Spatial Birth and Death Processes with Poisson Birth and Shot-Noise Deaths

F. Baccelli

Weierstrass Institute, Berlin

November 2-4

European Research Council

Established by the European Commission

SPATIAL POINT PROCESSES WITH POISSON BIRTHS AND DEATHS BY RANDOM CONNECTION

Geometry, Point Processes, Dynamics & Content Delivery Networking

I. Mathematical Analysis

1. Stochastic Model
2. Stochastic Stability
3. Repulsion
4. Birth Death Equations

II. TCP Example

1. Dimensional Analysis
2. Limiting Regimes
3. Simulation
4. Super-scalability

Joint work with F. Mathieu and I. Norros, Queueing Systems, 2017

PEER-TO-PEER CONTENT DISTRIBUTION

Content distribution is mainly:

- Filesharing
- Streaming
    - OnDemand (YouTube, Netflix)
    - Live (Sport events)

Things in common:

- Bandwidth is a key parameter
- Lot of stress for the network
- P2P solutions exist

- **P2P Principles**
  - **Peers join and leave aiming each at downloading a very large file**
  - **The file is cut in smaller chunks**
  - **Peers exchange chunks on a Tit for Tat basis**
  - **The swarm of peers solves the initialization problem**
  - **The bit rate between two peers is determined by their distance**

I. SPATIAL BIRTH AND DEATH STOCHASTIC MODEL

- **Nodes live in a subset D of the Euclidean space $\mathbb{R}^d$ or on a torus**
- **Dynamics: arrivals**
  - **Poisson rain: new nodes arrive according to a Poisson process with time space intensity $\lambda dxdt$ on $D \times \mathbb{R}$**
- **Service requirement: each node 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 nodes at locations x and y serve each other at rate $\mathbf{f}(||\mathbf{x}-\mathbf{y}||) \geq \mathbf{0}$**
  - **Service rate: the service rate of a node at x in configuration ϕ is**

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

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

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

SPATIAL BIRTH AND DEATH PROCESS

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

  – **a node has birth intensity** λ **at** $\mathbf{x}$

  – **a node located at** x **has death intensity** $\mu(\mathbf{x}, \phi_t)/\mathbf{F}$
- **Spatial birth-and-death process with a death rate defined as a shot-noise of the configuration**

CONSTRUCTION: FINITE CASE

- **Lemma**
  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: $M/M/p2p$ queue that is modified so that a lone customer cannot leave**
  - **Harris-recurrence techniques**

- **Remarks: in general**
  - **non monotonic dynamical system**

THE $M/M/p2p$ SYSTEM

Consider the following queueing system:

- **customers ('peers') arrive according to a Poisson process with parameter** λ
- **every peer has an independent** $\mathrm{Exp}(1)$ **distributed service requirement, and each peer serves every other peer at rate** μ**; that is, if there are** n **peers in the system, each of them has stochastic intensity** $(\mathrm{n}-1)\mu$ **to leave**

After arrival of first peer, the system is never empty

THE $M/M/p2p$ SYSTEM (*continued*)

The system is a birth-death process with balance equations

$$\pi_n \lambda = \pi_{n+1} \mu (n+1) n,$$

whose solution is the queue's stationary probability measure

$$\mathbb{P}[Q = n] = \pi_n = \frac{\frac{(\lambda/\mu)^{n-1}}{n!(n-1)!}}{\sum_{k=1}^{\infty} \frac{(\lambda/\mu)^{k-1}}{k!(k-1)!}}, \quad n \geq 1$$

The infinite sums are expansions of Bessel functions:

$$\sum_{n=0}^{\infty} \frac{x^n}{n!n!} = I_0(2\sqrt{x}), \quad \sum_{n=0}^{\infty} \frac{x^n}{(n+1)!n!} = \frac{1}{\sqrt{x}} I_1(2\sqrt{x})$$

In particular, the mean number of peers in system is

$$\mathbb{E}Q = \frac{I_0(2\sqrt{\lambda/\mu})}{I_1(2\sqrt{\lambda/\mu})} \sqrt{\frac{\lambda}{\mu}}$$

CONSTRUCTION: INFINITE CASE

- **Definition over finite time?**
- **Assume $F = 1$**
- **Proposition**
  If $D = \mathbb{R}^d$ and

$$\int_{1}^{\infty} \mathbf{f}(\mathbf{r})\mathbf{r}^{\mathbf{d}-\mathbf{1}}\, d\mathbf{r} < \infty$$

  then the spatial birth and death point process is uniquely defined on all finite time intervals $[t_0, t]$

PROOF: A–GRAPHICAL REPRESENTATION

- **$\mathbf{u}$ fixed, $\mathbf{t_0 < u}$**
- **Ψ_{t_0}: space-time arrival P.P.P. in $[t_0, u]$**
- **Point $p = (x_p, t_p)$ of Ψ_{t_0}**
- **Graphical representation of Shot-Noise**

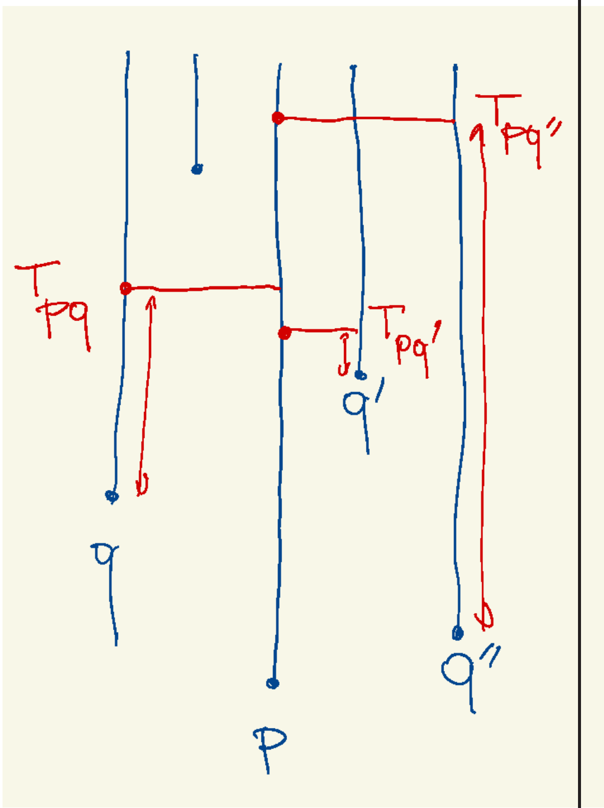

For all pairs $p \neq q \in \Psi_{t_0}$
Killing times T_{pq}

$$T_{pq} = T_{qp} \sim (t_p \vee t_q) + \mathrm{Exp}(2f(||x_p - x_q||))$$

Bernoulli directions of killing I_{pq}

$$I_{pq} = 1 - I_{qp} \sim \mathrm{Bernoulli}\,(\frac{1}{2})$$

PROOF: A–GRAPHICAL REPRESENTATION (*continued*)

- **Local finiteness in graphical construction**
- **Lemma If**

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

almost surely, none of the sets

$$N_p = \{T_{pq} : q \in \Phi_{t_0}\}$$

has accumulation points

- **Proof based on degree properties of the random connection model**

PROOF: A–GRAPHICAL REPRESENTATION (*continued*)

- **Let** $\mathbf{p} = (\mathbf{x}, \mathbf{t}) \in \mathbf{\Psi}_{\mathbf{t_0}}$

$$\begin{aligned}
\mathbb{E}^{\mathbf{p}}|\mathbf{N_p} \cap (\mathbf{t_0}, \mathbf{u}]| &= \mathbb{E} \int_{\mathbb{R}^{\mathbf{d}} \times (\mathbf{t_0}, \infty)} \mathbf{1}_{\mathbf{T_{pq}} \leq \mathbf{u}} (\mathbf{\Psi}_{\mathbf{t_0}} - \delta_{\mathbf{p}})(\,\mathrm{d}\mathbf{q}) \\
&= \mathbb{E} \int_{\mathbb{R}^{\mathbf{d}} \times (\mathbf{t_0}, \mathbf{u}]} \mathbf{1}_{\mathbf{T_{pq}} \leq \mathbf{u}} (\mathbf{\Psi}_{\mathbf{t_0}})(\,\mathrm{d}\mathbf{q}) \\
&= \lambda \int_{\mathbb{R}^{\mathbf{d}}} \int_{\mathbf{t_0}}^{\mathbf{u}} \mathbb{P}[\mathbf{Exp}(\mathbf{f}(\|\mathbf{x} - \mathbf{y}\|)) \leq \mathbf{u} - (\mathbf{t} \vee \mathbf{v})] \,\mathrm{d}\mathbf{v} \,\mathrm{d}\mathbf{y} \\
&\leq \lambda(\mathbf{u} - \mathbf{t_0}) \int_{\mathbb{R}^{\mathbf{d}}} \left(\mathbf{1} - \mathbf{e}^{-(\mathbf{u} - \mathbf{t_0})\mathbf{f}(\|\mathbf{y}\|)}\right) \,\mathrm{d}\mathbf{y} \\
&\leq \lambda(\mathbf{u} - \mathbf{t_0}) \left(\nu_{\mathbf{d}} + (\mathbf{u} - \mathbf{t_0}) \int_{\mathbb{R}^{\mathbf{d}} \setminus \mathbf{B}(\mathbf{0}, \mathbf{1})} \mathbf{f}(\|\mathbf{y}\|) \,\mathrm{d}\mathbf{y} \right) < \infty
\end{aligned}$$

PROOF: B–CONSTRUCTION ALGORITHM

- **Not all killing epochs lead to death: only living peers matter**
- **Death times solution of the infinite recursive equation**

$$\delta_{\mathbf{p}} = \inf \{\mathbf{T}_{\mathbf{pq}} : \mathbf{q} \in \mathbf{\Psi}_{\mathbf{t_0}}, \ \delta_{\mathbf{q}} \geq \mathbf{T}_{\mathbf{pq}}, \ \mathbf{I}_{\mathbf{pq}} = \mathbf{1}\}$$

- **The Construction Algorithm**
gives the solution of this recursive equation on compacts of time
- **Principle**

pick a node, check its earliest killing time;
determine whether the killer's death time is earlier or later than this time...

PROOF: B–CONSTRUCTION ALGORITHM (*continued*)

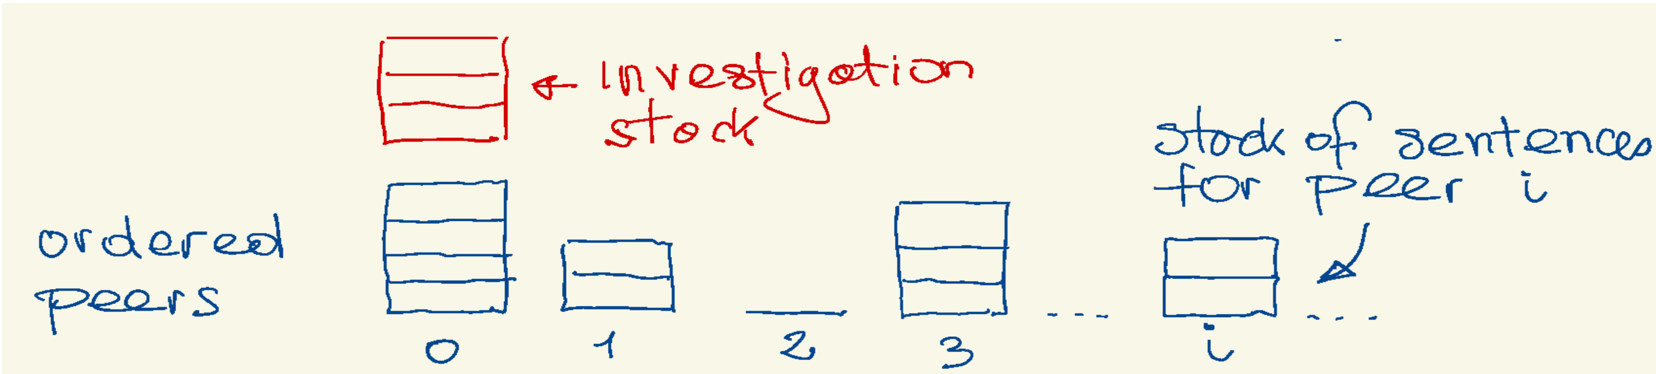

Algorithm : construction of death process on (t_0, ∞)

1. **Initialization:**
   - **each peer has a stack of ordered death sentences (peer, time, killer) - earliest on top**
   - **there is a global investigation stack, initially empty**

PROOF: B–CONSTRUCTION ALGORITHM (*continued*)

2. **If investigation stack empty:**
   pick first peer with top sentence with time $< u$ and no certificate;
   move this sentence to the investigation stack;
   if no such sentence exists, stop;

3. **Look at the top of investigation stack, say (x, s, y), and do**

   – **If y's stack has on top a death sentence or certificate later than s, then death happens: change the sentence (x, s, y) into death certificate with same date and return it to the top of x's stack;**

   – **If y has death certificate earlier than s, the sentence (x, s, y) is removed from investigation stack and deleted;**

   – **Otherwise move top sentence of y's stack to investigation stack;**

4. **Go to 2.**

PROOF: B–CONSTRUCTION ALGORITHM (*continued*)

- **Lemma If**

$$\int_0^\infty f(r)r^{d-1}\,dr < \infty,$$

for each sentence, the sequence of peers $p_1, p_2, \ldots$ produced by its investigation is such that $(T_{p_n p_{n+1}})$ is decreasing and a.s. finite

- **Proof For $u - t_0$ small enough, an upper-bound random connection model with connection function f does not percolate. For more general u, decompose $[t_0, u]$ in small intervals and apply the last observation**

- **Theorem For every peer p born in (t_0, ∞) the construction algorithm determines a unique death time $\delta_p \leq t$**

- **Corollary Almost surely, the death process is defined uniquely as a factor of Ψ_{t_0}**

PROOF: B–CONSTRUCTION ALGORITHM (*continued*)

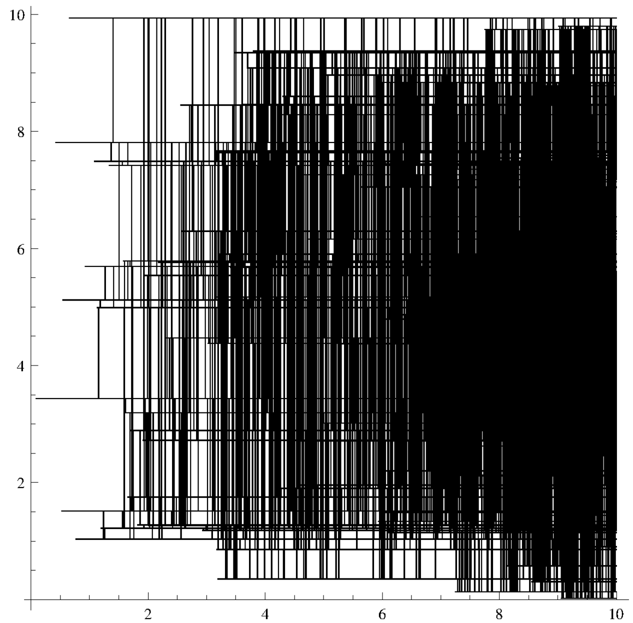

Graphical Representation

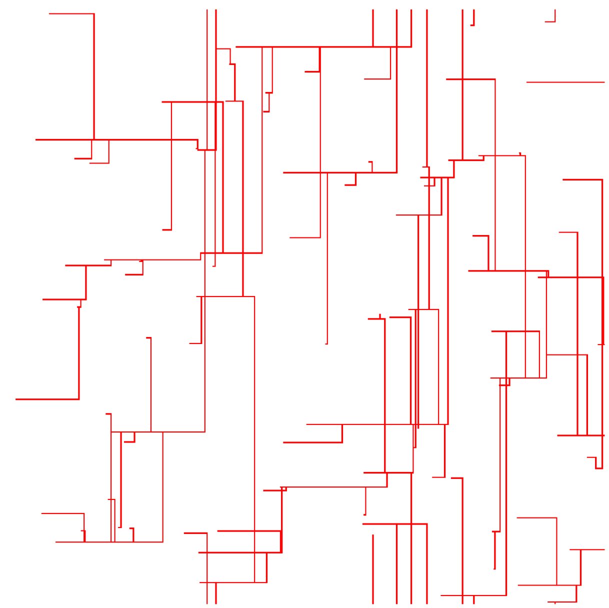

State Construction

PROOF: C–CONSTRUCTION PROPERTIES

- **When the process is a.s. well-defined, each peer p is a.s. killed by a uniquely determined peer $\kappa(\mathrm{p})$**
- **Since κ is non-cyclic, its graph is a forest of infinite trees**
- **Conjecture**
  When $\mathrm{D} = \mathbb{R}^2$, the directed graph $(\Psi_{\mathrm{t}_0}, \{(\mathrm{p}, \kappa(\mathrm{p})) : \mathrm{p} \in \Psi_{\mathrm{t}_0}\})$ is almost surely a tree
- **Proposition**
  For any peer p, the conditional distribution of the number of peers it kills, given the history of the process up to time t_p, is Geom($\frac{1}{2}$)
- **Observation**
  This does not show yet that the death process is well-defined when $\mathrm{t}_0 = -\infty$

II. CONSTRUCTION OF THE STATIONARY REGIME

- **Existence/uniqueness of stationary regimes**

- **Theorem** **Under the assumptions**

(i) $\mathbf{a := \int_{R^d} f(||x||)dx < \infty}$
(ii) f is non-increasing and bounded
there exists a unique stationary regime holding for all initial conditions made of a homogeneous Poisson point process of initial nodes.

II. CONSTRUCTION OF THE STATIONARY REGIME (*continued*)

- **Proof Structure: 2 coupling steps**
  - **Coupling 1: couple the two process defined on $(0,\infty)$ with and without the addition of a P.P. of initial nodes at time 0 and show that the effect of this addition vanishes at infinity**
  - **Coupling 2: Coupling from the past: $(0,\infty) \rightarrow (-\infty),0)$**
- **Class of additional point processes Z:**
  motion-invariant P.P. with a few additional properties, e.g.
  - **homogeneous P.P.P.**
  - **dependent thinning of homogeneous P.P.P.**

COUPLING 1

- **Two parallel histories**
  - **Black + Green without add. point**
  - **Black + Red with add. point**
- **Built by a variant of the graphical representation and construction algorithm**

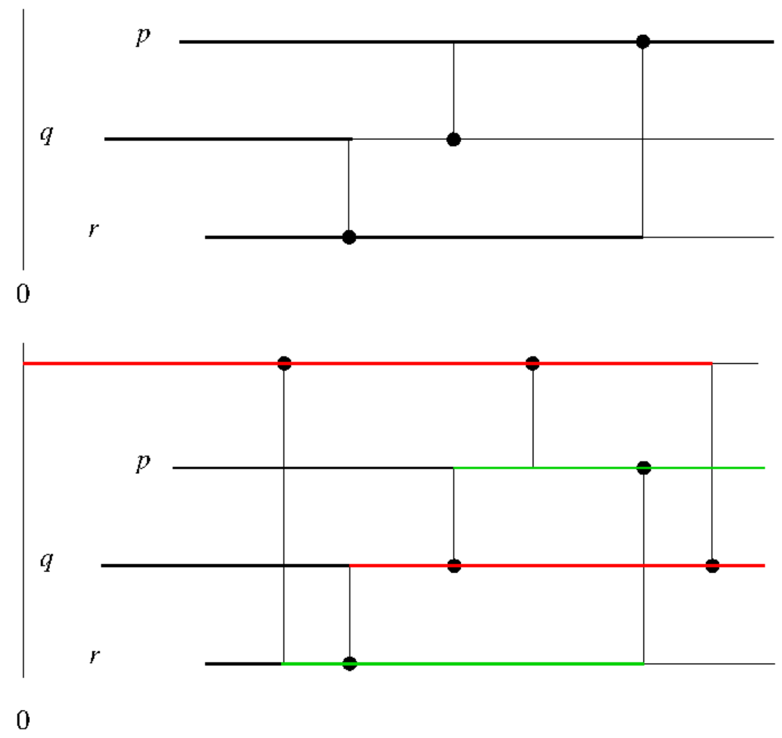

COUPLING 1 (*continued*)

Effect of the addition of an independent homogeneous P.P.P. of peers at time 0 on the two coupled histories on $(0, \infty)$

- **Lemma**
  Under the assumptions of the theorem, the number of special points (red and green) created by any single point of the additional P.P.P. is a.s. finite.
- **Proof supermartingale argument on cardinality of special progeny of this point**
- **Lemma**
  Under the assumptions of the theorem, the intensity of special points (red and green) created by the additional P.P.P. decreases exponentially fast with time
- **Proof Abel & Cain argument**

COUPLING 1 (*continued*)

Proof of First Lemma (Sketch)

- $\mathrm{S_n(z)}$ set of special points stemming from z at n-th duel concerning this population
- If duel between regular and special:

$$\mathrm{card(S_n(z)) := card(S_n(z)) + J_n}$$

with $\mathrm{J_n} \sim$ +1 w.p. 1/2 and -1 w.p. 1/2, cond. independent on the past

- If duel between special and special:

– either $\mathrm{card(S_n(z))}$ is unchanged (specials of the same kind)

– or it evolves as above (different kinds and same progeny)

– or its evolution is bounded above by the last equation (different kind and different progeny)

COUPLING 1 (*continued*)

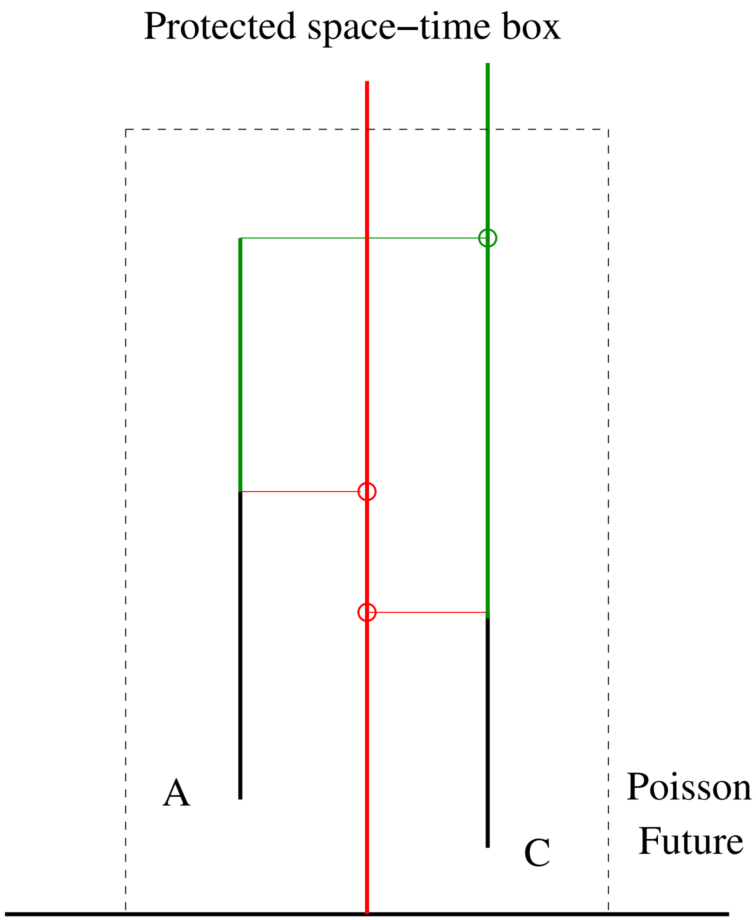

Proof of Second Lemma (Sketch)

uniformly over time and space, the offspring cardinality of a special point (here red) has a strictly negative drift through the Cain kills Abel scheme

- **This requires proving the Palm expectation of the death pressure on specials is uniformly bounded;**
- **This implies exponential decrease of the density of special points.**

COUPLING 2: FROM THE PAST

- $\Psi^{Z}_{s,t}$: **configuration at** t **built by Coupling 1 when initial time is** s **and initial condition is** Z
- **Corollary of second Lemma**
  For all compacts K **of space, there is a finite expectation time** $T_K(s)$ **such that for all** $t > T_K(s)$, $\Psi^{Z}_{s,t}$ **has no special points in** K
- $\Psi_{s,t}$: **configuration at** t **when initial time is** s **and initial condition is** $\emptyset$
- **From Lemma, for all** t

$$\exists \lim_{s \to -\infty} \Psi_{s,t} = \Phi_t$$

  with Φ_t **translation invariant w.r.t. space and time**
- **Theorem 1 follows**

3. f-REPULSION

- **Theorem**
  For all f such that there exists a non degenerate translation invariant stationary regime $\Phi = \Phi_0$, 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 $\mathbb{P}_0$ 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**

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

with

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

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

with n_0 the total population and

$$\mathbf{a} = \int_{\mathbf{T}} \mathbf{f}(||\mathbf{x}||)\mathbf{m}(\,\mathrm{d}\mathbf{x})$$

with T the torus of area $|\mathrm{D}|$

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:**

$$\mathbb{E}(\mathbf{n}_0)\frac{\mathbf{a}}{|\mathbf{D}|} = \frac{\mathbb{E}_0\left((\mathbf{A}_0(\mathbf{0}))^2\right)}{\mathbb{E}_0\left(\mathbf{A}_0(\mathbf{0})\right)}$$

- **Using the fact that**

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

we get

$$\mathbb{E}(\mathbf{n}_0)\frac{\mathbf{a}}{|\mathbf{D}|} \geq \mathbb{E}_0\left(\mathbf{A}_0(\mathbf{0})\right)$$

4. BIRTH AND DEATH EQUATIONS FOR MOMENT MEASURES

- **Factorial moment measure densities** e.g.

$$\mathbb{E}^0_{\Phi}[\Phi^!(B)] = \frac{1}{\beta} \int_B m^{[2]}_{mi}(\|x\|)dx$$

- **Theorem**

The time stationary SBD satisfies the following balance relations:

$$\int_{R^d} m^{[2]}_{mi}(\|x\|)f(\|x\|)dx = \lambda$$

- **RCP for first moment measure**

■ RCP Hierarchy of Equations for Moment Measures

For all $k \geq 2$, for all $x_1, \ldots, x_k$ in $\mathbb{R}^d$,

$$\mathbf{m}^{[k]}(\mathbf{x}_1, \ldots, \mathbf{x}_k) \left(\sum_{i=1,k} \sum_{j=1,k,\ j \neq i} \mathbf{f}(\|\mathbf{x}_i - \mathbf{x}_j\|) \right)$$
$$+ \int_{\mathbf{R}^d} \mathbf{m}^{[k+1]}(\mathbf{x}_1, \ldots, \mathbf{x}_k, \mathbf{z}) \left(\sum_{i=1,k} \mathbf{f}(\|\mathbf{x}_i - \mathbf{z}\|) \right) \mathbf{dz}$$
$$= \lambda \sum_{i=1,k} \mathbf{m}^{[k-1]}(\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \ldots, \mathbf{x}_k)$$

II. TCP EXAMPLE

- **TCP model**: D **is the Euclidean plane** $\mathbb{R}^2$

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

- **[Mathis, Semke, Mahdavi & Ott, 97]**
- **Most results extend to the general** f **case**

1. DIMENSIONAL ANALYSIS

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

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

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

- **Sketch of proof**
  - **choose R as a new distance unit, then**
    - **the arrival intensity becomes $l = \lambda R^2$**
    - **the download speed 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 intensity becomes $l = \frac{\lambda F R^3}{C}$**

- **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 node point process**
  - μ_o **the mean rate of a typical node**
  - W_o **the mean latency of a typical node**
  - N_o **the mean number of nodes in a ball of radius** R **around a typical node**

2. FIRST ORDER APPROXIMATION

- **Fluid, or Poisson heuristic: obtained when ignoring the Palm expectation of the rate, and using the mean rate at a typical location instead:**

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

with β_f the density of nodes in this heuristic

FIRST ORDER APPROXIMATION AS AN ASYMPTOTIC

- **Theorem**

  When ρ tends to infinity:

  – **The fluid heuristic is asymptotically tight:**
  $\beta_{\mathrm{o}} \to \beta_{\mathrm{f}}$, $\mathbf{W}_{\mathrm{o}} \to \mathbf{W}_{\mathrm{f}}$, $\mu_{\mathrm{o}} \to \mu_{\mathrm{f}} \cdots$

  – **The law of the latency of a typical node converges weakly to an exponential random variable of parameter** $\mathrm{W}_{\mathrm{f}} = \frac{\mathrm{F}}{\mu_{\mathrm{f}}}$

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

FIRST ORDER APPROXIMATION AS AN ASYMPTOTIC *(continued)*

- **In this heuristic/limit**

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

- **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

- 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 $\mathbf{W_f = \sqrt{F/(\lambda 2\pi CR)}}$**
  If two nodes are at a distance $\mathrm{r_0}$ such that

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

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

FIRST ORDER APPROXIMATION AS A BOUND

- **In the TCP case, the repulsion theorem 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**

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

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 node 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

$$\mathbf{N_f \ll 1}$$

$$\Downarrow$$

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

$$\Downarrow$$

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

$$\Downarrow$$

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

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

SECOND ORDER APPROXIMATION

- **Second order heuristic:**
  - **considers $\hat{\mu}$, the unique solution of**

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

  - **then defines**

$$\hat{\beta} = \lambda \mathbf{F} / \hat{\mu}, \qquad \hat{\mathbf{W}}_{\mathbf{h}} = \mathbf{F} / \hat{\mu}$$

SECOND ORDER APPROXIMATION (*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:**

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

SECOND ORDER APPROXIMATION (*continued*)

- **Then**

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

that is

$$\mathbf{m_{[2]}(x,y)} \approx \lambda\mathbf{F}\frac{\beta_o}{\frac{\mathbf{C1}_{||\mathbf{x}-\mathbf{y}||\le\mathbf{R}}}{||\mathbf{x}-\mathbf{y}||}+\mu_o}$$

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

SECOND ORDER APPROXIMATION (*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

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

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:**

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

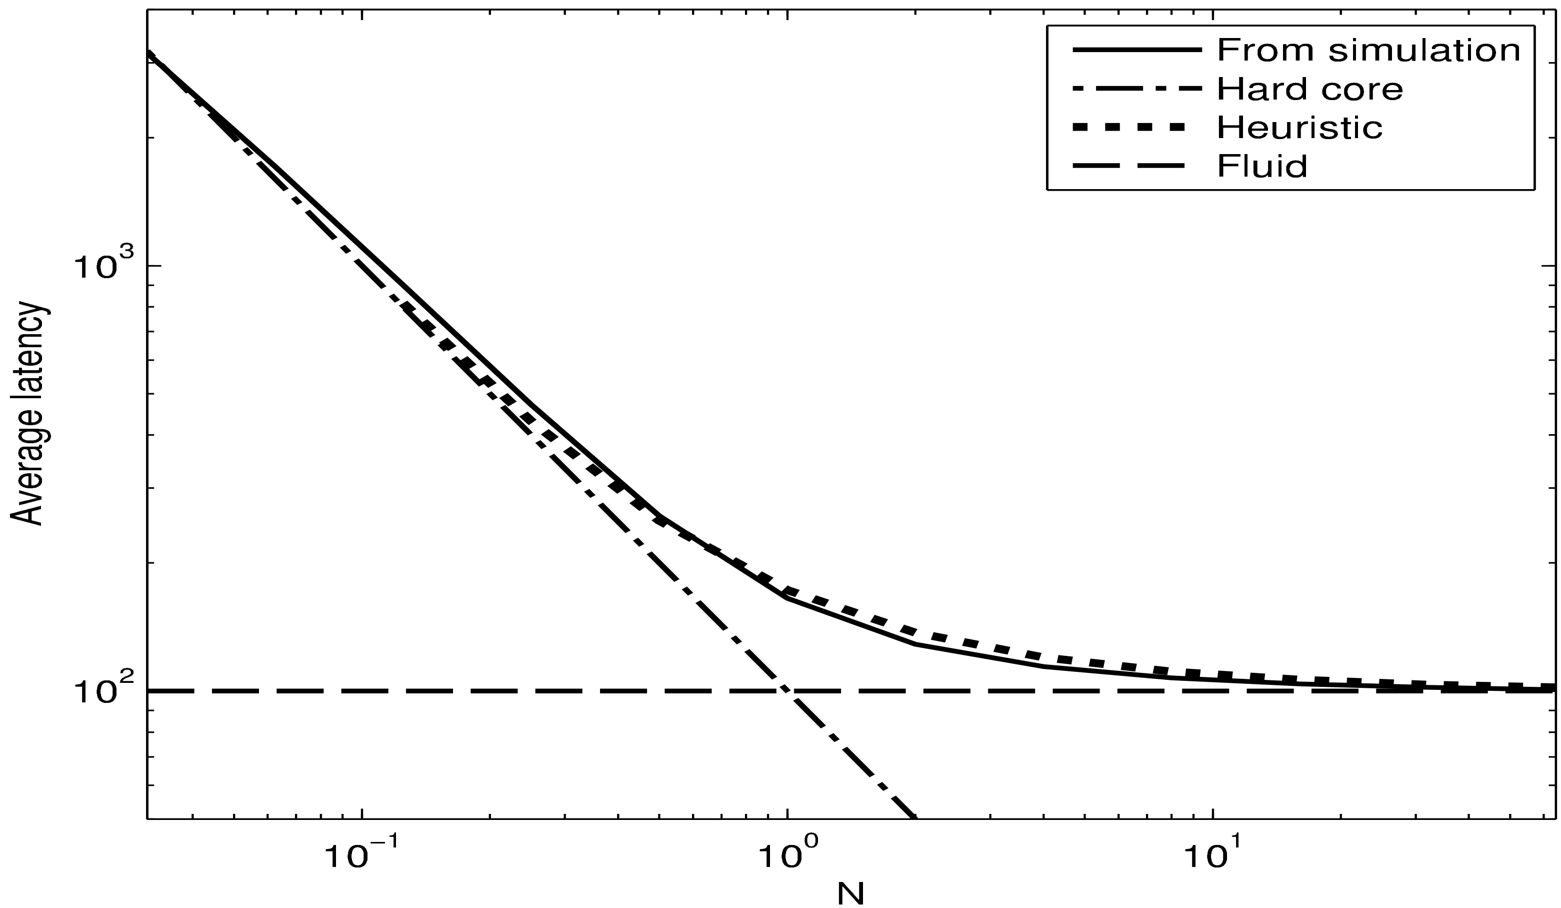

Latency in function of N_f.

4. SUPER-SCALABILITY

- **Dimensional analysis tells us that**

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

where M only depends on $\mathrm{N_f} = \sqrt{\frac{\pi\lambda\mathbf{F}\mathbf{R}^3}{2\mathbf{C}}}$ 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.**

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**
$\mathbf{W} = \frac{\mathbf{1}}{\mu - \lambda}, \lambda < \mu$	$\mathbf{W} = \frac{\mathbf{1}}{\mu}$	$\mathbf{W} = \frac{\mathbf{m}(\lambda)}{\sqrt{\lambda}}, \mathbf{m}(\cdot) \downarrow$

ADAPTING THE PEERING RADIUS

- **Mean Constant Number of Nearest Nodes: take as neighbors the nodes in a ball with a radius R such that the mean number of other nodes in the ball is L i.e. $\pi \mathrm{R}^2 \beta_{\mathbf{o}} = \mathrm{L}$, where $\beta_{\mathbf{o}}$ is the (unknown) steady state intensity of the point process ϕ_{t}. Then**

$$\mathbf{f}(\mathbf{r}) = \frac{\mathbf{C}}{\mathbf{r}} \mathbf{1}_{\mathbf{r} \le \mathbf{R}}, \quad \mathbf{R} = \sqrt{\frac{\mathbf{L}}{\pi \beta_{\mathbf{o}}}}$$

- **General Case**

$$\mathbf{f}(\mathbf{r}) = \frac{\mathbf{C}}{\mathbf{r}} \mathbf{1}_{\mathbf{r} \le \mathbf{R}}, \quad \mathbf{R} = \kappa \beta_{\mathbf{o}}^{-\alpha}$$

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

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

ASYMPTOTIC BEHAVIOR

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

EXTENSIONS & OPEN QUESTIONS

- **Extensions (P2P version to be presented at IEEE Infocom 13)**
  - **Seeders**
  - **Wireless**
- **Challenges (Applied Probability Journal Paper in preparation)**
  - **Hard core regime**
  - **Higher order approximations based on moment measures**
  - **Chunks**

CONCLUSION

- A **new point process dynamics** for handling P2P and CDN
- Has a unique stationary regime which exhibits a form of **repulsion**
- Satisfies a hierarchy of birth and death type integral relations which lead to **good approximations**