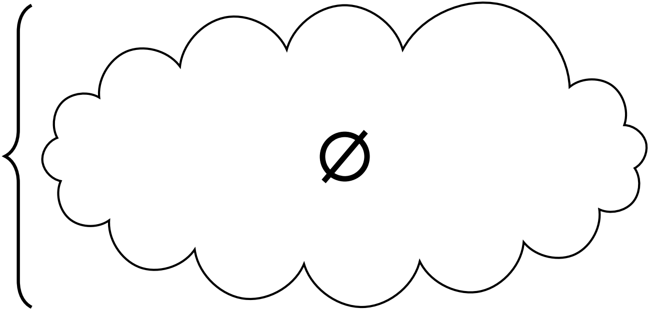

Back-up: λN

What is the decay rate of the network?

Model

$X_i(t)$: number of files with i copies at time t.

$(X_0(t), X_1(t), X_2(t))$: a transient Markov Process.

$$X_0(t) + X_1(t) + X_2(t) = \beta N.$$

A unique absorbing state $(\beta N, 0, 0)$.

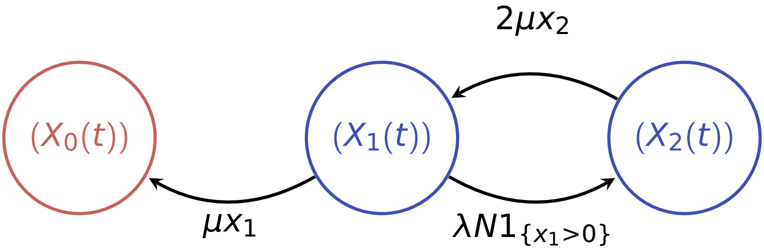

Model

$X_i(t)$: number of files with i copies at time t.

$(X_0(t), X_1(t), X_2(t))$: a transient Markov Process.

$$X_0(t) + X_1(t) + X_2(t) = \beta N.$$

A unique absorbing state $(\beta N, 0, 0)$.

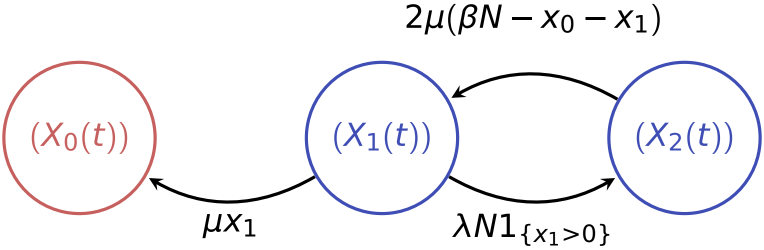

Different Behaviors

Three time scales:

$$\begin{cases} t & \to t/N \\ t & \to t \\ t & \to Nt \end{cases}$$

Three regimes:

Overload: $2\beta > \rho = \lambda/\mu$,

Critical load: $2\beta = \rho$,

Underload: $2\beta < \rho$.

Time scale: $t \to t/N$

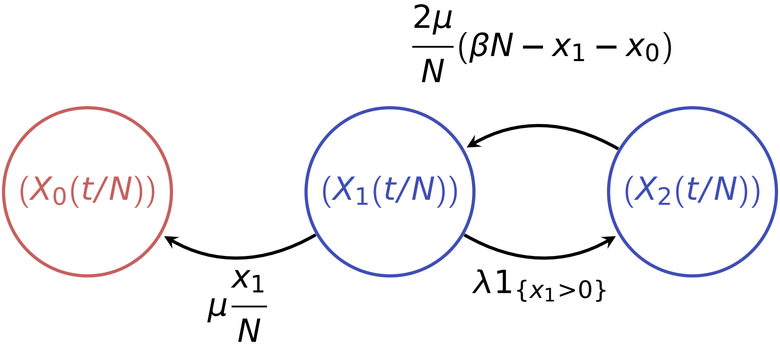

Time scale: $t \to t/N$

$(L_1(t))$: an $M/M/1$ queue $\begin{cases} \text{ergodic if } 2\beta < \rho, \\ \text{transient if } 2\beta > \rho. \end{cases}$

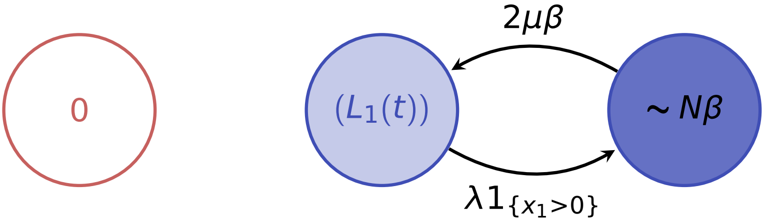

Time scale: $t \to t/N$

$(L_1(t))$: an $M/M/1$ queue $\begin{cases} \text{ergodic if } 2\beta < \rho, \\ \text{transient if } 2\beta > \rho. \end{cases}$

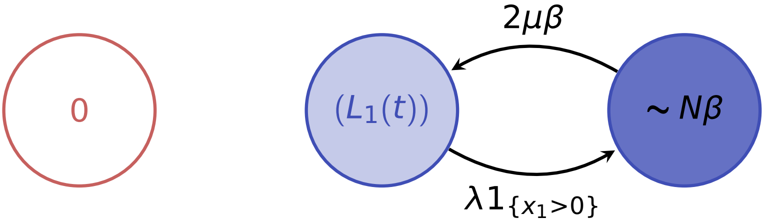

No loss!

Time scale: $t \to t$

Overloaded network

If $2\beta > \rho$, $(X_0(t)/N, X_1(t)/N, X_2(t)/N)$ converges to a deterministic process $(x_0(t), x_1(t), x_2(t))$.

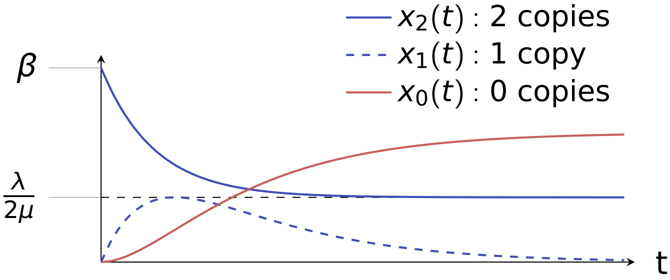

Time scale: $t \to t$

Overloaded network

If $2\beta > \rho$, $(X_0(t)/N, X_1(t)/N, X_2(t)/N)$ converges to a deterministic process $(x_0(t), x_1(t), x_2(t))$.

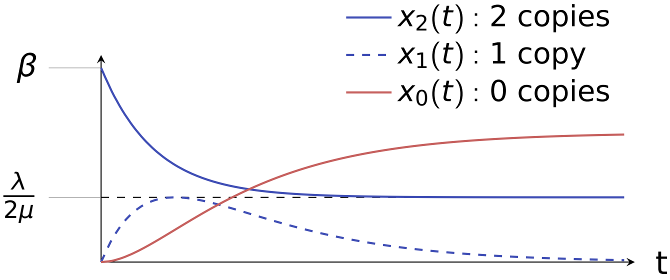

A fraction $N(\beta - \rho/2)$ is lost!

Time scale: $t \rightarrow t$
Underloaded network

If $2\beta < \rho$, $(X_0(t)/N, X_1(t)/N, X_2(t)/N)$ converges to

$$\begin{cases} x_0(t) & = 0, \\ x_1(t) & = 0, \\ x_2(t) & = \beta. \end{cases}$$

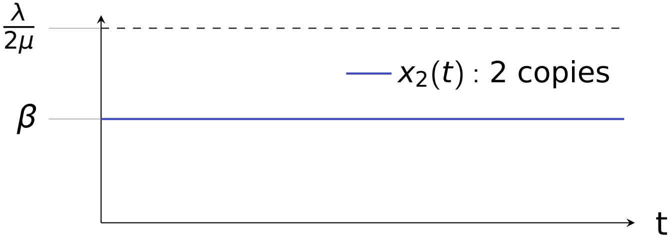

Time scale: $t \to t$
Underloaded network

If $2\beta < \rho$, $(X_0(t)/N, X_1(t)/N, X_2(t)/N)$ converges to

$$\begin{cases} x_0(t) & = 0, \\ x_1(t) & = 0, \\ x_2(t) & = \beta. \end{cases}$$

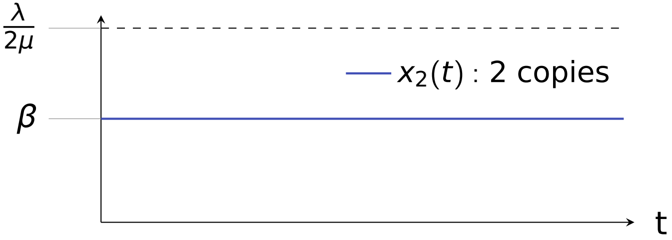

No significant loss!

Time Scale $t \to Nt$

$$\lim_{N \to +\infty} \left(\frac{X_0(Nt)}{N} \right) = \Psi(t),$$

where $\Psi(t)$ is the unique solution of

$$\Psi(t) = \mu \int_0^t \frac{2\mu(\beta - \Psi(s))}{\lambda - 2\mu(\beta - \Psi(s))} \, \mathrm{d}s.$$

Time Scale $t \to Nt$

$$\lim_{N\to+\infty}\left(\frac{X_0(Nt)}{N}\right)=\Psi(t),$$

where $\Psi(t)$ unique solution in $(0,\beta)$ of

$$(1-\Psi(t)/\beta)^{\rho/2}\,e^{\Psi(t)+t}=1.$$

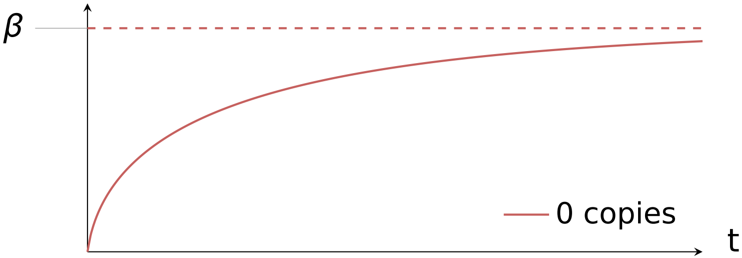

$t \to Nt$ is the “correct” time scale.

A Stochastic Averaging Phenomenon

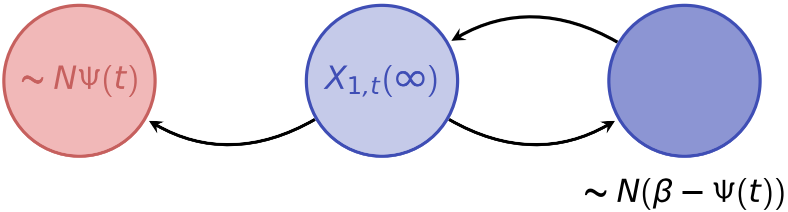

Fast time scale: At "time" Nt,
$(X_1(Nt{+}u/N), u \geq 0)$: an $M/M/1$ with transition rates:

$+1$ at rate $2\mu(\beta - \Psi(t))$

-1 at rate λ.

A Stochastic Averaging Phenomenon

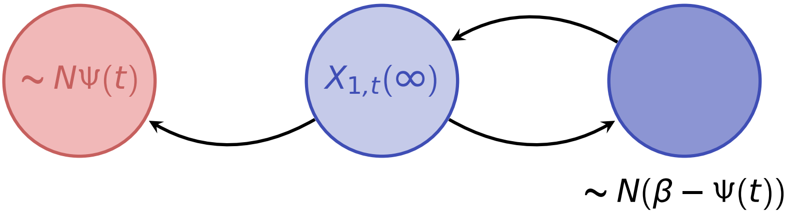

Slow time scale: $(X_0(Nt)/N)$ "sees" only X_1 at equilibrium:

$$\Psi(t)\text{"}=\text{"}\mu\int_0^t \mathbb{E}(X_{1,s}(\infty))\,\mathrm{d}s = \mu\int_0^t \frac{2\mu(\beta-\Psi(s))}{\lambda-2\mu(\beta-\Psi(s))}\,\mathrm{d}s.$$

Technical Corner

Step 1 Radon measures: tightness of (μ^N) with

$$\langle \mu^N, g \rangle = \frac{1}{N} \int_0^{Nt} g\left(X_1^N(s), s\right) \mathrm{d}s$$

Technical Corner

Step 1 Radon measures: tightness of (μ^N) with

$$\langle \mu^N, g \rangle = \frac{1}{N} \int_0^{Nt} g\left(X_1^N(s), s\right) \mathrm{d}s$$

Step 2 Control of limits of (μ^N):

$$\lim_{N\to\infty} \frac{1}{N} \int_0^{Nt} X_1^N(s)\, \mathrm{d}s = \Psi(t) = \int_0^t \langle \pi_s, I \rangle\, \mathrm{d}s$$

$$\int_0^t \pi_s(\mathbb{N})\, \mathrm{d}s = t$$

Technical Corner

Step 1 Radon measures: tightness of (μ^N) with

$$\langle \mu^N, g \rangle = \frac{1}{N} \int_0^{Nt} g\left(X_1^N(s), s\right) \mathrm{d}s$$

Step 2 Control of limits of (μ^N):

$$\lim_{N\to\infty} \frac{1}{N} \int_0^{Nt} X_1^N(s)\,\mathrm{d}s = \Psi(t) = \int_0^t \langle \pi_s, I \rangle \,\mathrm{d}s$$

$$\int_0^t \pi_s(\mathbb{N})\,\mathrm{d}s = t$$

Here: Proof by stochastic domination

Technical Corner

Step 1 Radon measures: tightness of (μ^N) with

$$\langle \mu^N, g \rangle = \frac{1}{N} \int_0^{Nt} g\left(X_1^N(s), s\right) \mathrm{d}s$$

Step 2 Control of limits of (μ^N):

$$\lim_{N\to\infty} \frac{1}{N} \int_0^{Nt} X_1^N(s)\,\mathrm{d}s = \Psi(t) = \int_0^t \langle \pi_s, I \rangle \,\mathrm{d}s$$

$$\int_0^t \pi_s(\mathbb{N})\,\mathrm{d}s = t$$

Here: Proof by stochastic domination

Step 3 Identification of π_s with martingale techniques and balance equations.

Decay Rate of the Network

$$T_N(\delta) = \inf\left\{t \geq 0 : \; X_0^N(t) \geq \delta\beta N\right\}$$

Theorem:

$$\lim_{N\to\infty} \frac{T_N(\delta)}{N} = -\frac{\rho}{2}\log(1-\delta) - \delta\beta.$$

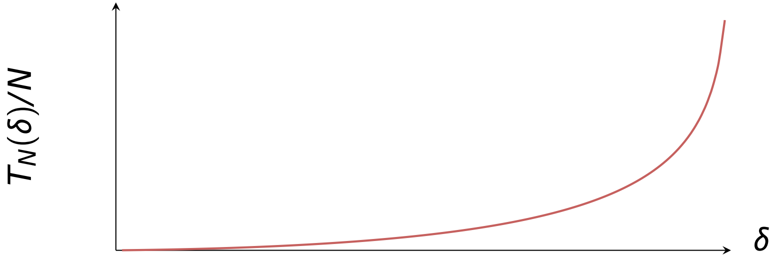

Conclusion

- Three different time scales
- A first example of stochastic averaging
- Asymptotics on a transitory property.

Extensions:

- Number of copies: $d > 2 \Rightarrow d - 1$ times scales
- Decentralized back-up (mean-field)

Open problem:

- Modeling a DHT: geometrical considerations

Example 2:

The Law of the Jungle

Context

Congestion control:

- Rate adjustment to limit packet loss
- Retransmission of lost packets

Context

Congestion control:

- Rate adjustment to limit packet loss
- Retransmission of lost packets

No congestion control:

- No rate adjustment
- Sources send at their maximum rate
- Coding to recover from packet loss

Context

Congestion control:

- Rate adjustment to limit packet loss
- Retransmission of lost packets

No congestion control:

- No rate adjustment
- Sources send at their maximum rate
- Coding to recover from packet loss

Does this bring congestion collapse?

Bandwidth Sharing Networks

[Massoulié Roberts 00]

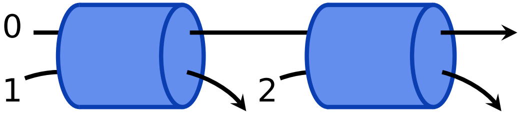

- A flow: a stream of packets
- Flows are considered as a fluid
- Users divided in classes/routes
- Poisson arrivals/Exponential sizes
- Resource allocation determined by congestion policy

Bandwidth Sharing Networks

[Massoulié Roberts 00]

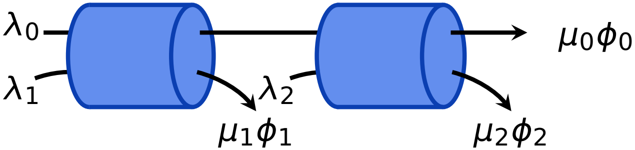

- A flow: a stream of packets
- Flows are considered as a fluid
- Users divided in classes/routes
- Poisson arrivals/Exponential sizes
- Resource allocation determined by congestion policy

Resource Allocation

Usually, α-fair policies are considered [MW00].

Here:

- Sources send at their maximum rate (1 or a)
- Tail dropping: At each link, output rates are proportional to input rates

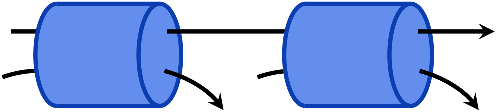

Resource Allocation

Usually, α-fair policies are considered [MW00].

Here:

- Sources send at their maximum rate (1 or a)
- Tail dropping: At each link, output rates are proportional to input rates

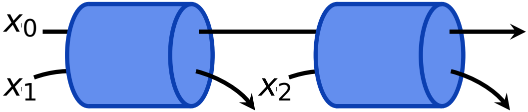

Resource Allocation

Usually, α-fair policies are considered [MW00].

Here:

- Sources send at their maximum rate (1 or a)
- Tail dropping: At each link, output rates are proportional to input rates

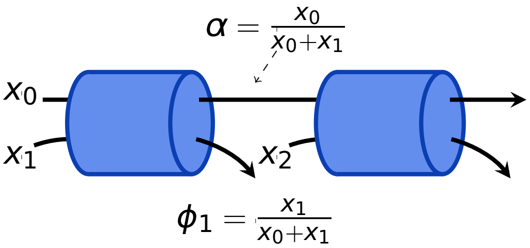

Resource Allocation

Usually, α-fair policies are considered [MW00].

Here:

- Sources send at their maximum rate (1 or a)
- Tail dropping: At each link, output rates are proportional to input rates

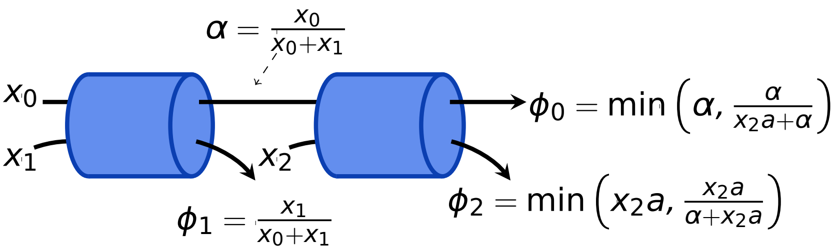

Ergodicity Condition

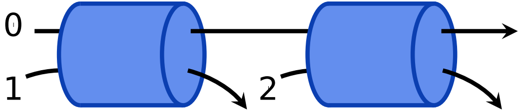

Optimal ergodicity condition:

$$\rho_0 + \rho_1 < 1, \quad \rho_0 + \rho_2 < 1$$

where $\rho_i = \lambda_i / \mu_i$.

We know α-fair policies are optimal [BM02].

Ergodicity Condition

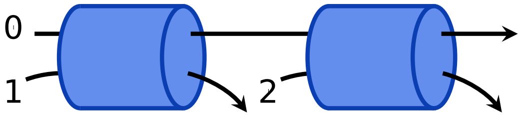

Optimal ergodicity condition:

$$\rho_0 + \rho_1 < 1, \quad \rho_0 + \rho_2 < 1$$

where $\rho_i = \lambda_i / \mu_i$.

We know α-fair policies are optimal [BM02].

What about our policy?

Fluid Limits

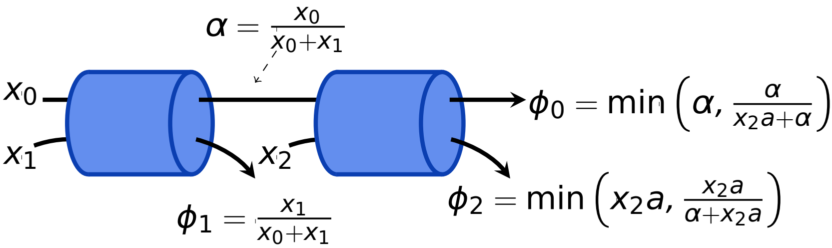

If $x_2 \gg 0$, class 2 uses virtually all the second link.
If $(z_0(t), z_1(t), z_2(t))$ is a fluid limit with $z_2(0) > 0$,

$$\begin{cases} \dot{z}_0(t) = \lambda_0, \\ \dot{z}_1(t) = \lambda_1 - \mu_1 \frac{z_1(t)}{z_0(t)+z_1(t)}, \\ \dot{z}_2(t) = \lambda_2 - \mu_2. \end{cases}$$

Fluid Limits

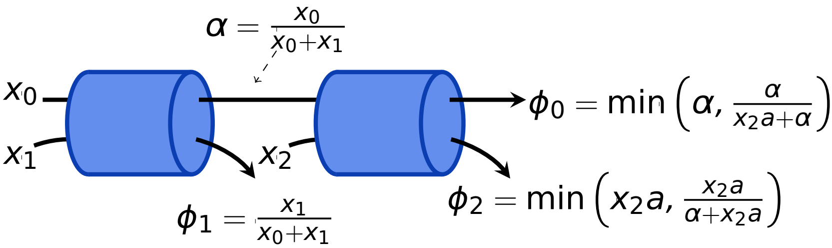

If $x_2 \gg 0$, class 2 uses virtually all the second link.
If $(z_0(t), z_1(t), z_2(t))$ is a fluid limit with $z_2(0) > 0$,

$$\begin{cases} \dot{z}_0(t) = \lambda_0, \\ \dot{z}_1(t) = \lambda_1 - \mu_1 \frac{z_1(t)}{z_0(t)+z_1(t)}, \\ \dot{z}_2(t) = \lambda_2 - \mu_2. \end{cases}$$

If $\rho_2 < 1$, $(z_2(t))$ reaches 0 in finite time.

Fluid Limits

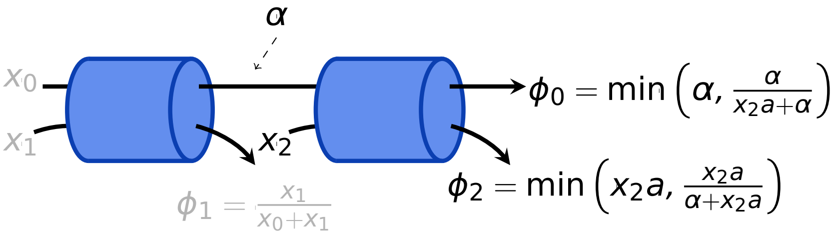

Classes 0 and 1 are frozen:

π_2^α is the stationary distribution of class 2

$$\bar{\Phi}_0(\alpha) = \mathbb{E}_{\pi_2^\alpha}\left(\Phi_0\left(\alpha, \frac{\alpha}{x_2 a + \alpha}\right)\right).$$

Fluid Limits

When $z_2(t) = 0$:

$$\begin{cases} \dot{z}_0(t) & = \lambda_0 - \mu_0 \bar{\phi}_0 \left(\frac{z_0(t)}{z_0(t) + z_1(t)} \right), \\ \dot{z}_1(t) & = \lambda_1 - \mu_1 \frac{z_1(t)}{z_0(t) + z_1(t)}, \\ \dot{z}_2(t) & = 0. \end{cases}$$

with

$$\bar{\Phi}_0(\alpha) = \mathbb{E}_{\pi_2^{\alpha}} \left(\Phi_0 \left(\alpha, \frac{\alpha}{x_2 a + \alpha} \right) \right).$$

Fluid Limits

When $z_2(t) = 0$:

$$\begin{cases} \dot{z}_0(t) & = \lambda_0 - \mu_0 \bar{\phi}_0 \left(\dfrac{z_0(t)}{z_0(t) + z_1(t)} \right), \\ \dot{z}_1(t) & = \lambda_1 - \mu_1 \dfrac{z_1(t)}{z_0(t) + z_1(t)}, \\ \dot{z}_2(t) & = 0. \end{cases}$$

with

$$\bar{\Phi}_0(\alpha) = \mathbb{E}_{\pi_2^\alpha} \left(\Phi_0 \left(\alpha, \frac{\alpha}{x_2 a + \alpha} \right) \right).$$

Stochastic averaging

Ergodicity Conditions

Ergodicity conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \bar{\phi}_0(1 - \rho_1)$$

Optimal conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \min(1 - \rho_1, 1 - \rho_2)$$

Ergodicity Conditions

Ergodicity conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \bar{\phi}_0(1 - \rho_1)$$

Optimal conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \min(1 - \rho_1, 1 - \rho_2)$$

But:

$$\bar{\phi}_0(1 - \rho_1) < \min(1 - \rho_2, 1 - \rho_1)$$

Ergodicity Conditions

Ergodicity conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \bar{\phi}_0(1 - \rho_1)$$

Optimal conditions:

$$\rho_1 < 1, \quad \rho_2 < 1,$$
$$\rho_0 < \min(1 - \rho_1, 1 - \rho_2)$$

But:

$$\bar{\phi}_0(1 - \rho_1) < \min(1 - \rho_2, 1 - \rho_1)$$

Not optimal!

Impact of Maximum Rate a

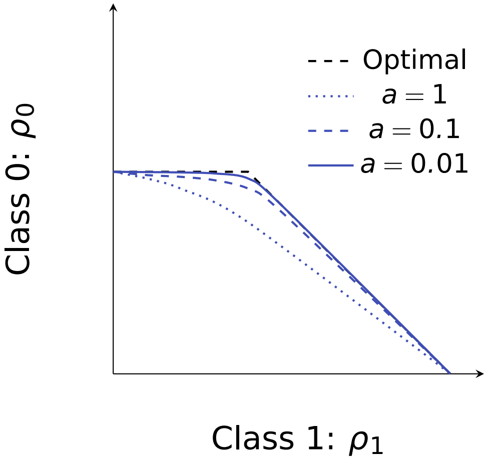

Impact of Maximum Rate a

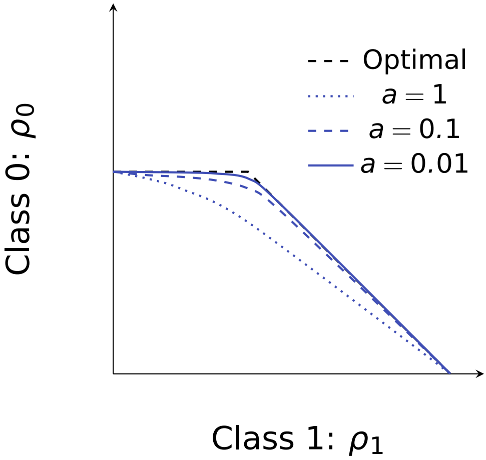

What happens when $a \to 0$?

Scaling the Maximum Rate a

We freeze α and consider the process $(X_2^S(t))$ with $\mathcal{Q}$-matrix:

$$q(x_2, x_2 + 1) = \lambda_2,$$

$$q(x_2, x_2 - 1) = \mu_2 \min\left(x_2 a, \frac{x_2 a}{\alpha + x_2 a}\right)$$

Scaling the Maximum Rate a

We freeze α and consider the process $(X_2^S(t))$ with $\mathcal{Q}$-matrix:

$$q(x_2, x_2+1) = \lambda_2,$$

$$q(x_2, x_2-1) = \mu_2 \min\left(x_2 \frac{a}{S}, \frac{x_2 a/S}{\alpha + x_2 a/S}\right)$$

Time-scale: $t \mapsto St$

Scaling the Maximum Rate a

We freeze α and consider the process $(X_2^S(t))$ with $\mathcal{Q}$-matrix:

$$q(x_2, x_2+1) = \lambda_2,$$

$$q(x_2, x_2-1) = \mu_2 \min\left(x_2 \frac{a}{S}, \frac{x_2 a/S}{\alpha + x_2 a/S}\right)$$

Time-scale: $t \mapsto St$

$(X_2^S(St)/S) \Rightarrow (x_2(t))$ with

$$\dot{x}_2(t) = \lambda_2 - \mu_2 \min\left(a x_2(t), \frac{x_2(t) a}{\alpha + x_2(t) a}\right)$$

Scaling the Maximum Rate a

We freeze α and consider the process $(X_2^S(t))$ with $\mathcal{Q}$-matrix:

$$q(x_2, x_2+1) = \lambda_2,$$

$$q(x_2, x_2-1) = \mu_2 \min\left(x_2\frac{a}{S}, \frac{x_2 a/S}{\alpha + x_2 a/S}\right)$$

Time-scale: $t \mapsto St$

$(X_2^S(St)/S) \Rightarrow (x_2(t))$ with

$$\dot{x}_2(t) = \lambda_2 - \mu_2 \min\left(a x_2(t), \frac{x_2(t)a}{\alpha + x_2(t)a}\right)$$

Fixed point:

$$x_2 = \frac{\rho_2}{a}\max\left(1, \frac{\alpha}{1-\rho_2}\right)$$

Scaling the Maximum Rate a

$$\begin{array}{ccc} (X_2^S(St)/S) & \xrightarrow{t\to\infty} & X_2^S(\infty)/S \\ {\scriptstyle S\to\infty}\Big\downarrow & & \Big\downarrow{\scriptstyle S\to\infty} \\ (x_2(t)) & \xrightarrow[t\to\infty]{} & x_2(\infty) \end{array}$$

Scaling the Maximum Rate a

$$\begin{array}{ccc} (X_2^S(St)/S) & \xrightarrow{t\to\infty} & X_2^S(\infty)/S \\ \Big\downarrow{\scriptstyle S\to\infty} & & \Big\downarrow{\scriptstyle S\to\infty} \\ (x_2(t)) & \xrightarrow[t\to\infty]{} & x_2(\infty) \end{array}$$

Convergence of processes

$\Downarrow$

Convergence of stationary distribution

Scaling the Maximum Rate a

$$\begin{array}{ccc} (X_2^S(St)/S) & \xrightarrow{t\to\infty} & X_2^S(\infty)/S \\ \Big\downarrow{\scriptstyle S\to\infty} & & \Big\downarrow{\scriptstyle S\to\infty} \\ (x_2(t)) & \xrightarrow[t\to\infty]{} & x_2(\infty) \end{array}$$

Convergence of processes

$\Downarrow$

Convergence of stationary distribution

$$\lim_{a\to 0} \bar{\Phi}_0(1-\rho_1) = \min(1-\rho_1, 1-\rho_2)$$

The policy is asymptotically optimal

Conclusion

- Analysis of equilibrium,
- Inversion of limits: scaling on stationary distributions
- Impact of access rates

Extensions:

- Linear networks with L links
- Second order scaling: speed of convergence.
- Upstream trees

Open problem:

- General acyclic networks

Example 3:

Flow-Aware CSMA

Model

The network is represented by a conflict graph

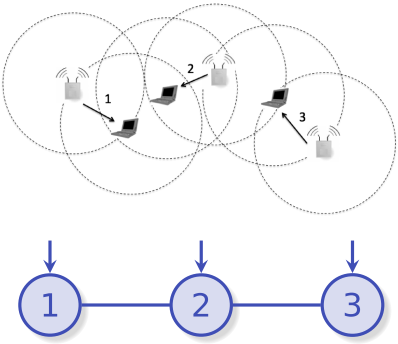

For each node i:

- $X_i(t) \in \mathbb{N}$: number of flows at time t
- $Y_i(t) = 1$ if node is active at time t, 0 otherwise.

Model

The network is represented by a conflict graph

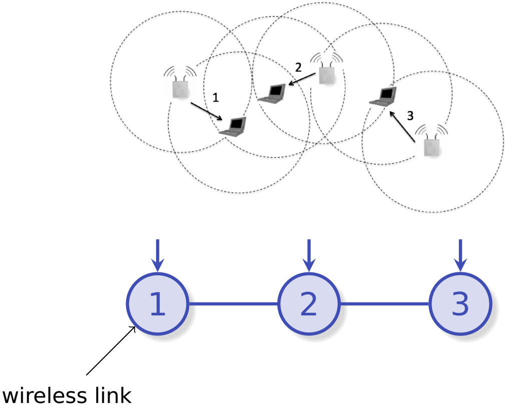

For each node i:

- $X_i(t) \in \mathbb{N}$: number of flows at time t
- $Y_i(t) = 1$ if node is active at time t, 0 otherwise.

Model

The network is represented by a conflict graph

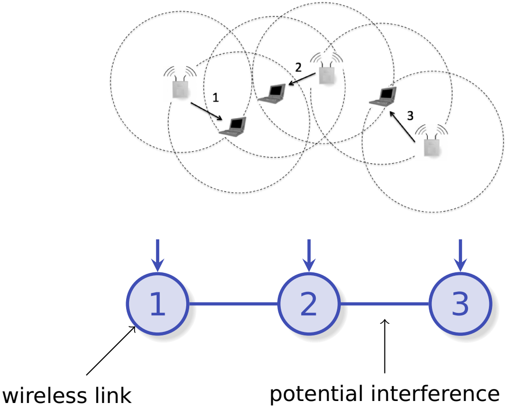

For each node *i*:

- $X_i(t) \in \mathbb{N}$: number of flows at time t
- $Y_i(t) = 1$ if node is active at time t, 0 otherwise.

Conflict Graph

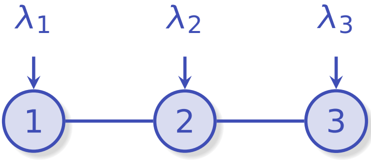

Schedules: $\varnothing$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Conflict Graph

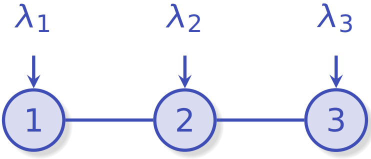

Schedules: $\varnothing$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Conflict Graph

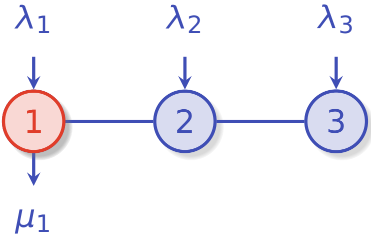

Schedules: $\varnothing$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Conflict Graph

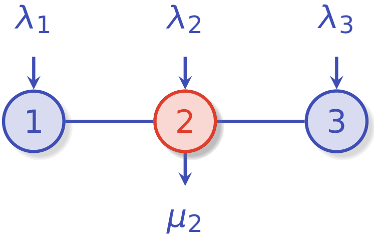

Schedules: ∅, {1}, {2}, {3}, {1, 3}.

Conflict Graph

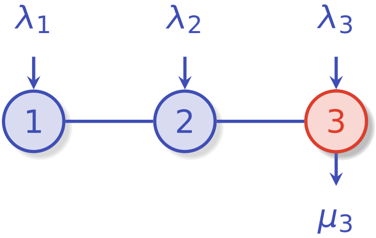

Schedules: $\varnothing$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Conflict Graph

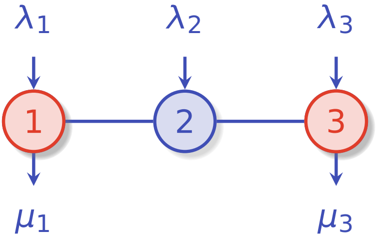

Schedules: $\emptyset$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Conflict Graph

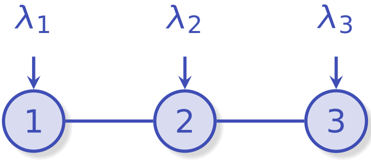

Schedules: $\emptyset$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Optimal stability region: convex hull of schedules

Conflict Graph

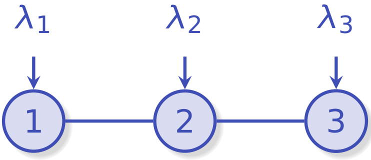

Schedules: $\emptyset$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1, 3\}$.

Optimal stability region: convex hull of schedules
In this example: $\{\rho_1 + \rho_2 \leq 1,\ \rho_2 + \rho_3 \leq 1\}$
with $\rho_i = \lambda_i/\mu_i$

Conflict Graph

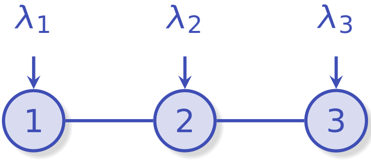

Schedules: $\emptyset$, $\{1\}$, $\{2\}$, $\{3\}$, $\{1,3\}$.

Optimal stability region: convex hull of schedules
In this example: $\{\rho_1 + \rho_2 \leq 1,\ \rho_2 + \rho_3 \leq 1\}$
with $\rho_i = \lambda_i/\mu_i$

Stability region?

Standard CSMA

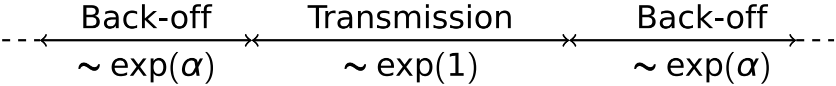

Standard CSMA

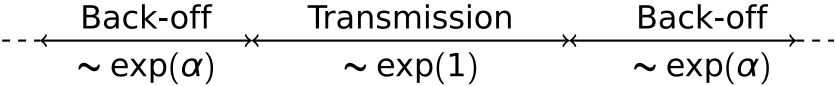

Optimal?

Standard CSMA

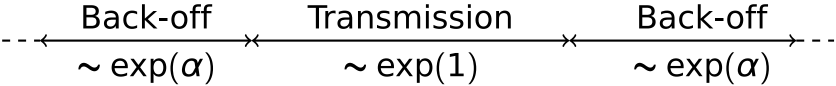

Optimal?

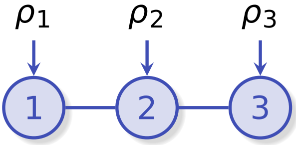

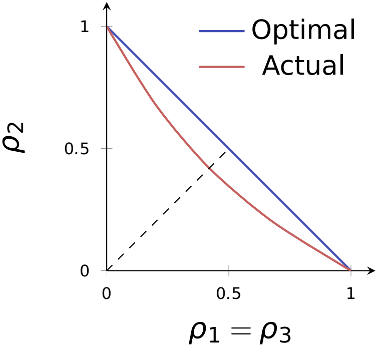

Flow-Aware CSMA

Proposed modification of CSMA:

Exponential backoff time for each flow

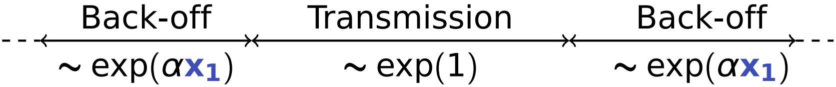

Flow-Aware CSMA

Proposed modification of CSMA:

Exponential backoff time for each flow

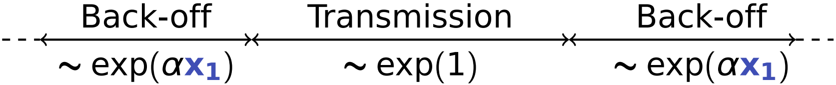

The process $(X(t), Y(t))$ is difficult to analyze:

Flow-Aware CSMA

Proposed modification of CSMA:

Exponential backoff time for each flow

The process $(X^N(t), Y^N(t))$ is difficult to analyze:

Idea: Separate network dynamics and flow dynamics.
When $N \to \infty$, $(Y^N(t))$: classical loss network.

Flow-Aware CSMA

Proposed modification of CSMA:

Exponential backoff time for each flow

The process $(X^N(t), Y^N(t))$ is difficult to analyze:

Idea: Separate network dynamics and flow dynamics.
When $N \to \infty$, $(Y^N(t))$: classical loss network.

Stochastic averaging

Optimality of Flow-Aware CSMA

Theorem:

Flow-aware CSMA algorithm is optimal for any network.

Sketch of proof:

- Asymptotically behaves as Max-Weight.
- Deduce a Lyapunov function and apply Foster's criterion.

Conclusion

- An optimal and fully distributed channel access mechanism
- Limiting process: jump process
- Simplification of the problem

Extension:

- Multi-channel

Open problem:

- Initial problem still open

General Conclusion

Three examples:

- Capacity of an unreliable file system
- Law of the Jungle
- Flow-Aware CSMA

General Conclusion

Three examples:

- Capacity of an unreliable file system
- Law of the Jungle
- Flow-Aware CSMA

Mathematical tools:

- Several examples of scalings
- A simpler proof for stochastic averaging

General Conclusion

Three examples:

- Capacity of an unreliable file system
- Law of the Jungle
- Flow-Aware CSMA

Mathematical tools:

- Several examples of scalings
- A simpler proof for stochastic averaging

and. . .

- Scalings: A set of powerful tools
- Stochastic averaging: a not so rare phenomenon

General Conclusion

Three examples:

- Capacity of an unreliable file system
- Law of the Jungle
- Flow-Aware CSMA

Mathematical tools:

- Several examples of scalings
- A simpler proof for stochastic averaging

and. . .

- Scalings: A set of powerful tools
- Stochastic averaging: a not so rare phenomenon

Many interesting open questions. . .

Thank you!