PERFORMANCE OF RANDOM MULTI-ACCESS ALGORITHMS IN LARGE NETWORKS

Charles Bordenave

CNRS & University of Toulouse

joint work with **David McDonald** (Univ. Ottawa) and **Alexandre Proutière** (Microsoft)

RANDOM MULTI-ACCESS NETWORKS

- N users share a common communication channel.
- Users are in interaction: two interacting users cannot successfully access the channel at the same time.
- An active user is a user accessing the channel.
- If two interacting users are active simultaneously, then there is a collision.
- Otherwise, an active user transmits successfully.
- At each time slot, if none of its interacting users is already active (Carrier Sense), user i becomes active with probability p_i independently of the others.
- A successful active user sends a packet in L slots.
- Colliding active users hold the channel for a duration L_c.

FULL INTERACTION - UNIT PACKET DURATION

- N independent arrival processes with intensity $(\lambda_i)_{1 \leq i \leq N}$.

- Full interaction: at most one successful transmission per slot.

- Unit packet duration: $L = L_c = 1$.

- User i is active with fixed probability p_i.

STABILITY REGION

What is the set Λ^N *of vectors* $(\lambda_i, p_i)_{1 \leq i \leq N}$ *such that the system is stable ?*

$\longrightarrow$ If $(\lambda_i, p_i) = (\lambda, p)$, the stability region is $\lambda < p(1-p)^{N-1}$.

$\longrightarrow$ If $N = 2$, the answer is known *Tsybakov-Mikhailov (79)*,...

$\longrightarrow$ If $N = 3$, the answer is known only for a Bernoulli arrival process, *Szpankowski (94)*.

$\longrightarrow$ If $N \geq 4$, the answer is unknown and it depends on the detailed characteristics of the arrival process.

$\longrightarrow$ *Anantharam (91)* has established the stability region for all N when the arrival processes have a particular joint dependent distribution.

ASYMPTOTIC STABILITY REGION

The analysis is very difficult: due to the interaction of the users, the queues are dependent.

$\Longrightarrow$ KEY IDEA: if N goes large, the evolutions of the queues become independent and depend only on the empirical distribution of the states of the queues.

We will try to compute the asymptotic stability region as N goes large. Consider a sequence $(\lambda_i^N, p_i^N)_{1\leq i\leq N}$ is the system stable for N large enough ?

We are interested in the regime $\inf_{i,j,N} p_i^N/p_j^N > 0$ and $\inf_{i,j,N} \lambda_i^N/\lambda_j^N > 0$: all traffic parameters have the same order of magnitude.

ASYMPTOTIC STABILITY REGION

We define the approximate stability region

$$\hat{\Lambda}^N = \left\{ \lambda \in \mathbb{R}_+^N : \exists \rho \in [0,1]^N : \forall i, \lambda_i \leq \rho_i p_i \prod_{j \neq i} (1 - \rho_j p_j) \right\}.$$

Assume that the users belong to a finite set of classes $\mathcal{V}$, $(\lambda_i^N, p_i^N) = (\lambda_v/N, p_v/N)$ and $\sum_i p_i^N \leq 1$. Let $\lambda^N = (\lambda_1^N, \cdots, \lambda_N^N)$ and $1^N = (1, \cdots, 1)/N$,

Theorem 1. *For all* $\epsilon > 0$*, there exists* N_ϵ *such that, for* $N > N_\epsilon$*,*
(a) if $\lambda^N + \epsilon \cdot 1^N \in \hat{\Lambda}^N$*, then the system is stable;*
(b) if $\lambda^N - \epsilon \cdot 1^N \notin \hat{\Lambda}^N$*, then the system is unstable.*

ASYMPTOTIC STABILITY REGION

$\longrightarrow$ Roughly speaking, the approximate stability region is obtained assuming that the evolutions of the queues of the various users are independent: decoupling heuristic.

$\longrightarrow$ The approximate stability region does not depend on detailed traffic characteristics.

$\longrightarrow$ Asymptotically, the set of traffic intensities $(\lambda_1, \ldots, \lambda_N)$ such that there exist transmission probabilities $(p_1, \ldots, p_N)$ stabilizing the system is the set

$$\mathcal{M} = \left\{ \lambda : \exists p_1, \ldots, p_N \in (0,1) : \forall i, \lambda_i \leq p_i \prod_{j \neq i} (1 - p_j) \right\}.$$

This result had been conjectured, *Tsybakov-Mikhailov (79)*.

HOW CLOSE IS THE APPROXIMATE STABILITY REGION ?

- For $N = 2$, $\Lambda^2 = \hat{\Lambda}^2$.

- For general N, using an argument borrowed from *Bonald et al. (04)* we can show that Λ^N and $\hat{\Lambda}^N$ coincide along curves where the traffic is homogeneous.

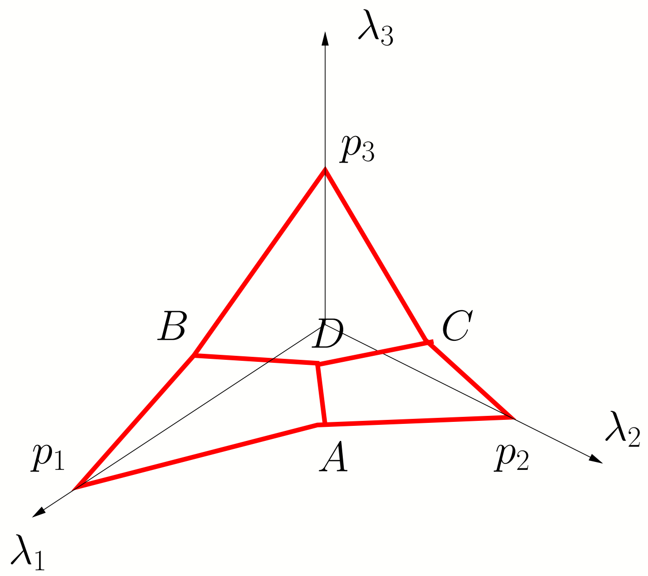

Curves where $\partial\hat{\Lambda}^3$ and $\partial\Lambda^3$ coincide. $A = (p_1(1-p_2), p_2(1-p_1), 0)$; $B = (p_1(1-p_3), 0, p_3(1-p_1))$; $C = (0, p_2(1-p_3), p_3(1-p_2))$; $D = (p_1(1-p_2(1-p_3), p_2(1-p_1(1-p_3), p_3(1-p_1)(1-p_2))$.

UNDERLYING MARKOV CHAIN

The state of a user i is

$$X_i^N(t) = (v_i, B_i^N(t), A_i^N(t)),$$

where v_i is the class of user i, $B_i^N(t)$ is the size of its buffer at time slot t and $A_i^N(t)$ is a finite state irreducible Markov chain describing its arrival process.

We assume for simplicity that the arrival processes are Bernoulli random variable (we remove A_i^N).

The empirical measure is

$$\nu^N(t) = \frac{1}{N} \sum_{i=1}^{N} \delta_{X_i^N(t)}.$$

$\Longrightarrow$ The process $(\nu^N(t)), t \in \mathbb{N}$, is a Markov chain on the set of probability measures on $\mathcal{V} \times \mathbb{N}$. We use a mean-field technique to analyse the convergence of the evolution as N goes to infinity.

MEAN-FIELD METHOD

As the number of particles N goes to infinity

- Propagation of chaos *($\Leftrightarrow$ decoupling)*.

- Convergence of the empirical measure process to a deterministic measure,

are nearly equivalent properties *Sznitman (91)*.

We aim at deriving the limiting evolution of the empirical measure of the buffers.

Examples of convergence of empirical measures in communication networks:

- TCP connections, *Baccelli-McDonald-Reynier (04-06)*.

- For exponential backoff algorithm, decoupling heuristic due to *Bianchi (00)*, made rigorous in *B-McDonald-Proutière (05-07), Sharma-Ganesh-Key (06)*.

- Multiuser receivers, *Tse and Hanly (99)* with random matrix theory.

PROBABILITY TRANSITIONS

If user/particle i in class v is in state $x = (v, k)$. The state becomes $(v, k+1)$ with probability:

$$\lambda_v/N + o(1/N).$$

and $(v, k-1)$ with probability:

$$\begin{aligned} & 1_{k>0} \frac{p_v}{N(1-\frac{p_v}{N})} \prod_{w \in \mathcal{V}} (1 - \frac{p_w}{N})^{\beta_w^N \nu_w^+ N} + o(1/N) \\ = \; & 1_{k>0} \frac{p_v}{N} \exp\left(-\sum_{w \in \mathcal{V}} \beta_w^N \nu_w^+ p_w\right) + o(1/N). \end{aligned}$$

where ν_v^+ is the proportion of class-v users with non-empty buffers and β_v^N is the proportion of class-v users.

$\longrightarrow$ The state of a user changes every $\Omega(N)$ slots.

LIMIT TIME EVOLUTION

Let $\beta_v = \lim_N \beta_v^N$. We accelerate time and define $q_i^N(t) = X_i^N(\lfloor Nt \rfloor)$. We prove a mean-field convergence theorem and deduce:

$\longrightarrow$ The evolutions of the users become asymptotically independent.

$\longrightarrow$ If $Q_{(v,k)}(t) = \lim_N \mathbb{P}(q_i^N(t) = k | v_i = v)$ then

$$\begin{aligned}\frac{\partial}{\partial t} Q_{(v,k)}(t) \quad = \quad & \lambda_v \left(1_{k>0} Q_{(v,k-1)}(t) - Q_{(v,k)}(t)\right) \\ & + p_v e^{-\gamma(t)} \left(Q_{(v,k+1)}(t) - 1_{k>0} Q_{(v,k)}(t)\right).\end{aligned}$$

with

$$\gamma(t) = \sum_v \beta_v p_v \nu_v^+(t) = \sum_w \beta_w p_w (1 - Q_{(w,0)}(t)).$$

For each $v \in \mathcal{V}$, these equations are the Kolmogorov equations of an M/M/1 queue with time-varying capacity equal at time t to $p_v \exp\left(-\gamma(t)\right)$.

SKETCH OF PROOF

Step 1. We prove the above mean-field convergence.

Step 2. We compute the stability region of this system of coupled differential equation $Q_{v,k}$, $(v, k) \in \mathcal{V} \times \mathbb{N}$. This is the stability region of $|\mathcal{V}|$ independent M/M/1 queues with time-varying capacity depending of their joint marginal distributions. Analyzing the existence of stationary distributions is easy, proving the global stability is harder.

Step 3. We come back to the stability region for finite N and prove the theorem by recursion on the number of classes. The argument is based on a geometric property of $\hat{\Lambda}^N$. We show that $\hat{\Lambda}^N$ is the region lying below one of N boundaries $\partial_j \hat{\Lambda}^N = \left\{ \lambda : \exists \rho \in \partial_j [0,1]^N, \forall i, \lambda_i = \rho_i p_i \prod_{k \neq i} (1 - \rho_k p_k) \right\}$, where $\partial_j [0,1]^N = \{\rho \in \mathbb{R}_+^N : \forall i, \rho_i \leq 1,\ \rho_j = 1\}$

$\longrightarrow$ we can saturate a class of users and use the recursion hypothesis.

FULL INTERFERENCE - ARBITRARY PACKET DURATION

For $\rho = (\rho_1, \ldots, \rho_N) \in \mathbb{R}_+^N$, define

$$\gamma_i(\rho) = \frac{P_i}{L \sum_j P_j + L_c C + E},$$

where

$$\begin{cases} P_i = \rho_i p_i \prod_{j \neq i} (1 - \rho_j p_j), \\ E = \prod_k (1 - \rho_k p_k), \\ C = 1 - E - \sum_j P_j. \end{cases}$$

The approximate stability region $\hat{\Lambda}_L^N$ is the set of points lying below one of the boundaries

$$\partial_j \hat{\Lambda}_L^N = \Big\{ \lambda : \exists \rho \in \partial_j [0,1]^N, \forall i, \lambda_i = \gamma_i(\rho) \Big\}.$$

Under the same assumptions than above

Theorem 2. $\hat{\Lambda}_L^N$ *is an asymptotically exact approximation of the stability region* Λ_L^N.

PARTIAL INTERACTION

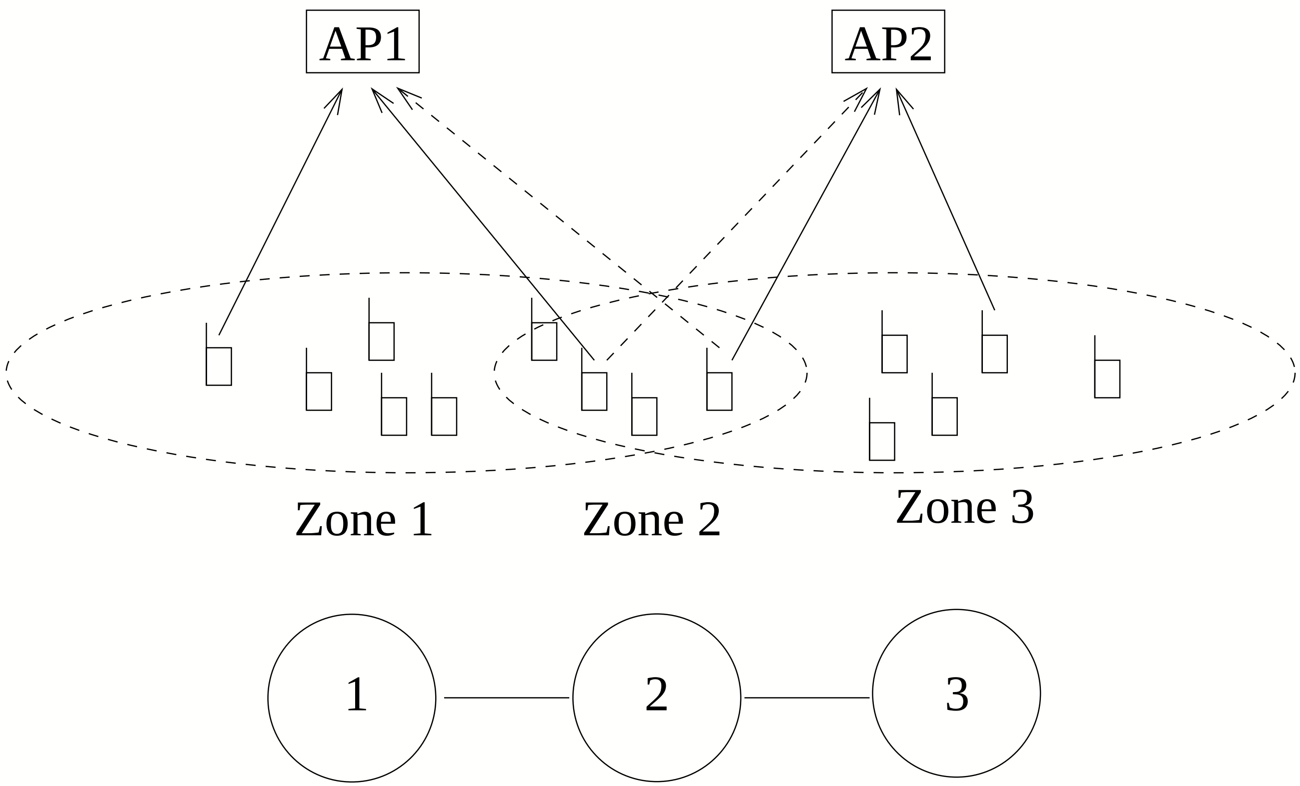

PARTIAL INTERFERENCE - UNIT PACKET DURATION

Assume $L = L_c = 1$. Let $\mathcal{I}_i^N$ be the set of users interfering with user i (including i).

$\hat{\Lambda}^N((\mathcal{I}_i^N)_{i=1}^N)$ is defined as the set of points lying below one of the boundaries

$$\partial_j \hat{\Lambda}^N((\mathcal{I}_i^N)_{i=1}^N) = \left\{\lambda : \exists \rho \in \partial_j [0,1]^N, \forall i, \lambda_i = \rho_i p_i \prod_{k \in \mathcal{I}_i^N \setminus \{i\}} (1 - \rho_k p_k)\right\}.$$

Assume that there is a finite number of classes of users, for all i

$$(\lambda_i^N, p_i^N, \mathcal{I}_i^N) = \left(\frac{\lambda_{v_i}}{N}, \frac{p_{v_i}}{N}, \{j : v_j \in \mathcal{I}_{v_i}\}\right)$$

Theorem 3. $\hat{\Lambda}^N((\mathcal{I}_i^N)_{i=1}^N)$ *is an asymptotically exact approximation of the stability region* $\Lambda^N((\mathcal{I}_i^N)_{i=1}^N)$*.*

SIMULATION

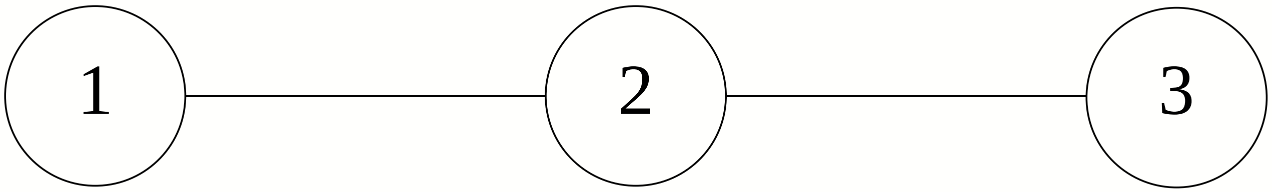

Maximum total rate compatible with stability

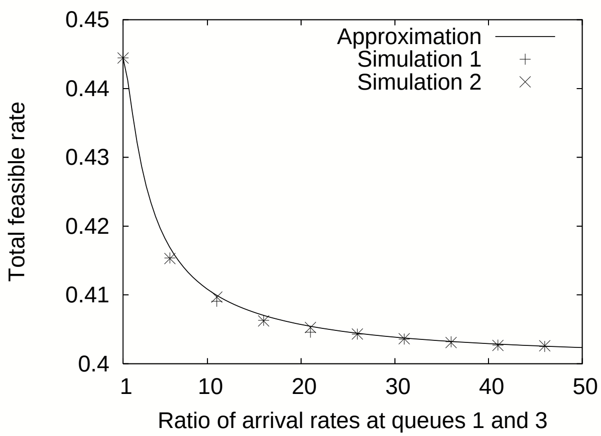

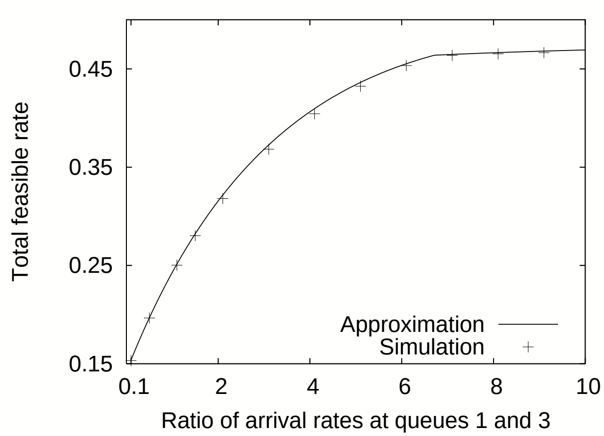

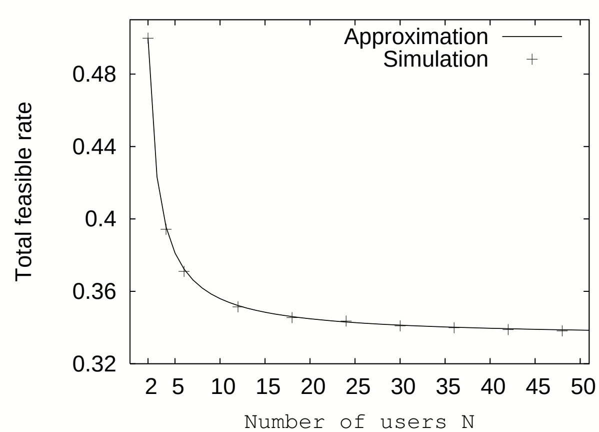

- *Left*: $N = 3$,

$(p_1, p_2, p_3) = (1/3, 1/3, 1/3),\ (\lambda_1, \lambda_2, \lambda_3) = \lambda(1, (1 + x^{-1})/2, x^{-1})$.

- *Center*: $N = 3$, $(p_1, p_2, p_3) = (0.6, 0.3, 0.1)$.

- *Right*: N increases, $p_i = 1/N$ and $\lambda^N = (\alpha_{v_1}, \cdots, \alpha_{v_N})\lambda/N$.

PARTIAL INTERFERENCE - ARBITRARY PACKET DURATION

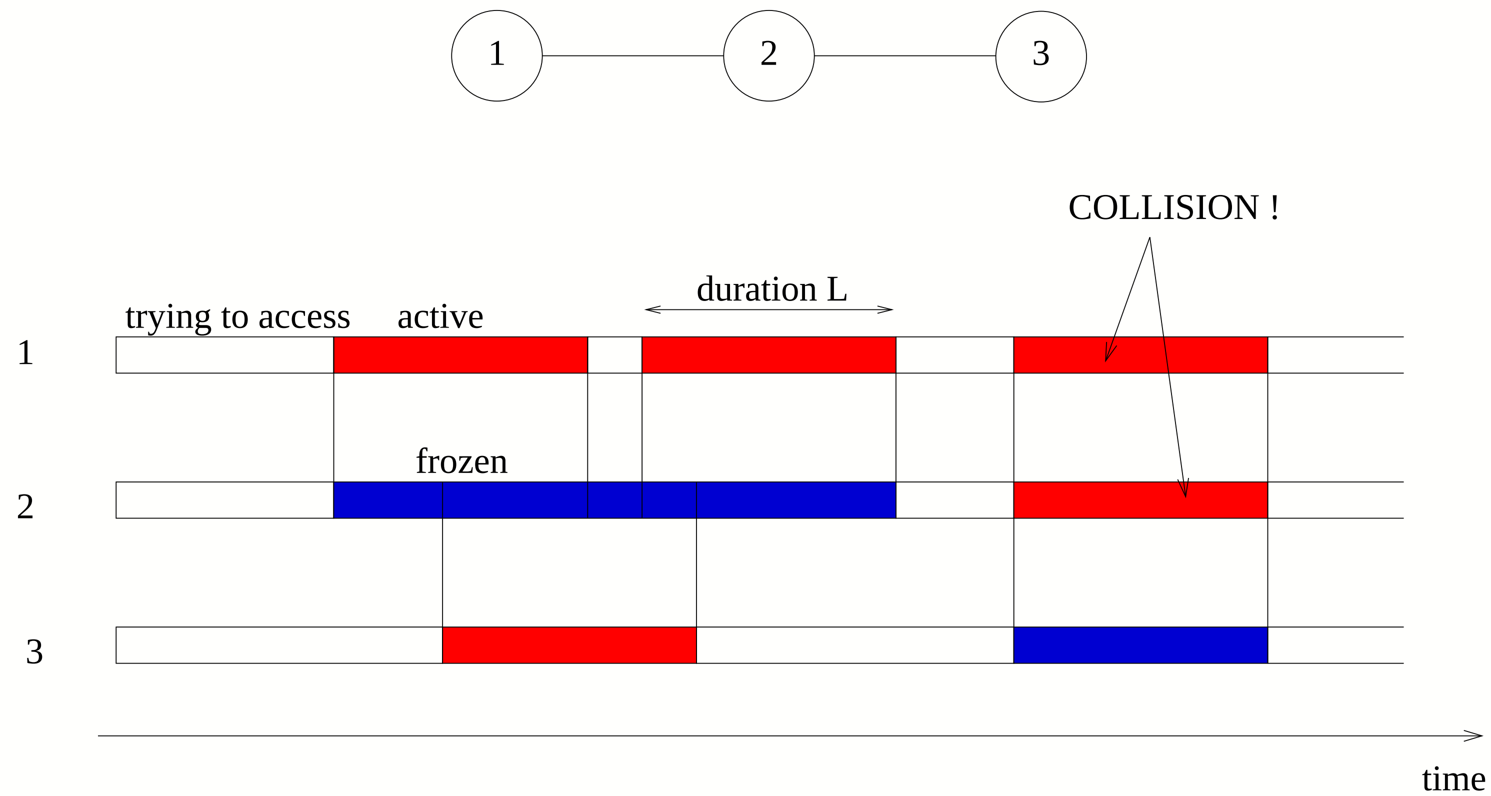

$\longrightarrow$ User 1 and user 3 do not see each other.

$\longrightarrow$ Users do not share the same history !

$\longrightarrow$ User 2 sees everything and is penalized ...

EVOLUTION OF A SINGLE USER

For a user i of class v. Define

$$Z_v^N(t) = \mathbf{1}_{\{\text{an interacting user is already active at time } t-\}}.$$

- If $Z_v^N(t) = 1$, then user i does no try to become active.
- If $Z_v^N(t) = 0$ and $B_i^N(t) \geq 1$ then user i becomes active with probability p_v/N.
- If user i becomes active, it will be a successful transmission with probability

$$(1 - \frac{p_v}{N})^{\beta_v^N \nu_v^+ N - 1} \prod_{w \in \mathcal{I}_v \setminus \{v\}} \left(\mathbf{1}_{Z_w^N(t)=1} + \mathbf{1}_{Z_w^N(t)=0} (1 - \frac{p_w}{N})^{\beta_w^N \nu_w^+ N} \right),$$

where ν_v^+ is the proportion of class-v users with non-empty buffers

EVOLUTION OF THE EMPIRICAL MEASURE

We define

$$\mu_t^N = \frac{1}{N}\sum_{i=1}^{N} \delta_{\{v_i, B_i^N(\lfloor Nt \rfloor)\}}$$

The probability transition of a single user depends on the environment variable $z = (z_v)_{v\in\mathcal{V}}$ with $z_v = \mathbf{1}_{\{\text{a user interacting with class } v \text{ is already active}\}}$.

$\Longrightarrow$ The environment variable and the transitions are mutually dependent !

The environment evolves at time scale $1/N$ and the empirical measure μ_t^N at scale 1.

$\longrightarrow$ in the limit, at time t the system should see the stationary mean π_{μ_t} of the environment.

$\longrightarrow$ in the limit, the environment and the transitions of each particle decouple !

ASYMPTOTIC STABILITY REGION

Let $\rho = (\rho_v)_{v \in \mathcal{V}}$, define π_ρ as the stationary distribution of the environment when the proportion of class-v users with non-empty buffers is ρ_v.

For $i \in v$, define $\gamma_i(\rho)$:

$$\gamma_i(\rho) = \sum_z \mathbf{1}_{\{z_v=0\}} \pi_\rho(z) \frac{\rho_i p_i}{1 - \rho_i p_i} \prod_{w \in \mathcal{I}_v} \left(\mathbf{1}_{\{z_w=1\}} + \mathbf{1}_{\{z_w=0\}} \prod_{j \in w} (1 - \rho_j p_j) \right).$$

Define $\hat{\Lambda}_L^N((\mathcal{I}_i)_{i=1}^N)$ as the set of points lying below one of the N boundaries

$$\partial_j \hat{\Lambda}_L^N((\mathcal{I}_i)_{i=1}^N) = \left\{ \lambda : \exists \rho \in \partial_j [0,1]^N, \forall i, \lambda_i = \gamma_i(\rho) \right\}.$$

Theorem 4. $\hat{\Lambda}_L^N((\mathcal{I}_i)_{i=1}^N)$ *is an asymptotically exact approximation of the stability region* $\Lambda_L^N((\mathcal{I}_i)_{i=1}^N)$*.*

A more sophisticated mean-field convergence theorem is needed.

SIMULATION

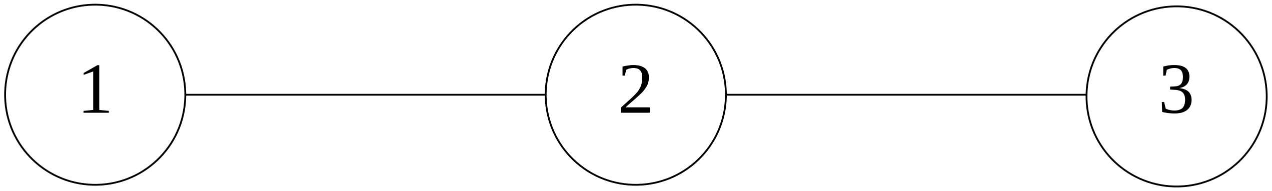

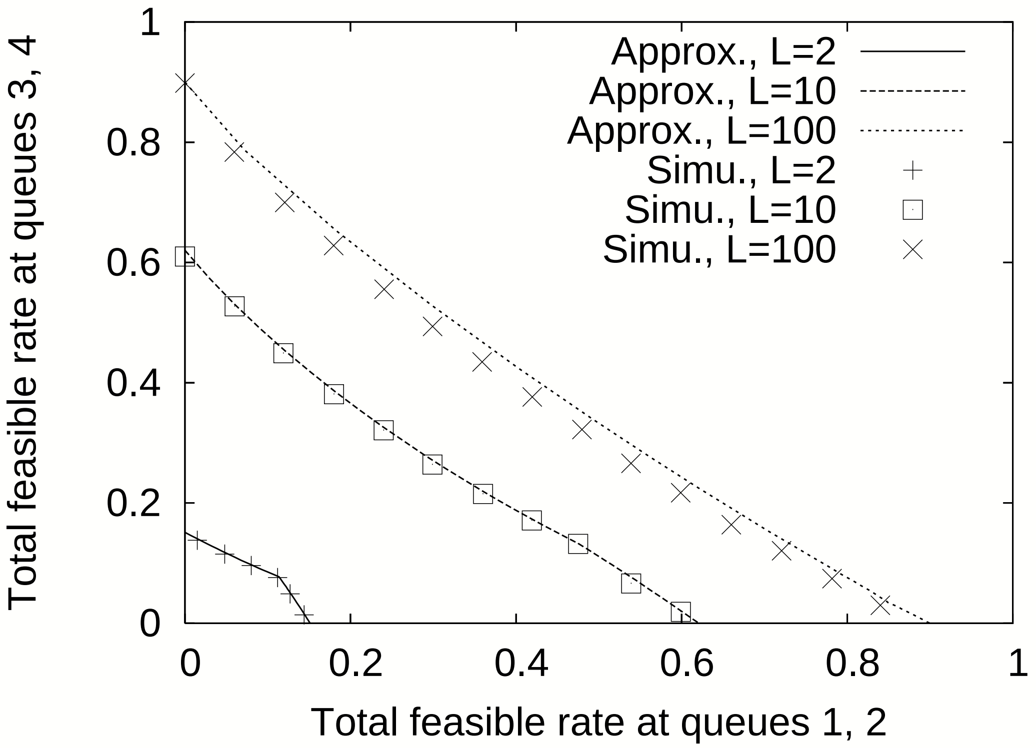

Stability region of the network with 2 users per class, $p_i = 0.1$ for all i, $L_c = L$.

FINAL COMMENTS

- Why is the approximation working so well for N small ? close to the stability limit, at least one class of queues is nearly saturated, thus these queues become nearly independent of the others.
- First application of a mean-field technique to compute the asymptotic stability region of a communication system.
- It works with any finite interaction graph.
- If users are saturated, it is also possible to analyse adaptive algorithms like exponential back-off with the mean-field technique *B-McDonald-Proutière (05-07)*. If users are not saturated, it is also mathematically feasible but the expressions become very complicated...
- The mean-field scaling does not explain everything, local time unfairness ?