P2P networks: a random infinite urn model

Philippe Robert and Florian Simatos

March 20, 2007, Alea 2007

Introduction

Emergence of P2P comes from the need to share large files efficiently

- Client/Server paradigm does not scale
- Neither with the size, nor with the load

P2P architectures:

- Rely on virtual and distributed networks (overlay network)
- A peer is a client and a server
  - Incentives for peers to upload
  - Files cut into chunks

Introduction

Examples of application:

- File sharing: Napster, Kazaa, BitTorrent, eDonkey, ...
- Distributed computation, telephony, gaming, ...

Problems triggered by P2P networks:

- Many CS problems:
  - Dynamic topology of the overlay network
  - Load balancing
  - ...
- Modeling P2P networks
  - As a graph...
  - ...or as a queueing system

P2P model

Description

Simulations

Random infinite urn model

Presentation

Chen Stein's inequality

Results

Description of the P2P model

The system aims at modeling the new paradigm:

- **A peer is a client and a server**
- **“Flash crowd” context**

Very simple model:

- **One file = one chunk**
- **Underlying graph is complete**
- **Peers do not leave the system**

Dynamic of the model

$t = 0$:

- N clients
- One server offers the file
- Each clients starts an exponential clock $\sim \exp(\lambda)$

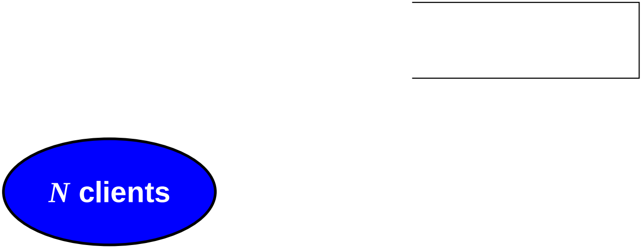

Dynamic of the model

First client arrives:

- **Time of arrival** $= \min_N \exp(\lambda) \sim \exp(N\lambda)$
- **Service requirements i.i.d.** $\sim \exp(\mu)$
- **Server serves only one client (FIFO)**
- **Server has infinite capacity**

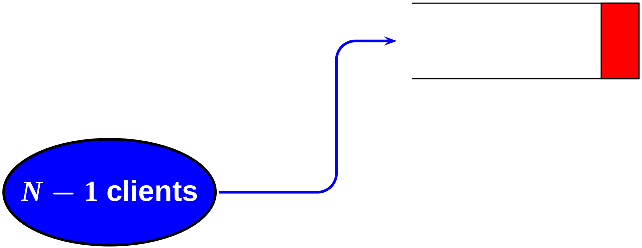

Dynamic of the model

$X(t)$ clients still in the pool

- **Next client** $\sim \exp\left(\lambda X(t)\right)$

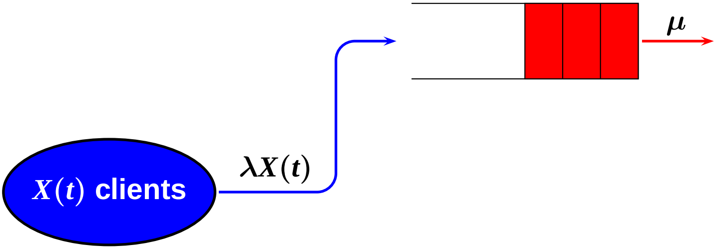

Dynamic of the model

$X(t)$ clients still in the pool

- Next client $\sim \exp\left(\lambda X(t)\right)$

$T_1 =$ time when the first client finishes its download

- $t < T_1$: clients arrive in the same server

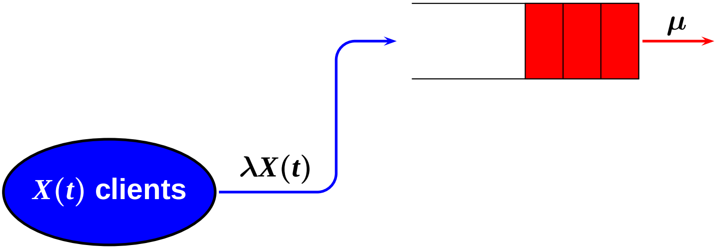

Dynamic of the model

T_1 = time when the first client finishes its download

- $t = T_1$: client finishes its download...
- ...and becomes a server

Dynamic of the model

For $T_1 < t < T_2$:

- Two possible servers: next server created at rate 2μ
- Policy on arrival: choose min? choose randomly? ...

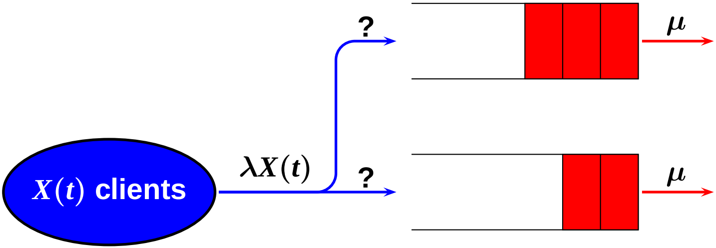

Dynamic of the model

At $t = T_2$:

- Creation of the third server
- Exit may be from one of the two servers

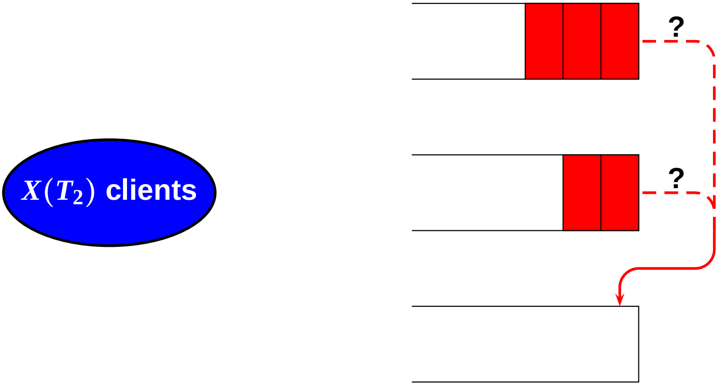

Dynamic of the model

$t > T_2$:

- ...

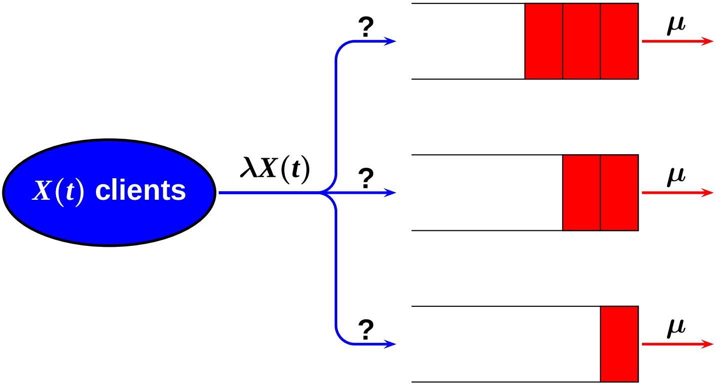

Discussion around simulations

Some simulations...

Discussion around simulations

Some simulations...

Two regimes

1. **No empty servers**
2. **More and more empty servers**

Discussion around simulations

Some simulations...

Two regimes

1. **No empty servers**
2. **More and more empty servers**

Why?

1. **Many arrivals before a new server is created**
2. **Many servers created before new arrival**

Discussion around simulations

From now on, we focus on the first regime

- **What is its duration?**

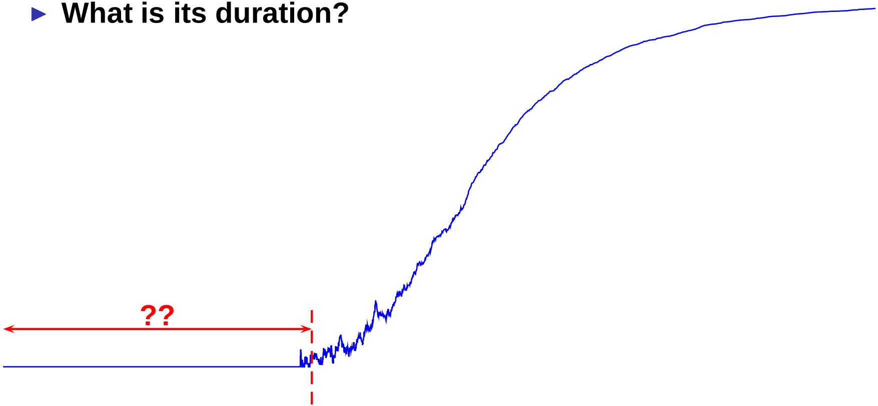

The random infinite urn model

Two stages:

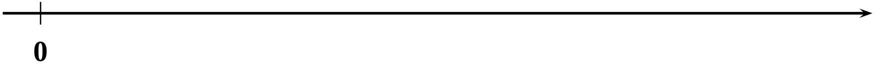

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp(i\mu)$ independently

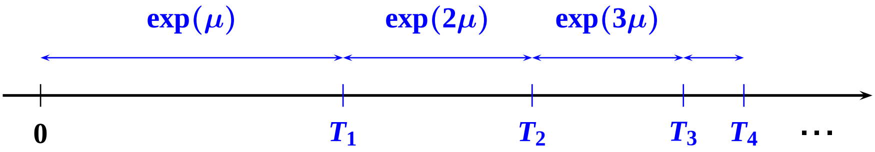

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i \geq 1}$ with $T_i - T_{i-1} \sim \exp(i\mu)$ independently
2. Throw N points $(E_\lambda^j)_{1 \leq j \leq N}$ with $E_\lambda^j \sim \exp(\lambda)$ i.i.d.

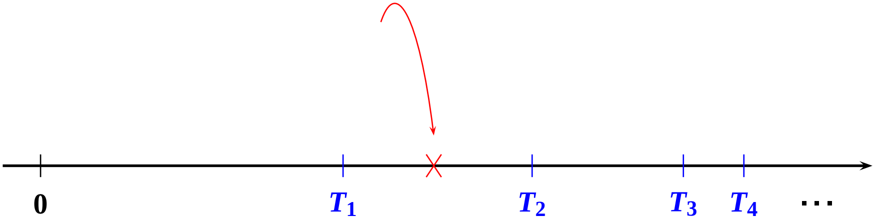

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp\left(i\mu\right)$ independently
2. Throw N points $(E_\lambda^j)_{1\leq j\leq N}$ with $E_\lambda^j \sim \exp(\lambda)$ i.i.d.

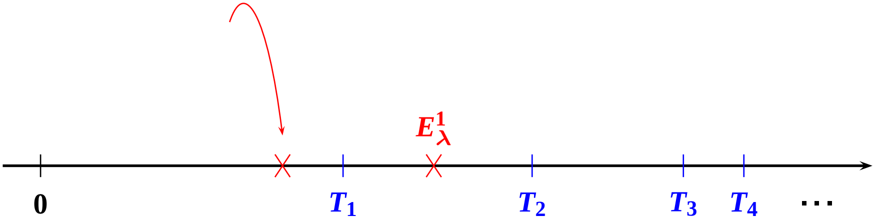

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp\left(i\mu\right)$ independently
2. Throw N points $(E_\lambda^j)_{1\leq j\leq N}$ with $E_\lambda^j \sim \exp(\lambda)$ i.i.d.

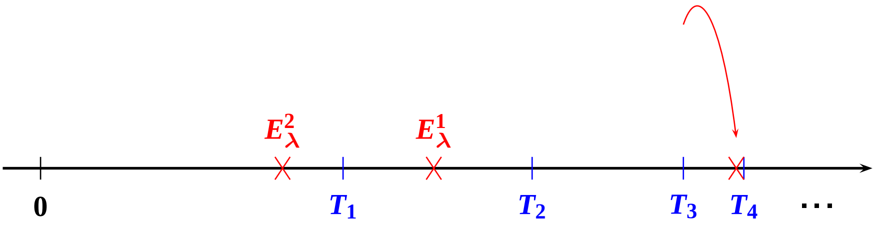

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp(i\mu)$ independently
2. Throw N points $(E_\lambda^j)_{1\leq j\leq N}$ with $E_\lambda^j \sim \exp(\lambda)$ i.i.d.

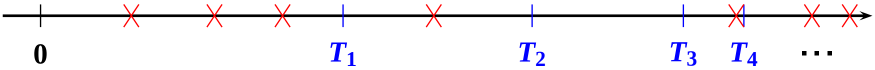

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp(i\mu)$ independently
2. Throw N points $(E^j_\lambda)_{1\leq j\leq N}$ with $E^j_\lambda \sim \exp(\lambda)$ i.i.d.

Back to the P2P model:

- $T_i =$ date of creation of the i-th server
- $E^j_\lambda =$ date of arrival of the client labelled j

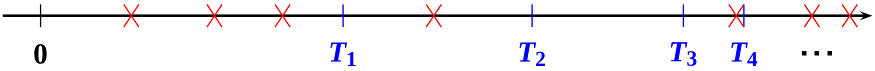

The random infinite urn model

Two stages:

1. Throw an infinite number of points $(T_i)_{i\geq 1}$ with $T_i - T_{i-1} \sim \exp\left(i\mu\right)$ independently
2. Throw N points $(E_\lambda^j)_{1\leq j\leq N}$ with $E_\lambda^j \sim \exp(\lambda)$ i.i.d.

Functionnal of interest:

$$\boxed{\nu = \text{index of the first empty interval}}$$

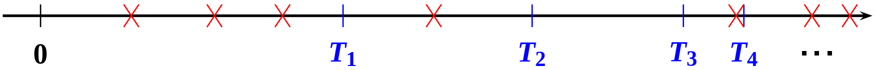

Remarks

The balls are not independent...

...but they are independent conditionnally on (T_i)

The number of empty urns W

Define W_N^x as the number of empty urns among the x first after throwing N balls

$$\nu > x \iff W_N^x = 0$$

The number of empty urns W

Define W_N^x as the number of empty urns among the x first after throwing N balls

$$\nu > x \iff W_N^x = 0 \quad \Longrightarrow \quad \boxed{\mathbb{P}\left(\nu/k(N) > x\right) = \mathbb{P}\left(W_N^{xk(N)} = 0\right)}$$

The number of empty urns W

Define W_N^x as the number of empty urns among the x first after throwing N balls

$$\nu > x \iff W_N^x = 0 \quad \Longrightarrow \quad \boxed{\mathbb{P}\left(\nu/k(N) > x\right) = \mathbb{P}\left(W_N^{xk(N)} = 0\right)}$$

Goal:

- **Find $k(N)$ such that $W_N^{k(N)}$ converges in distribution**
- **$k(N)$ is the order of magnitude of ν: $\nu \approx k(N)$**

Chen Stein's inequality

- W_N^x can be written as a sum of indicator functions
- Chen Stein's inequality applies, and gives for any n

$$\left|\left|\mathbb{P}(W_N^n \in \cdot) - \mathcal{P}_{\mathbb{E}W_N^n}(\cdot)\right|\right| \leq 1 - \frac{\mathbb{V}\mathrm{ar} W_N^n}{\mathbb{E}W_N^n}$$

- If for some $k(N)$, $\mathbb{E}W_N^{xk(N)} \sim \mathbb{V}\mathrm{ar} W_N^{xk(N)}$ and $\mathbb{E}W_N^{xk(N)} \to \zeta(x)$, then $W_N^{xk(N)}$ asymptotically has a Poisson distribution...
- ...and $\mathbb{P}\left(\nu/k(N) > x\right) \longrightarrow e^{-\zeta(x)}$

Preliminary results

Easy computations:

- $\mathbb{E}W_N^x = \sum_{i=1}^{x} \mathbb{E}\left[(1 - P_i)^N\right]$
- P_i is the random probability for a ball to fall in urn i conditionnally on (T_j)
- $P_i = \rho i^{-\rho-1} \times Z_i$, with $\rho = \lambda/\mu$
- Z_i converges in distribution

Does Chen Stein hold on simpler problems? 1/2

First try: deterministic urns

Theorem:

If $P_i \sim \alpha i^{-\rho-1}$, then $\nu/k(N) \to \beta$ in distribution, with $k(N) = (N/\ln N)^{1/(\rho+1)}$

Does Chen Stein hold on simpler problems? 1/2

First try: deterministic urns

Theorem:

If $P_i \sim \alpha i^{-\rho-1}$, then $\nu/k(N) \to \beta$ in distribution, with $k(N) = (N/\ln N)^{1/(\rho+1)}$

Closely related result:

If $P_i = 2^{-i-1}$, then $\mathbb{E}\nu \sim \log_2 N$ and $\mathbb{V}\mathrm{ar}\nu \sim \xi$

P. Flajolet and G. Nigel Martin, 1985

Does Chen Stein hold on simpler problems? 2/2

Second try: some randomness

Theorem:
If $P_i = \rho i^{-\rho-1} \widetilde{Z}_i$ with $\widetilde{Z}_i \to \exp(1)$, then $\nu/k(N)$ converges in distribution to a distribution of Weibull, with $k(N) = N^{1/(\rho+2)}$:

$$\mathbb{P}(\nu/k(N) > x) \longrightarrow e^{-(x/\alpha)^\beta}$$

Does Chen Stein hold on simpler problems? 2/2

Second try: some randomness

Theorem:
If $P_i = \rho i^{-\rho-1}\widetilde{Z}_i$ with $\widetilde{Z}_i \to \exp(1)$, then $\nu/k(N)$ converges in distribution to a distribution of Weibull, with $k(N) = N^{1/(\rho+2)}$:

$$\mathbb{P}(\nu/k(N) > x) \longrightarrow e^{-(x/\alpha)^\beta}$$

Remark:

- $N^{1/(\rho+2)} \ll (N/\ln N)^{1/(\rho+1)}$: randomness creates empty urns

Still…Chen Stein does NOT hold

Theorem:

For $\rho < 1/6$ and $k(N) = N^{1/(\rho+2)}$:

$$\begin{cases} \mathbb{E} W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)} \mathbb{E} X_\infty \\ \mathbb{V}\mathrm{ar} W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)} \mathbb{E} X_\infty - \left(\dfrac{x^{\rho+2}}{\rho(\rho+2)} \right)^2 \mathbb{V}\mathrm{ar} X_\infty \end{cases}$$

Still…Chen Stein does NOT hold

Theorem:

For $\rho < 1/6$ and $k(N) = N^{1/(\rho+2)}$:

Hard

$$\begin{cases} \mathbb{E} W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)} \mathbb{E} X_\infty \\ \mathbb{V}\mathrm{ar} W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)} \mathbb{E} X_\infty - \left(\dfrac{x^{\rho+2}}{\rho(\rho+2)}\right)^2 \mathbb{V}\mathrm{ar} X_\infty \end{cases}$$

Still...Chen Stein does NOT hold

Theorem:

For $\rho < 1/6$ and $k(N) = N^{1/(\rho+2)}$:

Hard
$$\begin{cases} \mathbb{E}W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)}\mathbb{E}X_\infty \\ \mathbb{V}\mathrm{ar}W_N^{xk(N)} \longrightarrow \dfrac{x^{\rho+2}}{\rho(\rho+2)}\mathbb{E}X_\infty - \left(\dfrac{x^{\rho+2}}{\rho(\rho+2)}\right)^2 \mathbb{V}\mathrm{ar}X_\infty \end{cases}$$

Remark:

- The condition $\rho < 1/6$ is partly an artefact, coming from integrability conditions
- But $\rho < 1$ is not an artefact: for $\rho \geq 1$, $\mathbb{E}X_\infty = +\infty$

Conclusion

Original P2P model far from being solved...

Interesting and original urn model arises:

- Urns both infinite and random
- Uncommon functional: first empty urn

Thank you