THE LAWS OF SUPER-SCALABILITY IN PEER TO PEER NETWORKS

F. Baccelli

UT–Austin and INRIA–ENS

Joint work with F. Mathieu and I. Norros

Asymptotics of Large-Scale Interacting Networks, BIRS 2013

STRUCTURE OF THE TALK

1. P2P Networking Motivations
2. Stochastic Model ∃!
3. Dimensional Analysis
4. Stochastic Analysis
5. Simulation
6. Scaling
7. Limitations
8. Extensions

Focus on the ongoing research part (in red) today.

PEER-TO-PEER CONTENT DISTRIBUTION

Content Distribution

– **Filesharing**

– **Streaming**

  * **OnDemand**
  * **Live**

Common Features

– **Lot of stress on the network**

– **P2P solutions**: large family of algorithms and implementations to cope with churn, load, heterogeneity ...

P2P STOCHASTIC NETWORK MODELING

State of the Art: Queuing Theory
[Yang and De Veciana 04], [Qiu and Srikant 04]

Three main types of nodes

- **Servers: provide, don't scale up**
- **Leechers: need, provide and scale**
- **Seeders: provide, scale**

Assumptions

- **Access-limited (physical/software)**
- **No network limitation**
- **Poisson arrivals**

This presentation: New models with network rate limitations

SPATIAL BIRTH AND DEATH STOCHASTIC MODEL

- **Peers live in a finite subset D of the Euclidean plane $\mathbb{R}^2$**
- **Dynamics: arrivals**
  - **Poisson rain: new peers arrive according to a Poisson process with time space intensity $\lambda dx dt$ on $D \times \mathbb{R}$**
- **Service requirement: each peer p is born with an individual service requirement $F_p > 0$ i.i.d. exponential with mean F.**

INTERACTION ?

- **Dynamics: service rate**
  - **Bit rate function: two peers at locations x and y serve each other at rate $f(||x-y||)$, where f is the bit rate function (BRF)**
  - **Service rate: the service rate of a peer at x in configuration ϕ is**

$$\mu(x,\phi) = \sum_{y \in \phi \setminus \{x\}} f(||x-y||).$$

  - **Service completion: for a system with state history $\{\phi_t\}_t$, a peer p born at point x_p at time t_p leaves at time**

$$\tau_p = \inf\{t > t_p : \int_{t_p}^{t} \mu(x_p, \phi_s) ds \geq F_p\}.$$

LARGE-SCALE?

- **Natural extensions to the case where D is**
  - **A torus (approximation of the whole plane);**
  - **The whole Euclidean plane;**
  - **General metric spaces (semantic spaces) e.g. $\mathbb{R}^d$.**

SPATIAL BIRTH AND DEATH PROCESS

- $\mathcal{N}(D)$: **the space of counting measures in** $(D, \mathcal{D})$
- **The state** ϕ_t **at time** t **is a Markov process living in the space** $\mathcal{N}(D)$**:**

  – **a peer has birth intensity** λ **at** x

  – **a peer located at** x **has death intensity** $\mu(x, \phi_t)/F$

- **(New?) class of spatial birth-and-death process with a death rate defined as a shot-noise of the configuration.**

EXISTENCE AND UNIQUENESS FINITE CASE

- **Lemma 1**
  If D is compact and f is bounded from below by a positive constant on some non-degenerate interval, then the Markov process $\{\phi_t\}_t$ is ergodic for any birth rate $\lambda > 0$.
- **Proof**
  - **stochastic domination: $M/M/\infty$ queue that is modified so that a lone customer cannot leave.**
  - **petite set technique à la Tweedie**
- **Remarks**
  - **non monotonic dynamical system**
  - **non reversible Markov process**
  - **non Gibbsian point process**

EXISTENCE AND UNIQUENESS INFINITE CASE

- **1. Definition over finite time?**
- **Lemma 2**
  If $D = \mathbb{R}^d$ **and**

$$\int_1^\infty f(r) r^{d-1} \, \mathrm{d}r < \infty$$

  then the spatial birth and death point process is uniquely defined on all finite time intervals $[t_0, t]$**.**
- **Proof: Random connection model definition of dynamics + existence and uniqueness of solution of a recursive equation.**

EXISTENCE AND UNIQUENESS INFINITE CASE (*continued*)

- Ψ_{t_0}: **space time arrival p.p. in** $[t_0, t]$
- **Random connection model definition of the SBD process:**

  – **exponential killing times** T_{pq}

  – **Bernoulli directions of killing** I_{pq}

between all pairs p, q **of points of the space time arrival p.p.** Ψ_{t_0}

- **Death times solution of an infinite recursive equation**

$$\delta_p = \inf \{T_{pq} : \; q \in \Psi_{t_0}, \; \delta_q \geq T_{pq}, \; I_{pq} = 1\} \, .$$

- **In the above setting, for all** $[t_0, t]$ **for all** p**, we give an algorithm determining whether** $\delta_p < t$ **or the value of** δ_p **otherw. in a.s. finite time.**

EXISTENCE AND UNIQUENESS INFINITE CASE (*continued*)

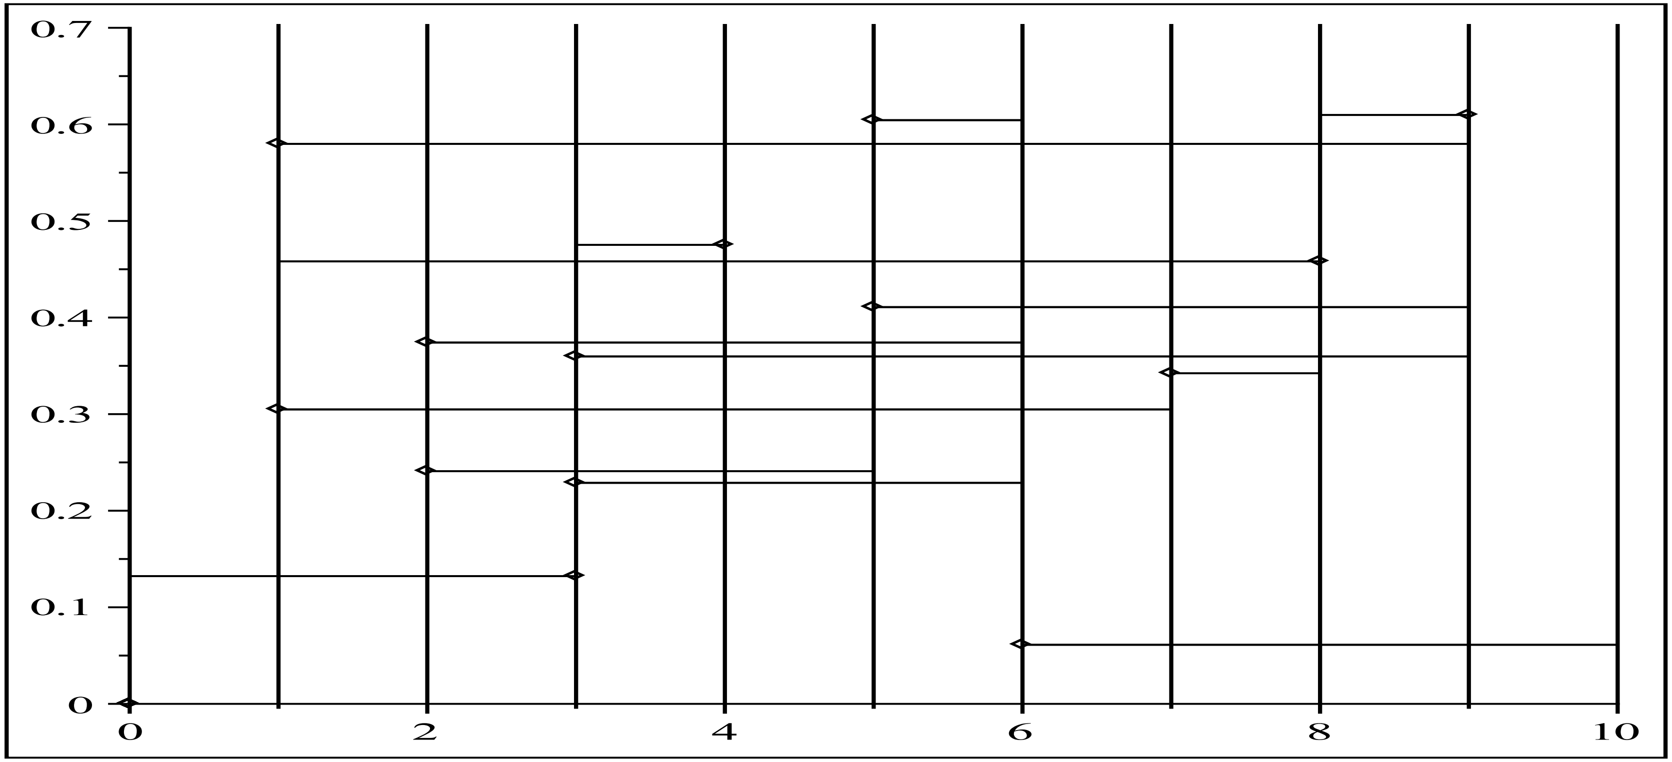

Randon Connection Model

EXISTENCE AND UNIQUENESS INFINITE CASE (*continued*)

- **2. Existence/uniqueness of stationary regimes**
- **Theorem 0**
  Under the assumptions of Lemma 2, there exists a unique stationary regime holding for all initial conditions made of a homogeneous Poisson point process of initial peers.
- **Proof: [quasi complete] based on coupling methods.**
  - **Analyze the effect on the process on $[t_0, \infty)$ of adding one peer (or a p.p. of peers) at t_0.**
  - **Coupling Algorithm building the two parallel universes with and without the additional peers.**

EXISTENCE AND UNIQUENESS INFINITE CASE (*continued*)

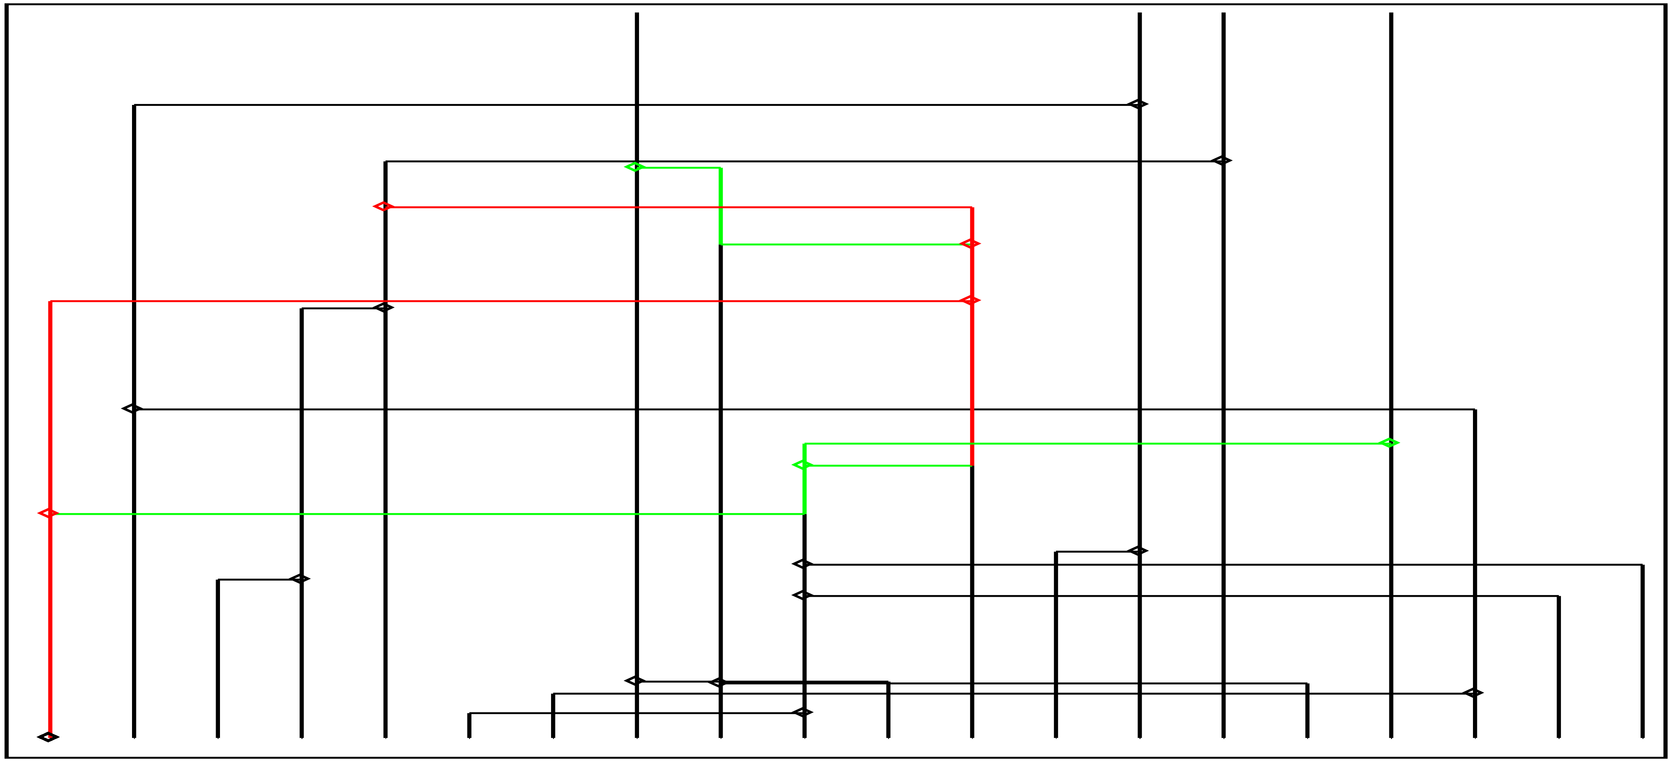

Two parallel universes: the familly of offsprings of the added point is a.s. finite.

EXAMPLES OF BRF: TCP

- **TCP model:** D **is the Euclidean plane** $I\!R^2$ **and**

$$f(r) = \frac{C}{r} 1_{r \leq R}.$$

- **Justification:**
  - **peers use TCP Reno**
  - **on the path between two peers, if the packet loss probability is** p **and the round trip time is** RTT**, then the rate obtained on this path is**

$$\frac{\eta}{\mathrm{RTT}\sqrt{p}}$$

    with $\eta =\sim 1.309$ **square root formula**
  - **the RTT is proportional to distance** r
  - **only peers at distance less than** R **are retained.**

EXAMPLES OF BRF: TCP (*continued*)

- **Variants**
  - **Affine RTT model:** $\mathrm{RTT} = ar + b$, **where** a **accounts for propagation delays in the Internet path and** b **for the mean access latency:**

$$f(r) = \frac{C}{r+q} 1_{r \leq R}$$

  - **Additional overhead cost:** c **bits per second:**

$$f(r) = \left(\frac{C}{r+q} - c \right)^{+} 1_{r \leq R}$$

  - **Upload (or Download) rate limitations:**

$$f(r) = \min \left(U, \left(\frac{C}{r+q} - c \right)^{+} \right) 1_{r \leq R}$$

with U **the individual rate limitation**

EXAMPLES OF BRF: UDP

- **UDP assumptions:**
  - D **is the Euclidean plane** $I\!R^2$
  - **only peers within distance** R **are retained**
  - **peers use UDP with prescribed rate** C **regardless of distance**

$$f(r) = C1_{r \leq R}.$$

EXAMPLES OF BRF: WIRELESS SNR

- **SNR model**: the rate between a transmitter and its receiver at distance r is

$$f(r) = \frac{1}{2} \log \left(1 + \frac{C}{r^\alpha} \right) 1_{r<R}$$

with

– $\alpha > 2$ the path loss exponent

– C the signal to noise power ratio at distance 1

– R the transmission range

- **Requirement**: all point-to-point channels are mutually orthogonal

DEFAULT MODEL

- **Default option model throughout the talk:**
  - **– D is the Euclidean plane or a large torus**
  - **– TCP Bit Rate Function:**

$$f(r) = \frac{C}{r} 1_{r<R}$$

- **+ comments on the other Bit Rate Functions**

DIMENSIONAL ANALYSIS

- **4 basic parameters:**
  - R **in meters (m),**
  - F **in bits,**
  - λ **in m^{-2} per second (s)**
  - C **in bit·m·s^{-1}.**
- **π-Theorem**
  In the TCP case, all system properties only depend on the parameter

$$\rho = \frac{\lambda F R^3}{C}.$$

- **Extension for more general f s.t. $\int f(r) r dr < \infty$.**

DIMENSIONAL ANALYSIS (*continued*)

- **Sketch of proof**
  - **choose R as a new distance unit, then**
    - **the arrival intensity becomes $l = \lambda R^2$**
    - **the download constant becomes $c = C/R$**
  - **now define F as an information unit, then**
    - **the download speed constant becomes $c = C/(RF)$**
  - **take a time unit such that the download speed constant is 1, then**
    - **all parameters are equal to 1**
    - **the arrival rate becomes $l = \frac{\lambda F R^3}{C}$**

DIMENSIONAL ANALYSIS (*continued*)

- **Terminology: Three cases**
  - $\rho \gg 1$ **is called fluid**
  - $\rho \ll 1$ **is called hard core**
  - ρ **inbetween is called intermediate**

NOTATION

- **In the steady state regime of the P2P dynamics:**
  - β_o **the density of the peer point process**
  - μ_o **the mean rate of a typical peer**
  - W_o **the mean latency of a typical peer**
  - N_o **the mean number of peers in a ball of radius** R **around a typical peer**

f-REPULSION

- **Theorem 1**
  For all BRF f, in the stationary regime,

$$\mathbb{E}[\sum_{x_i \in \phi} f(||x_i||)] \geq \mathbb{E}_0[\sum_{x_i \in \phi \setminus 0} f(||x_i||)],$$

  where $I\!P_0$ is the Palm probability w.r.t. Φ.

- **Proof:** **rate conservation principle + Papangelou theorem for point processes with stochastic intensity**

SKETCH OF PROOF - TORUS

- Φ_t: **state of the SBD at time t.**
- $\mathbb{A}_t$: **total rate**

$$\mathbb{A}_t = \sum_{X \in \Phi_t} A_t(X),$$

with, for all $X \in \Phi_t$:

$$A_t(X) = \sum_{Y \in \Phi_t, Y \neq X} f(||X - Y||)$$

SKETCH OF PROOF - TORUS (*continued*)

- **Rate conservation principle applied to $\mathbb{A}_t$:**
  - $\mathbb{E}^{\uparrow}$**: (time) Palm probability of the SBD at birth epochs**
  - $\mathbb{E}^{\downarrow}$ **at death epochs.**

$$r^{\uparrow}\mathbb{E}^{+}(\mathcal{I}) = r^{\downarrow}\mathbb{E}^{\downarrow}(|\mathcal{D}|)$$

with

- $\mathcal{I} = \mathbb{A}_{0+} - \mathbb{A}_0$ **the total rate increase,** $r^{\uparrow}$ **the inc. intensity**
- $\mathcal{D} = \mathbb{A}_{0+} - \mathbb{A}_0$ **the total rate decrease,** $r^{\downarrow}$ **the dec. intensity**

SKETCH OF PROOF - TORUS (*continued*)

- **Since $r^{\uparrow} = r^{\downarrow}$,**

$$\mathbb{E}^{\uparrow}(\mathcal{I}) = \mathbb{E}^{\downarrow}(\mathcal{D}).$$

- **From PASTA**

$$\mathbb{E}^{\uparrow}(\mathcal{I}) = 2\mathbb{E}(n_0)\frac{a}{|D|}.$$

with n_0 the total population and

$$a = \int_T f(||x||)m(\,\mathrm{d}x).$$

with T the torus of area $|D|$.

SKETCH OF PROOF - TORUS (*continued*)

- **The (total) death point process admits a stochastic intensity w.r.t. the filtration $\mathcal{F}_t = \sigma(\Phi_s, s \leq t)$ equal to $\mathbb{A}_t$.**
- **From Papangelou's theorem** $\frac{d\mathbb{P}^{\downarrow}}{d\mathbb{P}} |_{\mathcal{F}_{0-}} = \frac{\mathbb{A}_0}{\mathbb{E}(\mathbb{A}_0)}$.
- **Since the decrease (in state Φ_{0-}) is of magnitude $A_0(X)$ (w.r.t. Φ_{0-}) with probability $\frac{A_0(X)}{\mathbb{A}_0}$ (w.r.t. Φ_{0-}),**

$$\mathbb{E}^{\downarrow}(|\mathcal{D}|) = 2\mathbb{E}\left(\frac{\mathbb{A}_0}{\mathbb{E}(\mathbb{A}_0)} \sum_{X \in \Phi_0} \frac{A_0(X)}{\mathbb{A}_0} A_0(X)\right) = 2\frac{\mathbb{E}\left(\sum_{X\in\Phi_0} (A_0(X))^2\right)}{\mathbb{E}\left(\sum_{X\in\Phi_0} A_0(X)\right)}$$

$$= 2\frac{\mathbb{E}_0\left((A_0(0))^2\right)}{\mathbb{E}_0\left(A_0(0)\right)}$$

SKETCH OF PROOF - TORUS (*continued*)

- **Rate conservation principle for total rate:**

$$\mathbb{E}(n_0)\frac{a}{|D|} = \frac{\mathbb{E}_0\left((A_0(0))^2\right)}{\mathbb{E}_0\left(A_0(0)\right)}.$$

- **Using the fact that**

$$\mathbb{E}_0\left((A_0(0))^2\right) \geq \mathbb{E}_0\left(A_0(0)\right)^2,$$

we get

$$\mathbb{E}(n_0)\frac{a}{|D|} \geq \mathbb{E}_0\left(A_0(0)\right).$$

FLUID MODEL IN WHOLE PLANE

- **Fluid heuristic: obtained when approximating the Palm expectation of the rate, namely the mean rate obtained by a typical user, by the mean rate at a typical location:**

$$\mu_f = \beta_f 2\pi \int_{r=0}^{R} (C/r) r dr = \beta_f 2\pi C R.$$

with β_f the density of peers in this heuristic.

FLUID MODEL AS AN ASYMPTOTIC

- **Theorem 2**

  When ρ tends to infinity:

  – **The fluid heuristic is asymptotically tight:**

  $\beta_o \to \beta_f$, $W_o \to W_f$, $\mu_o \to \mu_f \cdots$

  – **The law of the latency of a typical peer converges weakly to an exponential random variable of parameter $W_f = \frac{F}{\mu_f}$**

- **Proof: fluid limit techniques extended to spatial processes**

FLUID MODEL AS AN ASYMPTOTIC (*continued*)

- **In this heuristic/limit**

$$\begin{aligned} \beta_f &= \sqrt{\frac{\lambda F}{2\pi C R}}, \\ \mu_f &= \sqrt{\lambda F 2\pi C R}, \\ W_f &= \sqrt{\frac{F}{\lambda 2\pi C R}}, \\ N_f &= \sqrt{\frac{\pi}{2}}\sqrt{\frac{\lambda F R^3}{C}} = \sqrt{\frac{\pi}{2}}\sqrt{\rho}. \end{aligned}$$

- **Proof:** $W_f = F/\mu_f$ **and** $\beta_f = \lambda W_f$ **(Little's law) and** $\mu_f = \beta_f 2\pi C R$**. Hence**

$$\beta_f \mu_f = \lambda F \quad \Leftrightarrow \quad \beta_f \beta_f 2\pi C R = \lambda F$$

COMMENTS ON FLUID ASYMPTOTIC

- ρ **is large when**
  - **either the arrival intensity, or the file size, or the range are large**
  - **or if the download speed constant C is small**
- **the time scale of a peer is $W_f = \sqrt{F/(\lambda 2\pi CR)}$.**
  If two peers are at a distance r_0 such that

$$\frac{F}{\frac{C}{r_0}} \ll W_f = \sqrt{\frac{F}{\lambda 2\pi CR}} \Leftrightarrow r_0 \ll \sqrt{\frac{C}{2\pi\lambda FR}} = \frac{R}{\sqrt{2\pi\rho}},$$

  then there is little chance to see these too peers in the steady state: hard exclusion below that scale.
- r_0 **tends to 0 in configurations where ρ tends to infinity and R is fixed**

FLUID REGIME AS A BOUND

- **In the TCP case, Theorem 1 is equivalent to saying that**

$$\beta_o 2\pi CR \geq \mu_o.$$

- **It follows from the relations** $W_o \geq F/\mu_o$ **and** $\beta_o = \lambda W_o$ **that**

$$\beta_o \geq \lambda \frac{F}{\beta_o 2\pi CR}$$

- **Corollary**

$$\beta_o^2 \geq \sqrt{\frac{\lambda F}{2\pi CR}} = \beta_f \quad \textbf{and} \quad W_o \geq W_f$$

HARD CORE REGIME

- **A stationary point process is hard–core for balls of radius R if there are no other points in a ball of radius R centered on any point.**
- **Conjecture When ρ tends to 0,**
  - **the stationary peer point process tends to a hard–core point process for balls of radius R with intensity β_h and latency W_h:**

$$\beta_h = \frac{1}{\pi R^2}, \quad W_h = \frac{1}{\lambda \pi R^2}.$$

  - **the cdf of the latency converges weakly to**

$$1 - \frac{e^{-\frac{t}{2W_h}}}{2}, \quad t > 0.$$

HARD CORE REGIME (*continued*)

Rationale

$$N_f \ll 1$$

$$\Downarrow$$

$$\sqrt{\frac{\lambda F R^3}{C}} \ll 1$$

$$\Downarrow$$

$$\sqrt{\frac{\lambda R C F^2 R^2}{F C^2}} \ll 1$$

$$\Downarrow$$

$$\frac{RF}{C} \ll \sqrt{\frac{F}{2\pi\lambda RC}} = W_f \leq W_o.$$

The latency of two peers within range is negligible w.r.t. the mean latency

GLOBAL HEURISTIC

- **Global Heuristic**:
  - **considers $\hat{\mu}$, the unique solution of**

$$\hat{\mu}^2 = \mu_f^2 \left(1 - \frac{C}{\hat{\mu}R} \ln\left(1 + \frac{\hat{\mu}R}{C}\right)\right),$$

  - **then defines**

$$\hat{\beta} = \lambda F/\hat{\mu}, \qquad \hat{W}_h = F/\hat{\mu}.$$

GLOBAL HEURISTIC (*continued*)

- **Factorization of the factorial moment measure of order 3**
- **Balance equation for the second order factorial moment density, which reads**

$$2\beta_o\lambda = 2m_{[2]}(x,y)\frac{C}{F}\frac{1_{||x-y||\le R}}{||x-y||} + \frac{C}{F}\int_D m_{[3]}(x,y,z)\left(\frac{1_{||x-z||\le R}}{||x-z||}+\frac{1_{||y-z||\le R}}{||y-z||}\right)dz,$$

for all x and y.

- **Approximations:**

$$m_{[3]}(x,y,z) \approx \frac{m_{[2]}(x,y)m_{[2]}(x,z)}{\beta_o}$$

$$m_{[3]}(x,y,z) \approx \frac{m_{[2]}(x,y)m_{[2]}(y,z)}{\beta_o}.$$

GLOBAL HEURISTIC (*continued*)

- **Then**

$$\beta_o \lambda \approx m_{[2]}(x,y)\frac{C}{F}\frac{1_{||x-y||\le R}}{||x-y||} + m_{[2]}(x,y)\frac{C}{F}\frac{1}{2}\int_D \frac{1_{||x-z||\le R}}{||x-z||}\frac{m_{[2]}(x,z)}{\beta_o}dz + m_{[2]}(x,y)\frac{C}{F}\frac{1}{2}\int_D \frac{1_{||y-z||\le R}}{||y-z||}\frac{m_{[2]}(y,z)}{\beta_o}dz,$$

that is

$$m_{[2]}(x,y) \approx \lambda F \frac{\beta_o}{\frac{C 1_{||x-y||\le R}}{||x-y||} + \mu_o}.$$

with $\mu_o =: C \int_{B(0,R)} \frac{m_{[2]}(0,z)}{\beta_o}\frac{1}{||z||}dz$**.**

GLOBAL HEURISTIC (*continued*)

So

$$\mu_o \approx \lambda F 2\pi C \int_0^R \frac{1}{\mu_o + \frac{C}{r}} dr$$
$$= \lambda F 2\pi C \left(\frac{R}{\mu_o} - \frac{C}{\mu_o^2} \ln(1 + \frac{\mu_o R}{C}) \right).$$

and

$$\hat{\mu}^2 = \mu_f^2 \left(1 - \frac{C}{\hat{\mu} R} \ln \left(1 + \frac{\hat{\mu} R}{C} \right) \right),$$

COMMENTS ON GLOBAL HEURISTIC

$$\hat{\mu}^2 = \mu_f^2 \left(1 - \frac{C}{\hat{\mu}R} \ln \left(1 + \frac{\hat{\mu}R}{C}\right)\right),$$

– **When $\hat{\mu}R/C$ tends to ∞, then it follows that $\hat{\mu} \sim \mu_f$, which is in line with Theorem 2.**

– **When $\hat{\mu}R/C$ tends to 0, then, expanding the log substantiates Conjecture 3.**

SIMULATION

- **Fix 3 independent parameters and use the 4-rth one to run through all possible scenarios.**
- **The two first fixed parameters are $R = .1$ and $C = 1$.**
- **Set W_f to 100. This implies that for all simulations, the fluid model will predict the same mean latency.**
- **Then, we use N_f as the variable parameter: We use N_f instead of ρ as main dimensionless parameter**
- **The remaining input parameters of the system are then completely defined:**

$$\lambda = \frac{N_f}{\pi R^2 W_f}, \qquad F = \frac{2 N_f C W_f}{R}$$

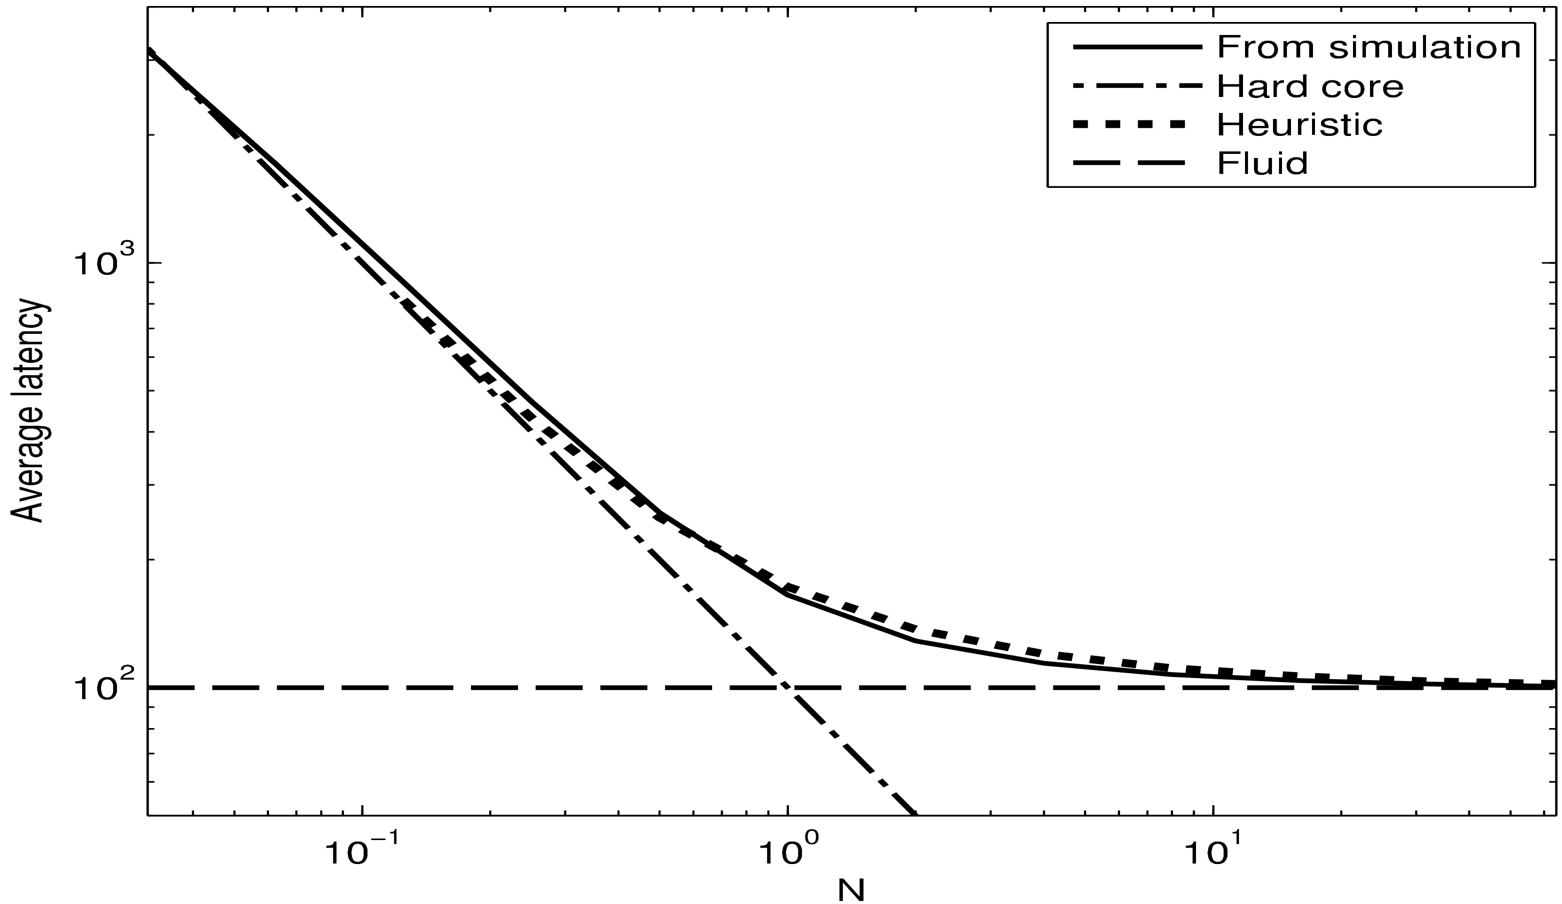

Latency in function of N_f.

SUPER-SCALABILITY

- **Dimensional analysis tells us that**

$$W_o(\lambda, F, C, R) = M\left(\sqrt{\frac{\pi\lambda F R^3}{2C}}\right) W_f(\lambda, F, C, R)$$
$$= M\left(\sqrt{\frac{\pi\lambda F R^3}{2C}}\right)\sqrt{\frac{F}{\lambda 2\pi C R}}$$

where M only depends on $N_f = \sqrt{\frac{\pi\lambda F R^3}{2C}}$ and is decreasing.

- **λ and R are both win-win parameters. As they increase, both terms in the RHS decrease and the mean latency hence tends towards 0, while the behavior of the system becomes more and more fluid.**
- **Super Scalability !**

SCALABILITY & SUPER SCALABILITY

Single Server M/M/1 Queue **Does not scale**	Infinite Server M/M/∞ Queue **Scales**	Network Limited P2P Spatial B & D P2P **Super Scales**
$W = \frac{1}{\mu - \lambda}, \lambda < \mu$	$W = \frac{1}{\mu}$	$W = \frac{m(\lambda)}{\sqrt{\lambda}}, m(\cdot) \downarrow$

SOME EXTENSIONS

- **Rate Limitations**
  - **Adapting** R
  - **Upload**
- **Seeders**

ADAPTING THE PEERING RADIUS

- **Mean Constant Number of Nearest Peers: take as neighbors the peers in a ball with a radius R such that the mean number of other peers in the ball is L i.e. $\pi R^2 \beta_o = L$, where β_o is the (unknown) steady state intensity of the point process ϕ_t. Then**

$$f(r) = \frac{C}{r} 1_{r \leq R}, \quad R = \sqrt{\frac{L}{\pi \beta_o}}$$

- **General Case**

$$f(r) = \frac{C}{r} 1_{r \leq R}, \quad R = \kappa \beta_o^{-\alpha}$$

- **(DA) All system properties only depend on the parameter**

$$\rho = \frac{\lambda F}{C} \kappa^{\frac{3}{1-2\alpha}}.$$

ADAPTING THE PEERING RADIUS (*continued*)

- **Fluid: in the general case $\mu_f = 2\pi C\kappa\beta_f^{1-\alpha}$, so that**

$$\begin{aligned}\beta_f &= \left(\frac{\lambda F}{2\pi C\kappa}\right)^{\frac{1}{2-\alpha}}\\ W_f &= \lambda^{-\frac{1-\alpha}{2-\alpha}} F^{\frac{1}{2-\alpha}} (2\pi C\kappa)^{-\frac{1}{2-\alpha}}\\ \mu_f &= (2\pi C\kappa)^{\frac{1}{2-\alpha}} (\lambda F)^{\frac{1-\alpha}{2-\alpha}}.\end{aligned}$$

This is obtained when choosing a radius of the form

$$R = \kappa\left(\frac{\lambda F}{2\pi C\kappa}\right)^{\frac{\alpha}{\alpha-2}}.$$

- **For instance in the constant number of nearest peers case**

$$\beta_f = \frac{\left(\frac{\lambda F}{2C}\right)^{\frac{2}{3}}}{(\pi L)^{\frac{1}{3}}},\ \mu_f = (2C)^{\frac{2}{3}}(\lambda F\pi L)^{\frac{1}{3}},\ W_f = \frac{\left(\frac{F}{2C}\right)^{\frac{2}{3}}}{(\lambda\pi L)^{\frac{1}{3}}}.$$

ASYMPTOTIC DESIGN

- **General α case: $R = \kappa\beta^{-\alpha}$.**
- **think of all parameters fixed and let λ tend to infinity.**
  - $d = \frac{1}{2-\alpha}$ **the density exponent: β is of the order λ^d**
  - $l = \frac{\alpha-1}{2-\alpha}$ **the latency exponent: W is of the order λ^l**
  - $r = \alpha/(\alpha - 2)$ **the radius exponent: r is of the order λ^r**
- **2 regimes, both compatible with fluid:**
  - **For $\alpha > 2$, we get a peer density and a latency which both tend to 0 when λ tends to ∞: Heaven's–flash**
  - **For $\alpha < \frac{1}{2}$, we get a peer density that tends to infinity and a latency which tends to zero when λ tends to ∞: swarm–flash**

UPLOAD AND NETWORK LIMITATIONS

- U: average upload capacity of a peer;
- The average rate in the fluid limit should be such that

$$\mu_f = \sqrt{\lambda F 2\pi C R} \leq U.$$

- **A natural dimensioning rule**: choose

$$R = \frac{U^2}{\lambda F 2\pi C}$$

in order to use all the available upload capacity and not more.

SEEDERS

- **When a leecher has obtained all its file, rather than leaving, it becomes a seeder and remains such for a duration** T_S
- **Fluid limit with seeders**

$$\mu_f = (\beta_f + \lambda T_S) 2\pi C R.$$

Using $F = W_f \mu_f$ **and** $\beta_f \mu_f = \lambda F$**, we get**

$$W_f^2 + W_f T_S = W_{f_0}^2, \text{ with } W_{f_0} = \sqrt{\frac{F}{\lambda 2\pi C R}}.$$

The positive solution of this equation is

$$W_f = \sqrt{W_{f_0}^2 + \left(\frac{T_S}{2}\right)^2} - \frac{T_S}{2}.$$

CONCLUSION

- A **new, non Gibbsian point process** model with many open challenges
  - Hard core regime
  - Intermediate regime
- Design implications
  - Laws of Super-Scalability for future P2P
  - First understanding of the assumptions for these laws to hold
- Ongoing work
  - Chunk level model → INFOCOM 13
  - Math paper in preparation
  - http://hal.inria.fr/inria-00615523/en