



HAL
open science

Performance of P2P Networks with Spatial Interactions of Peers

François Baccelli, Fabien Mathieu, Ilkka Norros

► **To cite this version:**

François Baccelli, Fabien Mathieu, Ilkka Norros. Performance of P2P Networks with Spatial Interactions of Peers. [Research Report] RR-7713, 2011, pp.23. inria-00615523v1

HAL Id: inria-00615523

<https://inria.hal.science/inria-00615523v1>

Submitted on 19 Aug 2011 (v1), last revised 18 Jun 2012 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Performance of P2P Networks with Spatial
Interactions of Peers*

François Baccelli — Fabien Mathieu — Ilkka Norros

N° 7713

August 2011

— Networks and Telecommunications —

*R*apport
de recherche

Performance of P2P Networks with Spatial Interactions of Peers

François Baccelli ^{*†}, Fabien Mathieu ^{* ‡}, Ilkka Norros[§]

Theme : Networks and Telecommunications
Networks, Systems and Services, Distributed Computing
Équipes-Projets GANG, TREC

Rapport de recherche n° 7713 — August 2011 — 23 pages

Abstract: We propose a new model for peer-to-peer networking which takes the network bottlenecks into account beyond the access. This model allows one to cope with the fact that distant peers often have a smaller rate than nearby peers. We show that the spatial point process describing peers in their steady state exhibits an interesting repulsion phenomenon. We study the implications of this phenomenon by analyzing two asymptotic regimes of the peer-to-peer network: the fluid regime and the hard-core regime. We get closed form expressions for the mean (and in some cases the law) of the peer latency and the download rate obtained by a peer as well as for the spatial density of peers in the steady state of each regime. The analytical results are based on a mix of mathematical analysis and dimensional analysis and have important design implications. The three key conjectures which underpin the analysis are substantiated by simulation.

Key-words: P2P; Point processes; Spatial birth and death processes; Congestion control

* INRIA

† ÉNS

‡ University Paris 7

§ VTT (Finland)

Systemes pair-à-pair à débits géo-dépendants

Résumé : Nous proposons une nouvelle modélisation des systèmes pair-à-pair pour les situations où le lien d'accès n'est plus l'unique goulot d'étranglement réseau. Ce modèle permet entre autres de prendre en compte la corrélation entre proximité et débit. Nous montrons que le processus ponctuel des pairs affiche un phénomène de répulsion dans son régime stationnaire.

Nous étudions plus particulièrement les deux régimes limites du système, qui sont un régime fluide et un régime de type *boules dures*. Nous donnons des expressions explicites de la valeur moyenne (voire de la loi) de la latence et du débit dans les différents régimes. Les résultats sont obtenus par un mélange d'analyse mathématique et dimensionnelle. Ils ont un impact significatif en terme de dimensionnement. Les conjectures servant de base aux résultats plus élaborés sont validées à l'aide de simulations.

Mots-clés : Pair-à-pair, processus de naissance et de mort, contrôle de congestion

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 1.1 | Contribution | 4 |
| 1.2 | Related work | 5 |
| 2 | Model | 6 |
| 3 | Mathematical Analysis | 9 |
| 3.1 | Dimensional Analysis | 9 |
| 3.2 | Repulsion | 10 |
| 3.3 | Fluid Model/Regime | 10 |
| 3.4 | Fluid as a Bound | 11 |
| 3.5 | Hard-Core Regime | 11 |
| 3.6 | Heuristic | 12 |
| 4 | Simulation Results | 13 |
| 4.1 | Average Performance Study | 14 |
| 4.2 | Fluid Model | 15 |
| 4.3 | Hard-Core Model | 16 |
| 4.4 | Intermediate Values | 17 |
| 4.5 | Summary of Simulations | 17 |
| 5 | Consequences on the Design of P2P Systems | 17 |
| 5.1 | Impact of Parameters | 17 |
| 5.2 | Asymptotic Behavior for Large Arrival Rates | 18 |
| 6 | Extensions of the Basic Model | 19 |
| 6.1 | Abandonment | 19 |
| 6.2 | More General Rate Functions | 19 |
| 6.2.1 | Affine RTT | 19 |
| 6.2.2 | Overhead | 19 |
| 6.2.3 | Per Flow Rate Limitation | 20 |
| 6.3 | Per Peer Rate Limitation | 20 |
| 6.4 | Leechers and Seeders | 20 |
| 6.4.1 | Fluid limit | 20 |
| 6.4.2 | Hard-core limit | 21 |
| 6.5 | Mixed extensions | 21 |
| 7 | Conclusion | 21 |

1 Introduction

Peer-to-peer (P2P) architectures have been widely used over the Internet in the last decade. The main feature of P2P is that it uses the available resources of participating end users. In the field of content distribution (file sharing, live or on-demand streaming), the P2P paradigm has been widely used to quickly deploy low-cost, scalable, decentralized architectures. For instance, the ideas and success of BitTorrent [Coh06] have shown that distributed file-sharing protocols can provide practically unbounded scalability of performance. Although there are currently many other architectures that compete with P2P (dedicated Content Distribution Networks, Cloud-based solutions, . . .), P2P is still unchallenged with respect to its low-cost and scalability features, and remains a major actor in the field of content distribution.

The Achilles' heel of today's P2P content distribution is the access upload bandwidth, as even high-speed Internet access connections are often asymmetric with a relatively low uplink capacity. Therefore, most theoretical models of P2P content distribution presented so far have been 'traditional' in the sense of assuming a common, relatively low access bandwidth, in particular concerning the upload direction, which functions as the main performance bottleneck. However, in a near future the development of very high speed access will ruin the justification of this assumption. This raises the need of new P2P models that describe what happens when the access is not necessarily the main/only bottleneck.

It makes sense to consider some type of distance in the analysis of the bandwidth in future networks, since end-to-end rates are in general (decreasing) functions of some distance. For instance, transport protocols like TCP provide transfer speeds as a decreasing function of round-trip time (RTT). Two nodes with short distance can communicate at very high rates, whereas long-distance transfers happen at much lower speeds.

1.1 Contribution

This paper presents a new kind of P2P content distribution model, where the peers are located in a metric space, and the dependence of transfer rates on distance plays a central role.

Our main focus is on a model presented in Section 2, which features the following characteristics:

Best effort content distribution

We focus on a basic filesharing QoS, with download time (latency) as the only performance metric. The typical P2P application we have in mind is a BitTorrent-like file-sharing system. However, the high abstraction level of our model allows also interpretations beyond this framework. Using proper QoS requirements, it could be extended to any service that is provided mutually between participating peers, like all P2P content distribution services (including live and on-demand streaming).

Fading bandwidth

We consider a scenario where peers randomly appear in some metric space, and download from their neighbors with a rate that decreases with distance. Many situations can fall within this framework. For instance, the space could be a representation of peers' interest, the position of a peer representing its own centers of interest. In such a space, two close peers share common interests, and therefore are likely to exchange more data. Our approach could also be used to model P2P behavior in an ad-hoc wireless sensor network (WSN), where nodes would randomly awake and try to get information from near awoken nodes. The distance-decreasing property would then come from the wireless physical layer. In this paper, we focus on a case study where the rate is defined by a function (typically TCP-like) of the end-to-end round-trip-time (RTT). For simplicity, we consider that RTTs are proportional to distances in some 2-dimensional space, although the results we give can straightforwardly be generalized to higher dimensions. Our framework could also be adapted to non-Euclidean models of the Internet, like a hyperbolic 'Internet space' [KPK⁺10], or inframetric [FLV08].

One theoretically novel feature of our models is the emergence of a *repulsion* phenomenon: as close peers get faster rates, they quit the system earlier, so a node "sees" less peers in its immediate vicinity than one would expect by considering the spatial entrance distribution alone. Assuming repulsion as an intuitively sound conjecture, we build a mathematical analysis of the model in Section 3. In details, we use dimensional analysis to characterize the system by one single dimensionless parameter. We show that two limit scenarios exist for the steady state regimes, a fluid one and a hard-core one. The performance law is explicitly given for both scenarios. While limit scenarios provide bounds for the intermediate states, we also provide a heuristic function that gives a reasonably good estimate of the average performance. All these results are validated through simulations in Section 4.

We discuss the impact of our results in Section 5 in terms of parameters influence and dimensioning rules. In particular, one interesting property of our systems is *supra-scalability*: their performance can improve with the number of users (as long as access bandwidth is not a bottleneck).

Finally, we give in Section 6 a first, non-exhaustive, list of the possible variants that can be built on our model, including taking into account the possible seeding behavior of users, an access bottleneck, or a protocol overhead.

1.2 Related work

Our main scenario is inspired by a BitTorrent-like file-sharing protocol. In BitTorrent [Coh06], a file is segmented into small chunks and each downloader (called *leecher*) exchanges chunks with its neighbors in a peer-to-peer overlay network. A peer may continue to distribute chunks after it has completed its own download (it is called a *seeder* then). Theoretical studies and modeling have already provided relatively good understanding of BitTorrent performance.

Qiu and Srikant [QS04] analyzed the effectiveness of P2P file-sharing with a simple dynamic system model, focusing on the dynamics of leechers and seeders. Massoulié and Vojnovic [MV05] proposed an elegantly abstracted stochastic chunk-level model of uncoordinated file-sharing. Their results indicated that stable continuous operation of a file-sharing torrent could be possible practically

without seeders. Following Mathieu and Reynier [MR06], Hajek and Zhu [HZ10] showed however that, in the same model, uncoordinated greedy downloading leads to the ‘missing piece syndrome’, where one chunk keeps existing in very few copies while the peer population grows unboundedly. On the other hand, this syndrome can be avoided by using more sophisticated download policies [Rei09, OAN11] at the cost of somewhat increased download times. Although these results are obtained in a different network model (homogeneous, potentially fully connected network), we neglect chunk-level modeling in this paper and leave the rare chunk problem in our model for future work, focusing on a bandwidth-centered approach similar to the one proposed by Benbadis *et al.* [BMHP08].

The natural feature of large variation of transfer speeds in P2P systems has been considered in a large number of papers. For example, part of the peers can rely on cellular network access that is an order of magnitude slower than fixed network access used by the other part. Such scenarios differ however substantially from our model, where the transfer speeds depend on pair-wise distances but not on the nodes as such.

There are some earlier papers considering P2P systems in a spatial framework. As an example, Susitaival *et al.* [SAV06] assume that the peers are randomly placed on a sphere, and compare nearest peer selection with random peer selection in terms of resource usage proportional to distance. However, the distance has no effect on transfer speed in their model. Our paper seems to be the first where a peer’s downloading rate is a function of its distances to other peers.

2 Model

Our peers live in a metric space (D, \mathfrak{d}) . The time parameter set is either \mathbb{R}_+ or, in the case of a stationary process, \mathbb{R} . Let m be a reference measure on the Borel σ -algebra \mathcal{D} of D (in most cases, m is a Lebesgue measure). We assume that new peers arrive according to a Poisson process with spatio-temporal intensity $\lambda dm dt$ (‘Poisson rain’), where λ describes the birth rate of peers. I.e., the number of peer arrivals taking place in the set $B \subset D$ in an interval $[s, t]$ is a Poisson random variable with parameter $\lambda m(B)(t - s)$. The set of the positions of peers present at time t is denoted by ϕ_t .

Assume further that each peer \mathfrak{p} is born with an individual service requirement $F_{\mathfrak{p}} > 0$. In the basic example where the service required by every peer consists of downloading one and the same file, $F_{\mathfrak{p}}$ would most naturally be modeled as a constant F describing the size of the file. For generality, we let the $F_{\mathfrak{p}}$ ’s to be independent random variables with finite expectation, denoting $F = EF_{\mathfrak{p}}$. On the other hand, to gain in mathematical tractability, we specify in this paper their common distribution as $Exp(1/F)$.

We assume that two peers at locations x and y serve each other at rate $f(\mathfrak{d}(x, y))$, where f is a non-increasing function (symmetric model). f can be seen as the feasible rate between x and y ¹. The services received from several

¹We implicitly assume that peers do not try to restrain their bandwidth capacities. In absence of natural altruism, many incentives can induce that behavior anyway, like the *Tit-for-Tat* algorithm [Coh06] or share-ratio policies [BMHP08].

peers are additive, so that the total download rate of a peer at x is

$$\mu(x, \phi_t) = \sum_{y \in \phi_t \setminus \{x\}} f(\mathfrak{d}(x, y)).$$

We call $\mu(\cdot, \cdot)$ the *bit rate function* of the model. A (leecher) peer \mathfrak{p} born at point $x_{\mathfrak{p}}$ at time $t_{\mathfrak{p}}$ leaves the system when its service requirement has been fulfilled, i.e. at time

$$\tau_{\mathfrak{p}} = \inf\{t > t_{\mathfrak{p}} : \int_{t_{\mathfrak{p}}}^t \mu(x_{\mathfrak{p}}, \phi_s) ds \geq F_{\mathfrak{p}}\}.$$

(The situation where finished leechers continue as seeders will be considered in Section 6.) A benefit of the memoryless property of the exponential distribution is that the remaining service requirements of peers do not need to be included in the system state description (see [QS04]). Instead, ϕ_t can be considered as a Markov process living in the space of spatial configurations, i.e. the space $\mathcal{N}(D)$ of counting measures in (D, \mathcal{D}) . Such processes are known as spatial birth-and-death processes, see [Pre75, GK06]. Thus, a peer located at x has death intensity $\mu(x, \phi_t)/F$.

A remarkable consequence of the peer-to-peer paradigm is the following

Lemma 1 If D is compact and $m(D) < \infty$, then the Markov process ϕ_t is ergodic for any birth rate $\lambda > 0$.

Proof. (Sketch) Choose a number $z_0 > 0$ such that $f(z_0) > 0$ and split D into cells with diameters at most z_0 . Then all peers in a cell with population higher than one receive service at least at rate $f(z_0)$. It follows that the population of each cell is stochastically dominated by an $M/M/\infty$ queue that is modified so that a lone customer cannot leave. Since such queues are stable with any input rate, the distribution of $(|\phi_t| : t \geq 0)$ is tight, whatever the initial state ϕ_0 . The ergodicity can now be shown by a standard coupling argument: two realizations with different initial states but same arrival process couple in finite time. \square

The domination argument above works also for unbounded domains, but the ergodicity cannot be established as easily. Garcia and Kurtz [GK06] proved the existence and ergodicity of wide classes of *attractive* spatial birth-and-death processes in infinite domains. Whether their approach is extendable to our repulsive case remains a challenge for future work; below, we will assume that the spatial birth and death processes admit a unique stationary regime in (infinite) Euclidean spaces.

Examples of Bit Rate Functions

Our basic model for the spatial space D is the Euclidean plane \mathbb{R}^2 . We assume that the model admits a steady state regime where the peers (in the basic model all leechers) form in \mathbb{R}^2 a stationary and ergodic point process with density β .

The examples below focus on the situation where the peers use a congestion control mechanism like TCP², and where the bandwidth limitation is not at

²In P2P systems, UDP is often used in place of TCP. P2P-over-UDP protocols try to be TCP-friendly [Nor09, Sha10]: they are designed to respect TCP flows and actually mimic TCP. For instance, TFRC (www.ietf.org/rfc/rfc3448.txt) recommends that UDP flows use the square root formula to predict the bit rate that a TCP flow would get and use this rate for throttling their traffic. The above model is hence directly applicable to such a setting.

the access but in the Internet. The case where both can be bottleneck will be considered in Section 6.

Consider first the case where peers use TCP Reno. On the path between two peers, let p denote the packet loss probability and RTT denote the round trip time. Then the square root formula [OKM92] stipulates that the rate obtained on this path is $\frac{\eta}{\text{RTT}\sqrt{p}}$, with $\eta \approx 1.309$. Assuming the RTT to be proportional to distance r yields a transfer rate of the form

$$g(r) = \frac{C}{r}. \quad (1)$$

We can refine (1) by assuming that RTT is not simply linear in distance r but some affine function of it, namely $\text{RTT} = ar + b$, where a accounts for propagation delays in the Internet path and b for the mean delay in the two access networks. Then the transfer rate between two peers with distance r becomes

$$g(r) = \frac{C}{r+q}, \quad \text{with } C = \frac{\eta}{a\sqrt{p}}, \quad q = \frac{b}{a}. \quad (2)$$

Another natural model is that where one accounts for an overhead cost of c bits per second. The rate between two peers at distance r is then

$$g(r) = \left(\frac{C}{r+q} - c \right)^+, \quad \text{with } (\cdot)^+ := \max(\cdot, 0). \quad (3)$$

In the main part of this paper, we assume that the transfer rate between two communicating peers is given by (1). The more general rate functions defined in (2) and (3) will be analyzed in Section 6, together with cases where some upload limitation is taken into account as well.

Having specified some transfer rate function g , we notice that a peer cannot interact with all other peers of the overlay network: it would result in a full mesh overlay, which is impossible to handle for large networks. Therefore, peers usually have to limit their neighborhood. For instance, a peer may select only peers within a certain distance³ and/or limit its total number of neighbors. This leads to the following two choices for bit rate function:

The *Peers in a Ball of Constant Radius* model: take $f(r) = g(r)1_{r \leq R}$ (R is called the *torrent range*), so that

$$\mu(x, \Phi) = \sum_{x_i \in \Phi, x_i \neq x} 1_{\|x_i - x\| \leq R} g(\|x_i - x\|). \quad (4)$$

The *Constant Number of Nearest Peers* model: take the L closest peers as the set of communicating neighbors. This rule is non-symmetric and difficult to deal with exactly, but an approximate version is obtained by taking as neighbors the peers in a ball with a radius R such that a mean number of other peers in the ball is equal to L : require $\pi R^2 \beta = L$, where β is the (unknown) steady state intensity of the point process ϕ_t . In this case, $\mu(x, \Phi)$ is as in (4) but with $R = \sqrt{\frac{L}{\pi\beta}}$.

Inspired by these two models, we will consider a general model with $R = \kappa\beta^{-\alpha}$ with α a real parameter. The constant radius ball corresponds to the case

³This approach is even more meaningful in the context of WSNs, as it can correspond to some transmission range.

Table 1: Table of Notation

| Name | Description | Units |
|---------------|----------------------|-----------------------------|
| Main Model | | |
| C | Speed parameter | $bits \cdot s^{-1} \cdot m$ |
| F | Mean file size | $bits$ |
| R | Peering range | m |
| λ | Leecher arrival rate | $m^{-2} \cdot s^{-1}$ |
| W | Mean latency | s |
| μ | Mean rate | $bits \cdot s^{-1}$ |
| Main variants | | |
| U | Upload bottleneck | $bits \cdot s^{-1}$ |
| T_S | Seeding time | s |

$\alpha = 0$ and the L nearest neighbor case to $\alpha = \frac{1}{2}$. Note that as β depends on R , $R = \kappa\beta^{-\alpha}$ has to be seen as a fixed point equation for $\alpha > 0$.

3 Mathematical Analysis

In this section, we focus on our basic model: D is the Euclidean plane and $g(r) = C/r$. Unless otherwise stated, we consider $\alpha = 0$ (constant radius). For this case, the system has 4 basic parameters: R in meters (m), F in bits, λ in m^{-2} per second (s) and C in $bit \cdot m \cdot s^{-1}$.

We denote by β_o the density of the peer (leecher) point process by μ_o the mean rate of a typical peer, by W_o the mean latency of a typical peer and by N_o the mean number of peers in a ball of radius R around a typical peer, all in the steady state regime of the P2P dynamics.

We will consider a fluid regime/limit, where the corresponding quantities will be denoted by β_f , μ_f , W_f and N_f , respectively as well as a hard-core regime/limit for which we will use the notation β_h , μ_h , W_h and N_h , respectively.

In any of these regimes, Little's law tells that the average density verifies $\beta = \lambda W$. Hence

$$\beta\mu = \lambda F. \quad (5)$$

3.1 Dimensional Analysis

We first use the π -theorem [Buc15] to strip our problem from redundant variables: if we choose R as a new distance unit, then the arrival intensity becomes $l = \lambda R^2$, the download constant becomes $c = C/R$ and the other parameters are unchanged. If we now define F as an information unit, then the download speed constant becomes $c = C/(RF)$ and the other parameters are unchanged. Finally, if we take a time unit such that the download speed constant is 1, we get a system where all parameters are equal to 1 but for the arrival rate which is equal to $l = \frac{\lambda FR^3}{C}$. As the system itself is not affected by the choice of measurement units, all its properties only depend on the parameter

$$\rho = \frac{\lambda FR^3}{C}. \quad (6)$$

By the same type of arguments, for all $\alpha \neq \frac{1}{2}$, all properties of the system only depend on the parameter

$$\rho = \frac{\lambda F}{C} \kappa^{\frac{3}{1-2\alpha}}. \quad (7)$$

For $\alpha = \frac{1}{2}$ (nearest peers), the parameter is $\rho = \kappa$ (or equivalently L).

3.2 Repulsion

We start with

Conjecture 1 For all f which is *decreasing*, and in particular for $f(r) = C/r$,

$$E \left[\sum_{x_i \in \phi \cap B(0,R)} f(\|x_i\|) \right] \geq E_0 \left[\sum_{x_i \in \phi \cap B(0,R)} f(\|x_i\|) \right], \quad (8)$$

where P_0 is the Palm probability [DVJ88] w.r.t. this point process.

The last property can be seen as a form of *repulsion* between the peers in the stationary regime of the P2P process, as it suggests there are less points in a ball of radius R around a typical peer than in a ball of the same radius around a typical location of the Euclidean plane.

3.3 Fluid Model/Regime

Let us start with the case $\alpha = 0$ for easy exposition. From Campbell's formula, the mean bit rate of a typical location of space (or equivalently of a newcomer peer) is

$$\mu_f = \beta_f 2\pi \int_{r=0}^R (C/r) r dr = \beta_f 2\pi C R. \quad (9)$$

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 as given in (9), we get that the mean latency of a peer is $W_f = \frac{F}{\mu_f}$. This heuristic will be referred to as the *fluid heuristic* below.

Conjecture 2 The fluid heuristic is asymptotically tight when ρ tends to infinity. In addition the law of the latency of a typical peer converges weakly to an exponential random variable of parameter W_f when ρ tends to infinity.

From (5), we have

$$\beta_f = \sqrt{\frac{\lambda F}{2\pi C R}}, \quad \mu_f = \sqrt{\lambda F 2\pi C R}, \quad W_f = \sqrt{\frac{F}{\lambda 2\pi C R}}. \quad (10)$$

In this heuristic, the mean number of peers in a ball of radius R around a typical peer is

$$N_f = \sqrt{\frac{\pi}{2}} \sqrt{\frac{\lambda F R^3}{C}} = \sqrt{\frac{\pi}{2}} \sqrt{\rho}. \quad (11)$$

For general α , similar computations give that $\mu_f = 2\pi C\kappa\beta^{1-\alpha}$, so that

$$\begin{aligned}\beta_f &= \left(\frac{\lambda F}{2\pi C\kappa}\right)^{\frac{1}{2-\alpha}} \\ W_f &= \lambda^{-\frac{1-\alpha}{2-\alpha}} F^{\frac{1}{2-\alpha}} (2\pi C\kappa)^{-\frac{1}{2-\alpha}} \\ \mu_f &= (2\pi C\kappa)^{\frac{1}{2-\alpha}} (\lambda F)^{\frac{1-\alpha}{2-\alpha}}.\end{aligned}\quad (12)$$

The algorithm which leads to this hence consists in 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, we get

$$\beta_f = \frac{\left(\frac{\lambda F}{2C}\right)^{\frac{2}{3}}}{(\pi L)^{\frac{1}{3}}}, \quad \mu_f = (2C)^{\frac{2}{3}} (\lambda F \pi L)^{\frac{1}{3}}, \quad W_f = \frac{\left(\frac{F}{2C}\right)^{\frac{2}{3}}}{(\lambda \pi L)^{\frac{1}{3}}}.\quad (13)$$

Here are a few comments on Conjecture 2:

- physical meaning: ρ is large when either the arrival intensity, or the file size, or the torrent range are large, or if the download speed constant C is small;
- Two peers that are at a distance r_0 get a mutual rate C/r_0 . If

$$r_0 \ll \sqrt{\frac{C}{2\pi\lambda FR}} = \frac{R}{\sqrt{2\pi\rho}},\quad (14)$$

one can verify using (10) that the corresponding latency is much smaller than W_f , so there is little chance to see r_0 -close peers in the steady state. This can be seen as a *hard exclusion* below that scale. We see that r_0 tends to 0 in configurations where ρ tends to infinity and R is fixed (for instance when λ tends to infinity).

3.4 Fluid as a Bound

Consider now a *free rider*, namely a peer that sends no bits to any other peer. The latter is invisible to other peers and the mean rate it receives is hence exactly μ_f as given by (9). Let μ_o denote the (exact) mean rate of a fair peer in the steady state regime, namely

$$\mu_o = E_0\left[\sum_{x_i \in \phi \cap B(0,R)} f(\|x_i\|)\right].$$

Conjecture 1 is equivalent to saying that $\mu_f \geq \mu_o$. Then it follows from $\mu_f \geq \mu_o$ and from the relations $W_o = F/\mu_o$ and $\beta_o = \lambda W_o$ that

$$W_o \geq W_f \quad \text{and} \quad \beta_o \geq \beta_f.\quad (15)$$

In words, repulsion implies that the fluid regime is actually a lower (resp. upper) bound for the mean latency and the peer density (resp. the mean rate).

3.5 Hard-Core Regime

A stationary point process is *hard-core* for balls of radius R if there are no points in a ball of radius R centered on any point of the point process (excluding the latter). We are now in a position to state our third conjecture:

Conjecture 3 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 , latency W_h and mean rate μ_h defined as follows:

$$\beta_h = \frac{1}{\pi R^2}, \quad W_h = \frac{1}{\lambda \pi R^2}, \quad \mu_h = \lambda F \pi R^2. \quad (16)$$

When ρ tends to 0, the cdf of the latency converges weakly to

$$1 - \frac{e^{-\frac{t}{2W_h}}}{2}, \quad t > 0. \quad (17)$$

Notice that ρ is small when either the arrival intensity, or the file size, or the torrent range are small, or if the download speed constant C is large.

We can complement Conjecture 3 by describing the limiting spatial birth-and-death process (with e.g. $C = \infty$): balls of radius R fall over the Euclidean plane as a Poisson rain of intensity λ . If a new ball $B(X, R)$ arriving at $X \in \mathbb{R}^2$ intersects no ball of the current hard-core configuration, say Φ , then it is accepted. If it does intersect the balls $B(X_1, R), \dots, B(X_n, R)$, with $X_1, \dots, X_n \in \Phi$ (notice that this implies that $n \leq 5$), then a new configuration of hard-core balls is chosen at random in the set $\{B(X, R), \cup_{S_n} B(X_i, R)\}$, where S_n is the set of all non empty subsets of $\{1, \dots, n\}$, and this new configuration replaces instantly the initial configuration $B(X_1, R), \dots, B(X_n, R)$. Namely, either we keep the ball at X and discard the n initial balls, or we discard the ball at X as well as all balls with centers not in the set. This new configuration is chosen with the probabilities $\pi(\emptyset)$ for the element $B(X, R)$ and $\pi(s)$ for the subset $s \in S_n$, with π solution of the following equations: $\pi(\{1, \dots, n\}) = 1/2$; for all $1 < k < n$,

$$\pi(\{i_1, \dots, i_k\}) = \frac{\sum_{p \notin \{i_1, \dots, i_k\}} \pi(\{i_1, \dots, i_k\} \cup \{p\}) f(\|X - X_p\|)}{2 \left(\sum_{q=1}^k f(\|X - X_{i_q}\|) + f(\|X - X_p\|) \right)}$$

and $\pi(\emptyset) = \sum_{p=1}^n \pi(\{p\}) \frac{1}{2}$.

It is worthwhile mentioning that according to (16), the *volume fraction* of the associated sphere packing model is $1/4$ (since we have a density $\frac{1}{\pi R^2}$ of non-intersecting balls of radius $R/2$). This volume fraction is hence the same as that of the Matérn hard-ball model in the so called jamming regime (see e.g. [BB09]).

3.6 Heuristic

The best approximation at hand is obtained from some heuristic factorization of the factorial moment measure of order 3 which is described below. It consists

in considering $\hat{\mu}$, the unique solution of

$$\hat{\mu}^2 = \mu_f^2 \left(1 - \frac{C}{\hat{\mu}R} \ln \left(1 + \frac{\hat{\mu}R}{C} \right) \right), \quad (18)$$

with μ_f defined in (10) and in defining $\hat{\beta} = \lambda F / \hat{\mu}$ and $\hat{W}_h = F / \hat{\mu}$. In order to derive (18), we use the balance equation for the second order factorial moment density, which reads

$$2\beta_o\lambda = 2m_{[2]}(x, y) \frac{C}{F} \frac{1_{\|x-y\| \leq R}}{\|x-y\|} + \frac{C}{F} \int_D m_{[3]}(x, y, z) \left(\frac{1_{\|x-z\| \leq R}}{\|x-z\|} + \frac{1_{\|y-z\| \leq R}}{\|y-z\|} \right) dz, \quad (19)$$

for all x and y . We then use the following approximations:

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

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

Then, we get from (19) that

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

that is

$$m_{[2]}(x, y) \approx \lambda F \frac{\beta_o}{\frac{C 1_{\|x-y\| \leq R}}{\|x-y\|} + \mu_o}, \quad (20)$$

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

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

Equation (18) follows directly.

Together with simulation, this heuristic is our best justification for Conjectures 2 and 3:

- When $\mu_o R / C$ tends to ∞ , then it follows from Equation (18) that $\mu_o \sim \mu_f$. This substantiates Conjecture 2.
- When $\mu_o R / C$ tends to 0, then, expanding the log in Equation (18) gives $\mu_o \sim \lambda F \pi R^2$, which substantiates Conjecture 3.

4 Simulation Results

In this section, we validate our results and substantiate our conjectures by means of simulations. We focus on the main studied case, ($\alpha = 0$). Other experiments are not presented due to space limitation, although they provide results similar to the one we expose here. For sake of computability, we approximate the infinite space by a torus of radius 1.

As stated by the dimensional analysis, we just have to fix three independent parameters and use the fourth one to run through all possible scenarios. The two first fixed parameters are $R = .1$ and $C = 1$.

For the third parameter, we set W_f defined by (10) 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, as defined by (11). We use N_f instead of ρ as main dimensionless parameter because it is strictly equivalent from the point of view of dimensional analysis, yet it gives a direct meaning to the variable (average number of neighbors in the fluid model).

The remaining input parameters of the system are then completely defined:

$$\lambda = \frac{N_f}{\pi R^2 W_f}, \quad F = \frac{2N_f C W_f}{R} \quad (21)$$

We choose to use a discrete time simulator, with elementary time step set to $\tau = 1$. With our settings, the resulting step transitions are empirically small enough for the discrete model to be a good approximation of the continuous model. In the end, we get a simulator that achieves the needed trade-off between speed and accuracy (an event-based simulator, for instance, would give exact rendering of the continuous model but would require a lot more of computation). For each considered setting, the simulation runtime was adjusted so that about 20000 peers could be observed in stationary state. All results presented are obtained through 10 runs per setting.

4.1 Average Performance Study

We propose to start with the mean latency, which is the main performance metric here. We made simulations for N_f varying from 1/32 to 64. Results are displayed Figure 1.

The empirical results are compared with 1) the fluid limit defined by (10) (constant within our setting), 2) the hard-core limit defined by (16), and 3) the heuristic formula (18).

Figure 1 allows us to check the conjectures in one look:

- the fluid limit latency is a lower bound of the actual mean latency (which is equivalent to Conjecture 1);
- as N_f goes to ∞ , the fluid bound becomes tight (this is Conjecture 2);
- as N_f goes to 0, the mean latency converges towards the hard-core mean latency (this is Conjecture 3).

Additionally, one checks that the heuristic (18) gives a good approximation of the mean latency for intermediate values of N_f , while converging to the hard-core and fluid limits when N_f goes to 0 and ∞ respectively.

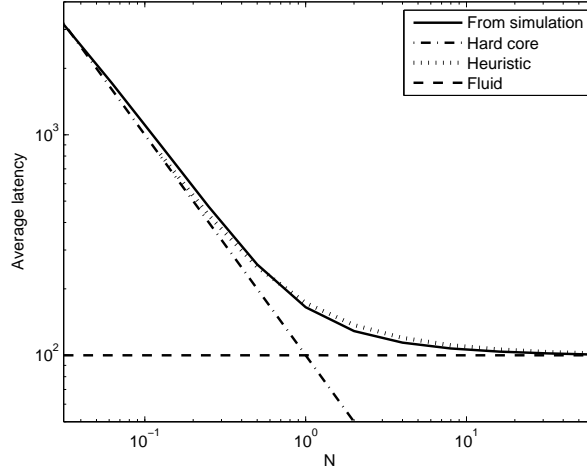


Figure 1: Mean latency in function of N_f .

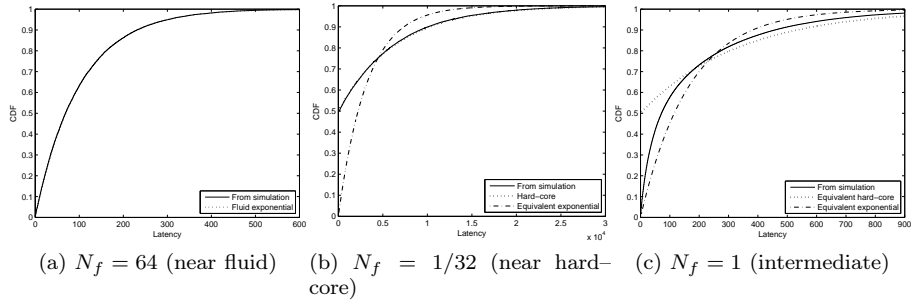


Figure 2: CDF of latency for several values of N_f

4.2 Fluid Model

We now propose to focus on the case $N_f = 64$, in order to analyze the system in detail when it reaches the fluid limit. The mean latency given by simulations is 100.7, which is higher than W_f yet very close to it, as predicted by Conjecture 2.

If one looks at the latency distribution, it is almost indistinguishable from an exponential distribution of mean W_f (Figure 2a) as predicted by Conjecture 2.

In the fluid model, it is quite difficult to distinguish the system from a spatial birth and death process of birth parameter λ and death parameter $1/W_f$, namely a Poisson point process of intensity β_f . Differences can only be spotted if small distances are involved. Indeed, Equation (14) predicts that $\frac{R}{2N_f}$ is a critical distance above which the Palm effects become less and less visible. For $N_f = 64$, this gives $\frac{R}{128} \approx 8.10^{-4}$.

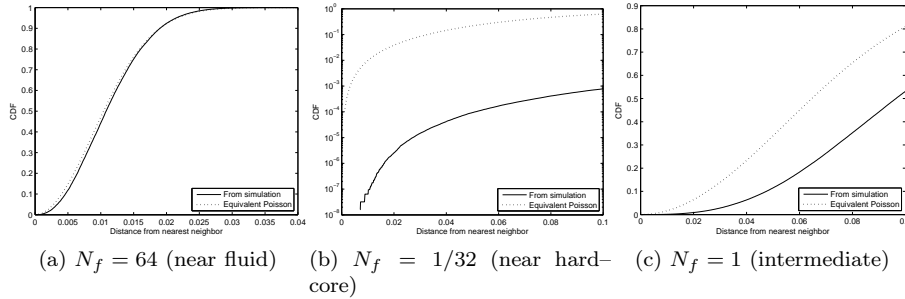


Figure 3: CDF of nearest neighbor distance for several values of N_f (reminder: $R = 0.1$)

In our case, the best way to differentiate the actual process from a Poisson process is to consider how far the closest neighbor of a peer is. While for a Poisson process the distance should be $\frac{1}{2\sqrt{\lambda W_f}} \approx 0.0111$ in average, simulation shows an actual average distance of 0.0115: the nearest neighbor is slightly farther away by about 4%. If we go into detail by comparing the two distributions, it appears that the main gap appears for small distances (cf. Figure 3a). Through this difference, we get a glimpse of how repulsion acts: if a peer gets a very close neighbor, both rates will be higher than usual, so one of them is likely to leave sooner, lowering the probability of finding very close neighbors in a random configuration. As N_f tends towards ∞ , we expect this difference to become negligible: the probability to get a neighbor so near that it will significantly affect the total rate becomes arbitrary low, so the repulsion effect becomes negligible.

4.3 Hard-Core Model

We conduct the same type of detailed study for $N_f = 1/32$. For these parameters, the mean latency obtained by simulation is now $W_o = 3160$. In comparison, the hard-core model predicts a latency $W_h = \frac{1}{\lambda\pi R^2} = \frac{W_f}{N_f} = 3200$; so the accuracy of the model is pretty good. The distribution itself is very close to the one predicted by the hard-core model: when a peer arrives, with probability one half, it disappears instantly; otherwise it follows an exponential distribution of average $2W_h$. This corresponds to the distribution function (17) of Conjecture 3.

Figure 2b displays the empirical distribution and the one described by (17). The exponential distribution of parameter W_f is also displayed for comparison, illustrating the key feature of the hard-core: while half of peers are instantly served, the other half needs to wait a rather long time. Not only the mean latency is much larger than in the fluid model (by a ratio $\frac{1}{N_f}$ according to the limit model), but half of the peers will get a service time much larger compared to the other half (as N_f goes towards 0).

The distribution of the closest neighbor is also of interest (cf. Figure 3b); the distribution has been truncated to the maximal distance R , as a peer does not “see” beyond R .

We see here the repulsion effect at its paroxysm: there are many orders of magnitude between the empirical distribution and the equivalent Poisson distribution. For instance, Poisson says that the probability to have at least one neighbor in range is $1 - e^{-\lambda\pi WR^2} \approx 62.6\%$. In the stationary regime, this probability is only 0.078%, whereas the hard-core conjecture tells us that it will continue to decrease as N_f goes to 0.

Let us conclude by mentioning that the observed distributions indicate that the hard-core process is already a good approximation at $N_f = 1/32$.

4.4 Intermediate Values

We have no good formal description of the laws observed for intermediate values of N_f . Therefore, the best thing we can currently do is to describe through simulations what happens between the hard-core and the fluid limits. To that purpose, Figure 2c shows the latency distribution for $N_f = 1$. It is a kind of compromise between the equivalent exponential and hard-core distributions obtained using the measured mean latency as parameter. Figure 3c shows the closest neighbor distribution. We can observe the repulsion effect rising.

4.5 Summary of Simulations

For both the fluid and hard-core limits, we have a good description of the performance of the system, including the latency distribution. For intermediate states, although the bounds still hold, it is better to rely on the heuristic, which gives quite accurate results on the mean latency, but with no details on the distribution. This allows us to describe the behavior of the system for all possible parameters.

5 Consequences on the Design of P2P Systems

5.1 Impact of Parameters

The simulations we conducted gave us the latency as a function of N_f . Let us define

$$M(N_f) = \frac{W_o(N_f)}{W_f}. \quad (22)$$

Dimensional analysis tells us that the function M is not affected by specific choices of the input parameters; it only depends on the resulting N_f . Hard core and fluid limits tell us that M is like $\frac{1}{N_f}$ when N_f is small, and tends to 1 as N_f goes to ∞ . Note that as $W_f = 100$ in the presented simulations, M is proportional (with a factor 1/100) to the latency displayed in Figure 1. Notice that our simulations hence show that the function M is *decreasing*.

The function M allows us to give the mean latency for any set of input parameters λ, F, C, R :

$$W_o(\lambda, F, C, R) = M \left(\sqrt{\frac{\pi\lambda FR^3}{2C}} \right) \sqrt{\frac{F}{\lambda 2\pi CR}}. \quad (23)$$

With Equation (23), we can give a complete description of the impact of the parameters on system performance. We focus here on the following question: what is the impact of each parameter on the system performance?

While having low mean latency is probably the most important aim, we should also keep in mind that the hard core regime should also be avoided, as it creates unfairness because of the two distinct service times separated by several orders of magnitude.

With that in mind, we deduce from a simple analysis of (23) based on the monotonicity of M that we can split the parameters into two categories:

- λ and R are both *win-win* parameters. As they increase, both terms in the RHS of (23) decrease and the mean latency hence tends towards 0, while the behavior of the system becomes more and more fluid. Therefore, if the values of these two parameters can be adjusted, they should be made as high as possible. Regarding λ , it means that, under our framework, P2P systems are *supra-scaling*: their performance increases with the number of users (λ is the mean number of users joining per unit time and space). To compare with, classical P2P systems exhibit *usual scaling*: their performance remains the same (or slowly decreases) with the number of users.
- F and C are *trade-off* parameters. If C increases, or if F decreases, the latency still decreases (as shown by simulations); however this is at the price of increasing the hard core behavior of the system. If these parameters can be adjusted, aiming at an intermediate state between the hard core and the fluid limits is probably the best thing to do: going too far in the fluid limit will only increase the latency with no noticeable effect on the distribution; going too far to the hard core limit will bring no improvement, as the latency reaches the asymptotic value $\frac{1}{\lambda\pi R^2}$, which is independent of F and C .

5.2 Asymptotic Behavior for Large Arrival Rates

Here, we focus on the general α case introduced in Section 2. We recall that this model makes of R a function of β of the form $R = \kappa\beta^{-\alpha}$. We think of all parameters fixed and let λ tend to infinity. We assume we are in the fluid regime (which will lead to some restrictions on the set of parameters).

In view of (12), we will call $d = \frac{1}{2-\alpha}$ the *density exponent*, and $l = \frac{\alpha-1}{2-\alpha}$ the *latency exponent*. We have the conservation rule $d - l = 1$, which is just a rephrasing of Little's law. Similarly $N_f = K\lambda^{\frac{1-2\alpha}{2-\alpha}}$, with K a constant. So, for λ tending to ∞ , the fluid regime requires that either $\alpha < \frac{1}{2}$ or $\alpha > 2$.

Hence, there are 2 regimes when $\lambda \rightarrow \infty$:

- For $\alpha < \frac{1}{2}$ (which encompasses our main case $\alpha = 0$), one gets $d > 0$ and $l < 0$, which means a peer density that tends to infinity and a latency which tends to zero when λ tends to ∞ . This is the *swarm-flash* regime.
- For $\alpha > 2$, one gets at the same time $d < 0$ and $l < 0$, which means a peer density and a latency which both tend to 0 when λ tends to ∞ . This is a rather surprising regime: the load per unit time and space tends to infinity; the density tends to 0 (there are no peers around for delivering

service); nevertheless, latency tends to 0 (i.e. when a peer arrives, it is instantly served by invisible peers located at infinity). We will call this regime *Heaven's-flash*.

Note that $\alpha = 1/2$ appears as a critical case where the structural properties of the system do not depend on λ anymore (up to a proper scaling), as shown by dimensional analysis.

6 Extensions of the Basic Model

The aim of this section is to show that our analysis can be extended in several ways and take important practical phenomena into account. Unless otherwise stated, we will place ourselves in the fluid regime, with $\alpha = 0$.

6.1 Abandonment

Here we consider the case where all leechers have some abandonment rate. Let a denote this rate. In the stationary state, we have $\lambda = (\frac{\mu_f}{F} + a)\beta_f$. From (9), we deduce $\mu_f^2 + \mu_f a F = 2\pi RC \lambda F$. The positive solution of this equation is

$$\mu_f = \sqrt{2\pi RC \lambda F + \left(\frac{aF}{2}\right)^2} - \frac{aF}{2}. \quad (24)$$

The analysis can hence be extended without difficulties. For instance, the abandonment ratio is given by $\frac{aF}{\mu_f + aF}$.

6.2 More General Rate Functions

All our results can easily be generalized to any rate function. We just have to replace (C/r) in (9) by a suitable rate function $g(r)$. This gives

$$\mu_f = \beta_f \gamma, \text{ with } \gamma = 2\pi \int_{r=0}^R g(r) r dr. \quad (25)$$

Once γ is known, we can generalize (10) by

$$\beta_f = \sqrt{\frac{\lambda F}{\gamma}}, \mu_f = \sqrt{\lambda F \gamma}, W_f = \sqrt{\frac{F}{\lambda \gamma}}. \quad (26)$$

Let us illustrate this method with a few concrete examples.

6.2.1 Affine RTT

If g is given by (2), then we have

$$\gamma = 2\pi \int_{r=0}^R \frac{C}{r+q} r dr = 2\pi C \left(R - q \ln \left(1 + \frac{R}{q} \right) \right). \quad (27)$$

6.2.2 Overhead

For g as in (3), after noticing the necessary condition $R \leq \frac{C}{c}$ (each connection needs to use a minimal bandwidth c for the overhead), we get

$$\gamma = 2\pi \int_{r=0}^R \left(\frac{C}{r} - c \right) r dr = 2\pi \left(RC - \frac{R^2 c}{2} \right). \quad (28)$$

The best value for R is $R = \frac{C}{c}$, which gives $\gamma = \pi C^2 / c$.

6.2.3 Per Flow Rate Limitation

The protocol or some physical constraints may limit the individual rates. If one assumes a maximal rate U for each flow, we have $g(x) = (C/x) \wedge U$. This gives

$$\gamma = 2\pi \int_{r=0}^R \left(\frac{C}{r} \wedge U \right) r dr = \begin{cases} \pi UR^2 & \text{if } C \geq UR \\ \pi \left(2CR - \frac{C^2}{U} \right) & \text{otherwise.} \end{cases} \quad (29)$$

We find back (10) as a special case of (26) for $U \equiv \infty$.

6.3 Per Peer Rate Limitation

Due to the asymmetric nature of certain access networks (e.g. ADSL), the uplink rate is often the most important access rate limitation. If U denotes the upload capacity of each peer (the meaning of U is different here and above), then in the fluid limit, the mean rate obtained by a typical peer is

$$\mu_f = \sqrt{\lambda F 2\pi C R} \wedge U. \quad (30)$$

A natural dimensioning rule is then be to choose $R = \frac{U^2}{\lambda F 2\pi C}$ in order for each peer to use all the available capacity (or equivalently to minimize its latency).

6.4 Leechers and Seeders

We now consider that when a leecher has obtained all its chunks, rather than leaving, it becomes a seeder and remains such for a duration T_S . In this setting, there is a density of seeders λT_S in the stationary regime.

6.4.1 Fluid limit

With seeders, (9) becomes

$$\mu_f = (\beta_f + \lambda T_S) 2\pi C R. \quad (31)$$

Using (5) and $F = W_f \mu_f$, we get

$$W_f^2 + W_f T_S = W_{f_0}^2, \text{ with } W_{f_0} = \sqrt{\frac{F}{\lambda 2\pi C R}}. \quad (32)$$

The positive solution of this equation is

$$W_f = \sqrt{W_{f_0}^2 + \left(\frac{T_S}{2} \right)^2} - \frac{T_S}{2}. \quad (33)$$

In particular, we have $W_f \approx W_{f_0}$ for $T_S \ll W_{f_0}$ and $W_f \approx \frac{W_{f_0}^2}{T_S}$ for $T_S \gg W_{f_0}$.

By comparing (33) and (24), one can interpret seeding as the exact opposite of abandonment: seeders, which improve the system, impact the latency the same way that abandonment, which degrades the system, impacts the rate.

We also remark that in the fluid model where rates are only determined by the upload access and the down link is infinite, we have (see [BMHP08] for details)

$$W_f + T_S = W_{f_0}. \quad (34)$$

We can see (32) as the spatial-dependent version of (34).

6.4.2 Hard-core limit

In the hard-core limit without seeder, a leecher may complete its download only if it finds a peer within range, and instant service happens with probability one half. With seeders, a leecher is certain to complete its download if there is another peer in its neighborhood, as the latter will not leave the system before the former finishes. We can then notice that the configuration of peers (leechers and seeders) includes a spatial Poisson distribution of density λT_S . In particular, the probability for a newcomer to find a peer within range R is at least $1 - e^{-\lambda T_S \pi R^2}$. Therefore, for any $\epsilon > 0$, if $T_S \geq \frac{-\log(\epsilon)}{\lambda \pi R^2}$, then leechers will get instant service with a probability greater than $1 - \epsilon$.

This indicates that seeders may be the perfect solution for systems where a hard-core behavior cannot be avoided: a seeding time of the same order of magnitude than the average latency in absence of seeders is enough to guarantee that most of the peers get instant download.

6.5 Mixed extensions

Although we presented them separately for sake of clarity, our extensions can be interleaved. For instance, if we consider a system with seeders and a rate function parameter γ (cf 6.2), one can verify that we obtain the average latency

$$W_f = \sqrt{\frac{F}{\lambda \gamma} + \left(\frac{T_S}{2}\right)^2} - \frac{T_S}{2}, \quad (35)$$

which is a sort of combination of (26) and (33).

7 Conclusion

There is a clear need for comprehensive P2P models beyond the general assumption that rates are shaped by the access link only. The present paper is a first step trying to answer this need. This leads to a new class of models which we find quite challenging in terms of mathematics and quite intriguing in terms of their practical implications.

The mathematical challenges are well summarized by the three conjectures listed in the paper. Properly proving them will require to establish the existence of the stationary regime, and to develop bounding and asymptotic analysis techniques.

In terms of practical results, we gave a first, yet detailed, insight of the behavior of spatial-dependent P2P systems, with some unexpected findings like for instance the possibility of supra-scalability.

Such regimes are of course to be understood in an asymptotic sense, but if the conjectures supported by our simulations are confirmed, new guidelines and dimensioning rules will clearly be at hand for future P2P overlays.

References

- [BB09] François Baccelli and Bartłomiej Błaszczyszyn. *Stochastic Geometry and Wireless Networks, Volume I — Theory*, volume 3, No 3–4 of *Foundations and Trends in Networking*. NoW Publishers, 2009.
- [BMHP08] Farid Benbadis, Fabien Mathieu, Nidhi Hegde, and Diego Perino. Playing with the bandwidth conservation law. In *IEEE P2P*, pages 140–149, 2008.
- [Buc15] E. Buckingham. The principle of similitude. *Nature*, 96 (2406):396–397, 1915.
- [Coh06] B. Cohen. BitTorrent specification, 2006. <http://www.bittorrent.org>.
- [DVJ88] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer, 1988.
- [FLV08] Pierre Fournier, Emmanuelle Lebhar, and Laurent Viennot. The inframetric model for the internet. In *IEEE INFOCOM*, pages 1085–93, 2008.
- [GK06] N.L. Garcia and T.G. Kurtz. Spatial birth and death processes as solutions of stochastic equations. *ALEA Lat. Am. J. Probab. Math. Stat.*, 1:281–303, 2006.
- [HZ10] B. Hajek and J. Zhu. The missing piece syndrome in peer-to-peer communication, 2010. <http://arxiv.org/abs/1002.3493>.
- [KPK⁺10] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguńa. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82(3):036106, Sep 2010.
- [MR06] Fabien Mathieu and Julien Reynier. Missing piece issue and upload strategies in flashcrowds and P2P-assisted filesharing. In *AICT/ICIW'06*, 2006.
- [MV05] L. Massoulié and M. Vojnovic. Coupon replication systems. *IEEE/ACM Trans. Networking*, 16(3):603–616, 2005.
- [Nor09] A. Norberg. Bittorrent enhancement proposals on uTorrent transport protocol, 2009. http://bittorrent.org/beps/bep_0029.html.

- [OAN11] B. Oğuz, V. Anantharam, and I. Norros. Stable, scalable, decentralized P2P file sharing with non-altruistic peers. 2011. arXiv:1107.3166v1.
- [OKM92] T. Ott, J.H.B. Kemperman, and M. Mathis. The stationary behavior of ideal TCP congestion avoidance. *Internetworking: Research and Experience*, 11:115–156, 1992.
- [Pre75] C. Preston. Spatial birth-and-death processes. *Bull. Inst. Internat. Statist.*, 46(2):371–391, 405–408, 1975.
- [QS04] D. Qiu and R. Srikant. Modeling and performance analysis of BitTorrent-like peer-to-peer networks. *ACM SIGCOMM Computer Communication Review*, 34(4):367–378, 2004.
- [Rei09] H. Reittu. A stable random-contact algorithm for peer-to-peer file sharing. In *IFIP IWSOS*, pages 185–192, 2009.
- [SAV06] Riikka Susitaival, Samuli Aalto, and Jorma Virtamo. Analyzing the dynamics and resource usage of P2P file sharing systems by a spatio-temporal model. In *P2P-HPCS06, in conj. with ICCS*, pages 420–427, May. 2006.
- [Sha10] S. Shalunov. Low extra delay background transport (LEDBAT). IETF Draft, 2010. <http://datatracker.ietf.org/wg/ledbat/charter/>.



Centre de recherche INRIA Paris – Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399