

Strategic Reasoning About Bundling in Swarming Systems

Daniel Sadoc Menasche, Giovanni Neglia, Don Towsley, Shlomo Zilberstein

Abstract—The objects of study of this paper are swarming systems, a special kind of peer-to-peer systems where users interested in the same content at the same time cooperate with each other. In particular, we consider the problem of how to combine files into bundles in such systems. First, we analyze the case of a monopoly where a single publisher decides how to aggregate its files so as to satisfy user demands while mitigating its serving costs. We establish conditions for the existence and uniqueness of an equilibrium and how the publisher’s bundling strategy affects its profit. Then, we consider the competitive case where bundling decisions of one publisher affect the outcome of other publishers. Using normal form games we analyze the impact of different system parameters on the Nash equilibrium.

I. INTRODUCTION

Content dissemination using swarming technology is in vogue. Nowadays, virtually any user can run a BitTorrent tracker and publish files to the rest of the world. In applications implementing swarming technology (e.g., BitTorrent [6] and Slurpie [20]), each user interested in a file (or collection of files) joins a swarm and exchanges blocks with other users that happen to be interested in the same content at the same time. As in any peer-to-peer system, cooperation among peers guarantees high scalability and tolerance to flash crowds since system capacity scales with the number of peers.

Reliance exclusively on intra-swarm exchanges leads to a simple yet robust incentive mechanism where a peer does not need to store history about others and can immediately stop sending packets to those that do not reciprocate. In the realm of peer-to-peer systems, this constitutes one of the advantages of swarming systems over their credit based counterparts. In credit based systems users need to either rely on a bank (single point of failure) or store information in local files about credits received or given (prone to hacking). In addition, the fact that in existing swarming systems trackers don’t exchange control information with each other allows any user to, independently, at any time, publish its contents.

Daniel Sadoc Menasche, Don Towsley and Shlomo Zilberstein are with the Department of Computer Science, University of Massachusetts, Amherst, {sadoc, towsley, shlomo}@cs.umass.edu
Giovanni Neglia is with the Maestro Team at INRIA Sophia Antipolis, giovanni.neglia@sophia.inria.fr

Despite all the advantages of the swarming technology, for unpopular or very small files publishers still incur costs for every received request. A critical mass of online peers is needed to leverage cooperation. In general, it is assumed that swarming technology is not suited for unpopular content.

Neglecting the long tail of the popularity curve and using swarming systems only for very popular content may not be economic efficient for some businesses. For instance, [2] points out that, in 2004, products not available in traditional stores were responsible for 22%, 57% and 20% of the sales of Rhapsody, Amazon.com and Netflix. In November of 2008, [16] reported that 20% of the revenue of Rhapsody came from songs that are not in the top 52,000. This reduction in the blockbuster effect indicates that in the field of content delivery, enterprises that can “make everything available”, at a low cost, have a significant advantage.

One option available to a content provider is to bundle (or group) unpopular files so that users downloading any one of the files cooperate with each other in order to get the desired content [14]. In the context of a typical PC user, tools such as BitTorrent already support bundles. A torrent can contain either a single file or a bundle of files. If a torrent consists of a bundle of files, BitTorrent’s interface allows users joining the swarm to give priorities to different files or even to download a subset of them. Nevertheless, many users can be expected to download more files than they sought out of curiosity. This option is referred to as mixed bundling, as opposed to pure bundling, where publishers create an archive (.zip, .rar, etc.) and users interested in any one of the files have to download all the package.

Enterprises can also take advantage of bundling in order to leverage swarming systems for the dissemination of content. For that purpose, the swarming system (e.g., BitTorrent) needs to be coupled with credit and copyright mechanisms (e.g., DRM [24]), which ensure that users gain access to a file only after paying for its access. In this case, users download bundles but have limited access to the files. In order to obtain full access, they need a passcode, which is made available only after payment.

The idea of using bundling to promote unpopular content is not new. In the economics literature bundling was proposed decades ago as a mechanism to increase sales and extend monopoly power. For information goods, bundling has been shown to be advantageous because it permits firms

to smooth demand for multiple goods. Traditionally it has been assumed that information goods have zero marginal cost for production and dissemination. In this paper, however, we are particularly interested in the impact of these costs on the publishers, which may be non negligible if, for instance, we consider publishers of a large number of unpopular contents.

When multiple enterprises are involved in the market, the bundling strategy of one enterprise impacts the outcome of the others. Note that different firms may offer overlapping, partially overlapping (or weak substitutable), or non overlapping contents. In this multi-firm scenario, game theory emerges as a natural tool to study the possible market outcomes.

In this paper, we show how and when bundling, coupled with a swarming mechanism, helps a publisher increase utility. In particular, we are interested both in the monopoly and the duopoly cases. We ask the following questions:

- 1) in the case of a monopoly, for a given bundling strategy does there exist an equilibrium and, if so, what is it?
- 2) how does the monopoly equilibrium change as functions of the bundling strategy?
- 3) in the case of a duopoly, what is the impact of the parameters of the system on the Nash Equilibrium?

In this work we assume that users have uniform valuations over files and that each user, when accessing the system, is interested in a single file. Users are reluctant to download bundles since the download of the bundle may lead to increased delays. Publishers, on the other hand, are willing to push bundles to the users since this increases the availability of the contents. We propose expressions to capture the utility of users and publishers as functions of the prices and bundling strategies. Using this model, we derive properties of the system equilibrium and perform a sensitivity analysis of the publishers' utility as a function of the bundling strategy.

For the duopoly, we consider a special parameterization of the utility functions for users and publishers and plug them into normal form game matrices (for which the Nash equilibrium is guaranteed to exist). Using these matrices, we show how different system parameters such as willingness of users to download bundles and level of overlap of contents between publishers affect the Nash equilibrium.

We provide the following answers to our initial questions:

- for the monopoly, we show conditions for existence and uniqueness of an equilibrium;
- we present scenarios under which the optimal bundling level can be easily determined. In particular, if the cost to the publisher is proportional to the number of bundles it offers and does not depend on demand, the publisher's utility as a function of the bundling level has a unique local maximum which can be easily determined using a gradient descent strategy;

- for the duopoly, we show that different games emerge as a function of the system parameters. This multitude of possible games may be regarded as a sign of the complexity of the problem in hand.

II. THE MODEL AND PROBLEM STATEMENT

We consider S publishers that offer C contents (files) to an infinite population of users. Each publisher s provides N_s files through T_s torrents, in bundles of size $K_s = N_s/T_s$. Each file provided by a publisher is a member of exactly one of its bundles.

We do not restrict K_s to take integer values. If the optimal value K^* is not an integer, one way to implement it is to let the publisher's bundling strategy change over time. The optimal value K^* can be used to decide the fraction of time that $\lceil K^* \rceil$ and $\lfloor K^* \rfloor$ are adopted.

Each user accessing the system is interested in a single file. Users requesting content i arrive with rate $\Lambda^{(i)}$ but only a fraction $x^{(i)}$ of the users download the desired content.

Let $i_s(l)$ be the index of the l^{th} content provided by s . If two publishers p and q offer the same content, $i_p(m) = i_q(n)$ for some $1 \leq m \leq N_p$ and $1 \leq n \leq N_q$. Publisher s experiences demand $\lambda_s^{(i_s(l))} = \Lambda^{(i_s(l))} x_s^{(i_s(l))}$ for its l^{th} content. Note that $\sum_{s=1}^S x_s^{(i)} = x^{(i)}$.

The vector $\underline{\lambda}_s = (\Lambda^{(i_s(1))} x_s^{(i_s(1))}, \dots, \Lambda^{(i_s(N_s))} x_s^{(i_s(N_s))})$ is referred to as the *publisher s demand* vector. $\underline{\lambda} = (\underline{\lambda}_1, \dots, \underline{\lambda}_S)$ is the *population demand* vector and $\underline{\Lambda} = (\Lambda^{(1)}, \dots, \Lambda^{(C)})$ is the vector of *popularities*. The vector of *bundling strategies* is given by $\underline{K} = (K_1, \dots, K_S)$.

Our key assumptions are the following:

- (A1) (reluctance to bundling) users are reluctant to download unrequested content;
- (A2) (minimize download time) users want to minimize the time to download requested content;
- (A3) (self sustaining swarms) users that download the same file or bundle from the same publisher help each other; if enough users collaborate with each other, the swarm sustains itself even in the absence of the publisher;
- (A4) (bandwidth costs) if the swarm is not able to sustain itself, the publisher needs to allocate a minimum bandwidth to each user;
- (A5) (rewards for serving users) the publisher receives a reward (which may be monetary) for each user served;
- (A6) (infinite capacity) the publisher bandwidth capacity is infinite.

A. Users' Utility

To capture assumptions (A1) and (A2), we model the cost function of a typical user as the sum of its download and

monetary costs. For a user that receives content at rate μ bps, the download cost of a bundle of K_s files from publisher s , $D(K_s)$, is given by

$$D(K_s) = \alpha K_s / \mu \quad (1)$$

Here the term K_s/μ captures the average download time and α the download cost to a user in units of currency per second. α accounts for the reluctance of users to download bundles. The smaller the value of α the longer the users are willing to wait in order to download the desired files. Note that both μ and α can in principle vary across users, but here we assume a homogeneous population.

Publishers charge a flat price for each file that they provide. Publisher s sets price p_s for each of its contents and the *vector of prices* is $\underline{p} = (p_1, \dots, p_S)$.

The user's utility function for the download of a bundle of K_s files (one of them being the file requested) from publisher s , $U_c(s)$, is

$$U_c(s) = V - D(K_s) - p_s \quad (2)$$

Here V is a random variable uniformly distributed in the range $[v_{min}, v_{max}]$, $0 \leq v_{min} < v_{max} < \infty$, which characterizes valuations of files by users. The valuations are assumed to be independent across users.

Users download the desired content from s if (a) that publisher has the content, (b) $U_c(s) > 0$ and (c) $U_c(s) \geq U_c(s')$ for all $s' \neq s$. In case of a tie, demand is equally split among publishers. Clearly, the pricing and bundling strategies of a publisher impacts the demand (hence utility) experienced by the others. This naturally characterizes a game between publishers, who are coupled through users demands. To describe this game, we introduce the function $q_s^{(i_s(l))}(\cdot)$ which determines the demand experienced by publisher s for its l^{th} content as a function of all publishers' strategies

$$q_s^{(i_s(l))}(\underline{p}, \underline{K}) = \lambda_s^{(i_s(l))} = \Lambda^{(i_s(l))} x_s^{(i_s(l))} \quad (3)$$

We present two possible definitions of $q_s^{(i_s(l))}(\cdot)$ in Sections III and IV (equations (8) and (28)).

Note that even though users download bundles, they only have access to the requested file. Hence, the download cost in (2) is proportional to K_s but the price charged by the publisher is for the single file.

B. Publishers' Utility

The aggregate costs incurred by publisher s to make all of its N_s/K_s bundles available must capture assumptions (A3) and (A4) and are given by $G_s(\underline{\lambda}_s, K_s)$. In some cases it may be convenient to make explicit the cost for publisher s to make its j^{th} bundle available, $C_s^{(j)}(\underline{\lambda}_s, K_s)$. The definition of $C_s^{(j)}(\underline{\lambda}_s, K_s)$ is postponed to Section III.

The publishers' revenue, $F_s(\underline{\lambda}_s, K_s)$, related to assumption (A5), is the standard one used in the economics literature:

$$F_s(\underline{\lambda}_s, K_s) = p_s \sum_{l=1}^{N_s} \Lambda^{(i_s(l))} x_s^{(i_s(l))} = p_s \sum_{l=1}^{N_s} \lambda_s^{(i_s(l))} \quad (4)$$

The publisher's utility, $U_s(p_s, K_s)$, is given by the difference between its revenue and its costs

$$U_s(p_s, K_s) = p_s \sum_{l=1}^{N_s} \lambda_s^{(i_s(l))} - G_s(\underline{\lambda}_s, K_s) \quad (5)$$

$$= p_s \sum_{l=1}^{N_s} \lambda_s^{(i_s(l))} - \sum_{j=1}^{T_s} C_s^{(j)}(\underline{\lambda}_s, K_s) \quad (6)$$

C. Problem statement

The objective of publisher s is to maximize its utility,

$$\max_{K_s, p_s} U_s(p_s, K_s) = \max_{K_s, p_s} F_s(\underline{\lambda}_s, K_s) - G_s(\underline{\lambda}_s, K_s) \quad (7)$$

subject to the constraints posed by (3).

Variable	Description
Client Parameters	
V	value of a file, uniform random variable [0,1]
$D(K)$	cost of download for the user
μ	download rate
Publisher Parameters	
$U(\lambda, K)$	utility (profit)
$C(\lambda, K)$	cost per bundle
$p(\lambda, K)$	price
K	number of files per bundle
N	number of files
Population Parameters	
Λ	total arrival rate per file
x	fraction of users that buy a file
$x_0 = 1 - (\alpha K / \mu)$	fraction of users that buy a file when $p = 0$
$\lambda = \Lambda x$	demand per file

TABLE I

TABLE OF NOTATION FOR THE MONOPOLY.

III. MONOPOLY

In this section we study the single publisher case. We are interested in understanding the impact of the publisher pricing strategy on the equilibrium, defined as the point where the publisher's profit is maximized. Henceforth, the index s is dropped for notational convenience.

A. Optimization Problem

Our goal is to pose a simplified version of the general problem (7). For that purpose, we assume that the popularities of all the files are the same ($\Lambda^{(i)} = \Lambda$ for all i).

A user downloads a bundle and purchases a file if $v - D(\lambda, K) - p > 0$ (equation (2)), where v is a realization of the random variable V . Letting $v_{min} = 0$ and $v_{max} = 1$, V is uniformly distributed in $[0,1]$ (Section II-A).

Hence, the fraction x of the population for which $v > D(K) + p$ is $1 - g$ where $g = D(K) + p = (\alpha K/\mu) + p$. The demand as a function of the price is given by

$$x(p, K) = \frac{q(p, K)}{\Lambda} = [1 - \alpha K/\mu] - p = x_0 - p \quad (8)$$

where $q(\cdot)$ was introduced in (3). We have dropped all the superscripts present in (3) because the popularities of all the files are now assumed to be the same. Also, we explicitly account for the dependence of x on p and K . Note that when the price is zero the demand is maximized and $x_0 = 1 - \alpha K/\mu$. Alternatively, the price as a function of the demand is

$$p(x, K) = (1 - \alpha K/\mu) - x, \quad (9)$$

which implies that the maximum feasible bundle size is μ/α . Since demand x and price p are coupled through (8) and (9) we take K and x as control variables rather than K and p . Substituting (9) into (7) yields the following objective function for the publisher

$$\max_{K, x} U(x, K) \text{ where } U(x, K) = N\Lambda p(x, K)x - \frac{N}{K}C(\lambda, K) \quad (10)$$

To solve the problem, we must determine the cost function $C(\lambda, K)$. To illustrate, in the following section we consider three possible functions, the first being the most realistic but also the most complex.

B. Cost Functions

The goal of this section is to determine the cost function $C(\lambda, K)$. For this purpose we rely on the model proposed and analyzed in [14] for content availability in swarming systems. The model proposed in [14] captures, in particular, the benefits of bundling and for the sake of completeness we briefly introduce it.

1) *Impact of Bundling:* Aggregating K contents together and offering them as a single bundle has a twofold impact on the system parameters: (a) the arrival rate of peers requesting the bundle is K times larger than the arrival rate of requestors for each individual file and (b) the active download time experienced by each peer is K times larger than the one experienced by peers downloading individual files. Assuming that peers arrive according to a Poisson process with rate $K\lambda$ and remain online for an exponentially distributed period of time equal to K/μ , and denoting by X the number of peers in the system [14]:

$$P\{X = n\} = \frac{(K^2\lambda/\mu)^n e^{-K^2\lambda/\mu}}{n!} \quad (11)$$

2) *Cost Functions:* We now briefly present the three cost functions studied in this paper. Please, refer to [14] for a full treatment.

Active publisher: The publisher must be active when the number of peers in the system is small so as to provide missing blocks. In particular, assuming that the publisher is active only when there is one peer in the system, the probability of finding an *active publisher* is obtained by setting $n = 1$ in (11). The publisher's cost function per swarm is given by

$$C(\lambda, K) = cK^2 \frac{\lambda}{\mu} e^{-K^2\lambda/\mu} \quad (12)$$

where c represents the cost per time unit incurred by an active publisher.

Busy period: The previous cost function is applicable if the server is active when there is *exactly one* customer in the system. For a second cost function, we assume that the publisher will be active whenever there is *at least one* leecher in the system [14].

$$C(\lambda, K) = cP\{X > 0\} = c(1 - e^{-K^2\lambda/\mu}) \quad (13)$$

Note that this cost function is concave in λ . For further discussion on the importance of the busy period as a metric in swarming systems, please refer to [14].

Number of torrents: Finally, the last cost function we consider is where the per swarm cost to the publisher is constant

$$C(\lambda, K) = c \quad (14)$$

3) *Numerical Evaluation:* In Figure 1 we illustrate the first two cost functions described above. In all cases $c = \mu$, which means that the costs are proportional to the allocated bandwidth, which in turns equals that of a typical peer. We start with the configuration $N = 10$, $K = 1$ and $\mu = 8$ Mbps. Figure 1(a) shows the *active publisher* cost (Section III-B.2) as a function of λ , as K varies from 1 to 4. For a fixed value of K , the costs for the server first increase but then sharply decreases, which captures the fact that when the population size is above a certain threshold the swarm is self sustaining. Note that the cost decreases as K increases: bundling leads to an increase in the arrival rate and in the residence time of peers, which maps into reduced costs to the publisher. The same rationale applies to Figure 1(b) where we consider the *busy period* cost function (Section III-B.2). However, in Figure 1(b) costs are bounded away from zero as λ increases. Finally, in Figure 1(c) and (d) we observe that as μ increases server costs also increase. That's because the server contributes with a capacity equal to that of a typical peer ($c = \mu$).

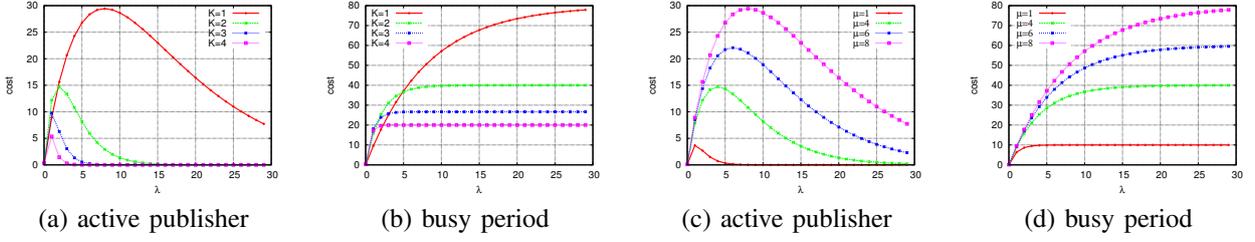


Fig. 1. Cost for the publisher as a function of the arrival rate. (a) and (b) varying K [$\mu = 1$]; (c) and (d) varying μ [$K = 1$].

C. Equilibrium for Fixed Bundling Strategy

We now establish the existence and uniqueness of the equilibrium in the monopoly case. Since in this section we consider a fixed value of K , we don't make explicit the dependency of any of the functions on K .

Definition: An equilibrium is a point x^* that maximizes the utility (10), i.e.:

$$x^* = \arg \max_{x \in [0, x_0]} U(x) = \arg \max_{x \in [0, x_0]} \{N\Lambda p(x)x - \frac{N}{K}C(x\Lambda)\} \quad (15)$$

In this section we assume that $C(x)$, as well as its first three derivatives, are continuous, $C(0) = 0$ and $C(x) > 0$ for $x > 0$. Moreover we assume that there exists at least one value $y \in [0, x_0]$ such that $U(y) > 0$ and that $U(x)$ is not constant in any interval contained in $[0, x_0]$.

Proposition 3.1: The monopoly has an equilibrium with positive utility in $(0, x_0)$.

Proof: $U(x)$ is continuous in the interval $[0, x_0]$, so it has a global maximum x^* . At the global maximum x^* , the utility is positive ($U(x^*) > 0$). As a consequence x^* belongs to $(0, x_0)$ because $U(0) = 0$ and $U(x_0) < 0$. ■

The *publisher's revenue*, $F(x)$, is given by (4), $F(x) = N\Lambda p(x)x = N\Lambda[(1 - \alpha K/\mu) - x]x$ and the *publisher's marginal revenue* is $\frac{dF(x)}{dx} = N\Lambda[-2x + 1 - \alpha \frac{K}{\mu}]$. The *publisher's cost per swarm*, $C(x\Lambda)$, is given by one of the equations (12)-(14), the *publisher's cost* is $\frac{N}{K}C(x\Lambda)$ and the *publisher's marginal cost* is $\frac{N}{K} \frac{dC(x\Lambda)}{dx}$.

We observe that a necessary condition for y to be an equilibrium is that $U'(y) = 0$ since the equilibrium falls inside $[0, x_0]$. This means that at the equilibrium the marginal revenue has to equal the marginal cost:

$$\frac{dF(x)}{dx} \Big|_{x=x^*} = N\Lambda[-2x^* + 1 - \alpha \frac{K}{\mu}] = \frac{dC(x\Lambda)}{dx} \frac{N}{K} \Big|_{x=x^*}. \quad (16)$$

Among the roots of $U(x) = 0$ there are local minima and local maxima that are not equilibria. In what follows we provide some sufficient conditions for the existence of a single local

maximum and then of a unique equilibrium¹.

Proposition 3.2: If $U'(x) = 0$ has at most two roots in $[0, x_0]$ then there is a unique equilibrium.

Proof: If $U'(x) = 0$ has exactly one root within $[0, x_0]$ then this root is the equilibrium (and the unique maximum). If $U'(x) = 0$ has exactly two roots within $[0, x_0]$, they cannot be both local maxima points of $U(x)$. For the sake of contradiction, assume that both roots are local maxima. Then, there would have to exist a third root that minimizes $U(x)$ within $[0, x_0]$. $U'(x) = 0$ would have to have three roots within $[0, x_0]$ (contradiction). So, if $U'(x) = 0$ has exactly two roots within $[0, x_0]$, one of them is a maximum and the other is a minimum. ■

Proposition 3.3: If the publisher marginal cost function $\frac{N}{K}C'(\Lambda x)$ is convex in x then there is a unique equilibrium.

Proof: The publisher marginal cost function $\frac{N}{K}C'(\Lambda x)$ being convex in x corresponds to $C'''(\Lambda x) \geq 0$. Since $U'''(x) = -\frac{N}{K}C'''(\Lambda x)$ it follows that $U'''(x) \leq 0$. Hence $U''(x)$ is non increasing and $U'(x)$ can change sign at most twice. As we assumed that $U(x)$ is not constant in any interval, $U'(x) = 0$ has at most two roots and the result follows from Proposition 3.2. ■

We observe that the hypothesis of Proposition 3.3 is satisfied by the cost function (13), as proven in the Appendix, but not by function (12) which increases and then decreases. For this case the following result can be applied (hypotheses check for the function (12) in the Appendix).

Proposition 3.4: If the publisher cost $\frac{N}{K}C(\Lambda x)$ is concave in $[0, \bar{x}]$ and convex in $(\bar{x}, x_0]$ and the corresponding marginal cost $\frac{N}{K}C'(\Lambda x)$ is convex in $[0, \bar{x}]$ then there is a unique equilibrium.

Proof: Due to Proposition 3.3, $U'(x) = 0$ has at most two roots in $[0, \bar{x}]$. Cost convexity in $(\bar{x}, x_0]$ (i.e., marginal cost being non decreasing in the interval $(\bar{x}, x_0]$) and the fact that the utility cannot be constant in any interval implies that $U'(x) = 0$ has at most one root in $(\bar{x}, x_0]$.

We now prove that if $U'(x) = 0$ has two roots in $[0, \bar{x}]$ it has

¹Note that even in presence of multiple local maxima, there could be a unique global maximum point and so a unique equilibrium.

no roots in $(\bar{x}, x_0]$. By assumption, $C'(\Lambda x)$ is decreasing and convex in $[0, \bar{x}]$. Also, $F'(x) = N\Lambda[-2x + 1 - \alpha \frac{K}{\mu}]$. Therefore, if $U'(x) = F'(x) - C'(\Lambda x) = 0$ has two roots in $[0, \bar{x}]$ it follows that $C'(\Lambda \bar{x}) > F'(\bar{x})$. Since in $(\bar{x}, x_0]$ the marginal cost is non decreasing, $C'(\Lambda x) \geq C'(\Lambda \bar{x}) \forall x \in (\bar{x}, x_0]$. The marginal revenue $F'(x)$ is always decreasing. Therefore, $U'(x) = 0$ has no roots in $(\bar{x}, x_0]$.

Note that if $U'(x) = 0$ has at most one root in $[0, \bar{x}]$ it could have one root in $(\bar{x}, x_0]$.

It has been shown that $U'(x) = 0$ has at most two roots in $[0, x_0]$. The rest of the proof then follows from Proposition 3.2. ■

D. The Impact of Bundling

We now study the impact of bundling on the publisher's utility in equilibrium. The three cost functions introduced in Section III-B are considered.

1) *Active Publisher*: Given the active publisher cost function (12),

$$\frac{\partial}{\partial x} F(x, K)|_{x=x^*} = N\Lambda \left[-x + [(1 - \alpha K/\mu) - x] \right] |_{x=x^*} \quad (17)$$

and

$$\frac{\partial}{\partial x} \frac{N}{K} C(x, K)|_{x=x^*} = cNK \frac{\Lambda}{\mu} \left[1 - \frac{xK^2\Lambda}{\mu} \right] e^{-K^2x\Lambda/\mu} |_{x=x^*} \quad (18)$$

Substituting the two equations above into (16) yields

$$\left[-2x^* + 1 - \alpha K/\mu \right] - \frac{cK}{\mu} \left[1 - \frac{x^*K^2\Lambda}{\mu} \right] e^{-K^2x^*\Lambda/\mu} = 0 \quad (19)$$

The roots of (19), which correspond to candidate values for the demand in equilibrium, x^* , can be obtained using the following fix point equation:

$$x^{(i+1)} = \frac{1}{2\mu} \left(\left(\frac{K^2\Lambda x^{(i)}}{\mu} - 1 \right) cK e^{-\frac{K^2x^{(i)}\Lambda}{\mu}} + \mu - \alpha K \right) \quad (20)$$

Now we get insights on how α affects the system equilibrium through a numerical experiment. We set the parameters of the model as follows: $\Lambda = 1$, $\mu = 1$, $c = 12$ and $N = 100$. Adopting 0.9 as the initial condition of (20), we set α equal to 0.02, 0.01, 0.0002 and 2×10^{-7} and obtain the results shown in Table II. The utility values presented in the table were normalized and are given by $U(x^*, K) - 24$, where $U(x^*, K)$ is obtained through (10).

The same trend is observed in all the columns of Table II. For small values of K , an increase in the bundling level leads to a significant reduction in the costs for the publisher and a corresponding utility increase. However, for larger values of K increasing the bundling level leads to a decrease in demand which negatively impacts the utility.

TABLE II
EQUILIBRIUM VARYING K AND α (ACTIVE PUBLISHER COST).

α	0.02		0.001	0.0002	2×10^{-7}
K	x^*	$U(x^*, K)$	$U(x^*, K)$	$U(x^*, K)$	$U(x^*, K)$
3	0.5007	-19.18514	-18.80194	-18.49608	-18.41979
4	0.5321	-0.041375	0.171757	0.342649	0.385383
5	<u>0.4963</u>	<u>0.490159</u>	<u>0.738963</u>	0.938880	0.988931
6	0.4940	0.403533	0.700839	<u>0.939980</u>	0.999885
7	0.4930	0.304900	0.651225	0.930049	<u>0.999930</u>
8	0.4920	0.206400	0.601600	0.920064	0.999920
9	0.4910	0.108100	0.552025	0.910081	0.999910
10	0.4900	0.010000	0.502500	0.900100	0.999900

For $\alpha = 0.002$ we also show the equilibrium value of x . When K is small, the publisher sets high prices to cope with its costs, hence demand is small. As K increases, prices decrease and demand increases. Finally, increasing K further leads to a demand decrease due to the reluctance of users to download bundles.

For a fixed value of K , as α decreases we note an increase in the utility perceived by the publisher. This happens because a higher fraction of the population joins the system. The optimal bundle size for $\alpha = 0.02, 0.01, 0.0002$ and 2×10^{-7} is equal to 5, 5, 6, 7, respectively (underlined elements in Table II). Finally, note that we haven't included in Table II the results for $K = 1, 2$ because the system admits no equilibrium in those cases. For the results presented in the previous section regarding the existence of the equilibrium, we assumed that there is at least one value of x in the range $[0, x_0]$ for which $U(x) > 0$. This condition is not satisfied for $K = 1, 2$ in this example.

2) *Busy Period*: When the cost function is given by the busy period, we can proceed as in the last section and solve the problem using a fix point algorithm. However, in this case we can express the equilibrium demand, x^* , as a function of K in terms of the Lambert $W(\cdot)$ function. The Lambert $W(\cdot)$ function is defined as the inverse of $f(w) = we^w$. Since $f(w)$ is not injective $W(\cdot)$ is multivalued; nevertheless, $W(\cdot)$ takes real values only on two of its branches, referred to as $W_0(\cdot)$ and $W_{-1}(\cdot)$ [12].

The demand in equilibrium, x^* , is defined implicitly as a function of K by substituting (13) into (16). First, we define the function $H(x, K)$ as follows

$$\begin{aligned} H(x, K) &= \frac{\partial}{\partial x} F(x, K) - \frac{\partial}{\partial x} \frac{N}{K} C(x, K) \\ H(x, K) &= N\Lambda \left[-2x + 1 - \alpha \frac{K}{\mu} \right] - \frac{N}{K} c e^{-K^2x\Lambda/\mu} K^2 \frac{\Lambda}{\mu} \end{aligned}$$

Then, solving $H(x^*, K) = 0$ for x^* yields

$$x^* = \frac{\left(2W \left(-\frac{cK^3 \Lambda e^{1/2 \Lambda K^2 (-\mu + \alpha K)}}{2\mu^2} \right) \frac{\mu^2}{K^2} + \Lambda \mu - \Lambda K \alpha \right)}{2\Lambda \mu} \quad (21)$$

We make three observations about this solution. First, the argument of the function W decreases if $\frac{2\mu}{3\alpha} \leq K \leq \frac{\mu}{\alpha}$. But $W_0(z)$ and $W_{-1}(z)$ take real values if and only if $z \in [-1/e, +\infty)$ and $z \in [-1/e, 0)$, respectively [12, Lemma 2.10]. Therefore, once we find (using $W_0(z)$) a value $K' \geq \frac{2\mu}{3\alpha}$ for which x^* is an imaginary number we don't need to further search for solutions $K^* > K'$.

Second, substituting x^* into (10) and taking the derivative of (10) with respect to K yield

$$\lim_{K \rightarrow 0} \frac{dU}{dK} = -\frac{\Lambda(\alpha + c)N}{2\mu} < 0 \quad (22)$$

Third, we can compute $\frac{dx^*}{dK}$ using the implicit function theorem,

$$\frac{dx^*}{dK} = -\frac{\partial H / \partial K}{\partial H / \partial x^*} = \frac{(c\mu - 2\Lambda K^2 x^* c) e^{-\frac{K^2 \Lambda x^*}{\mu}} + \alpha \mu}{-2\mu^2 + cK^3 \Lambda e^{-\frac{K^2 \Lambda x^*}{\mu}}} \quad (23)$$

which leads to

$$\lim_{K \rightarrow \infty} \frac{dx^*}{dK} = -\frac{\alpha}{2\mu} < 0 \quad \lim_{K \rightarrow \infty} \lim_{x^* \rightarrow 0} \frac{dU}{dK} = -\infty \quad (24)$$

In all the numerical experiments we conducted we found at most three roots for $\frac{dU}{dK} = 0$. Based on the second observation above, if $\frac{dU}{dK} = 0$ has at most three roots in $[0, K'']$, $K'' \leq \mu/\alpha$, and x^* takes only real values in that interval then $U(x^*, K)$ has at most one local maximum in $(0, K'')$.

3) *Number of Torrents*: Motivated by the cost function (13) we consider the case where $e^{-K^2 x \Lambda / \mu}$ is negligible (equation (14)). Therefore

$$\frac{\partial U}{\partial x} = \Lambda - \alpha \frac{K}{\mu} \Lambda - 2x\Lambda; \quad x^*(K) = \frac{1 - \alpha \frac{K}{\mu}}{2}; \quad \frac{dx^*}{dK} = \frac{-\alpha}{2\mu}$$

where the last expression agrees with (24). Also,

$$\frac{dU}{dK} = \frac{N(-\alpha x^* \Lambda + cK^{-2} \mu)}{\mu} \quad (25)$$

$$\frac{d^2 U}{dK^2} = \frac{N}{2} \left(\frac{\alpha^2}{\mu^2} \Lambda - \frac{4c}{K^3} \right); \quad \frac{d^3 U}{dK^3} \geq 0 \quad (26)$$

We conclude that, restricting to $K > 0$, $\frac{dU}{dK}$ is convex (note that $\frac{dU}{dK}$ is not defined for $K = 0$), $\lim_{K \rightarrow 0+} \frac{dU}{dK} = +\infty$ and

$\lim_{K \rightarrow +\infty} \frac{dU}{dK} = +\infty$ (compare to (24)). The convexity of $\frac{dU}{dK}$ implies that U has at most two critical points and at most one local maximum in $(1, \mu/\alpha)$. To find the optimal value of K , K^* , one may search for (at most two) roots of $\frac{dU}{dK}$ in the interval $(1, \mu/\alpha)$ and include them in set I . The optimal value of K is given by $K^* = \arg \max_{K \in I \cup \{1, \mu/\alpha\}} U(x^*(K), K)$.

IV. COMPETITION

We now consider the case where multiple publishers compete in the market. The publishers are coupled through the demand. The actions of one publisher affect the demand perceived by the others.

We assume publishers are price-takers. The market price for each file is $p > 0$. The only strategic decision is the bundling level. In addition, we consider the simplified scenario where each publisher has only two options: either to bundle all its files or not to bundle at all.

Users' utilities are given by (2). Each user requests a single content and its valuation for the requested file, V , is uniformly distributed in the range $[\alpha/\mu + p, M + \alpha N/\mu + p]$ where M is a constant, and N is the number of files offered by each publisher ($N_s = N$ for all s). Note that each user is willing to download the requested file and pay its market price. A fraction x of users tolerates downloading bundles,

$$x = 1 - \frac{\frac{\alpha N}{\mu} - \frac{\alpha}{\mu}}{M + \frac{\alpha N}{\mu} - \frac{\alpha}{\mu}} = \frac{M\mu}{M\mu + \alpha(N-1)} \quad (27)$$

If $N > 1$, for any $M > 0$ there is one single value of α which leads to demand x , $0 \leq x \leq 1$. The demand x decreases as the reluctance of users to bundles α increases.

Let $\mathcal{O}^{(i)}$ be the set of publishers that offer content i . The sets of publishers that offer bundled and unbundled versions of content i are $\mathcal{B}^{(i)} = \{t \in \mathcal{O}^{(i)} : K_t = N\}$ and $\mathcal{N}^{(i)} = \mathcal{O}^{(i)} \setminus \mathcal{B}^{(i)}$, respectively. We now define $q(\cdot)$ (equation (3)). When multiple publishers offer content i , demand is equally split across those publishers that do not bundle. This happens because if $s, t \in \mathcal{N}^{(i)}$ and $u \in \mathcal{B}^{(i)}$ then $K_s = K_t = 1$, $K_u = N$ and $U_c(s) = U_c(t) > U_c(u)$ (equation (2)). If all the publishers offering content i bundle, demand splits equally among them since $U_c(s) = U_c(t)$ for all $s, t \in \mathcal{B}^{(i)}$. Therefore,

$$q_s^{(i)}(\underline{K}) = \begin{cases} \frac{\Lambda^{(i)}}{|\mathcal{N}^{(i)}|} & s \in \mathcal{N}^{(i)} \\ \frac{x\Lambda^{(i)}}{|\mathcal{B}^{(i)}|} & s \in \mathcal{B}^{(i)} \text{ and } \mathcal{N}^{(i)} = \emptyset \\ 0 & \text{otherwise} \end{cases} \quad (28)$$

where $q(\cdot)$ does not depend on p because all publishers are price takers and all users are willing to pay the market price.

The utility of publisher s is given by (5). We now focus on publisher s costs. If the average demand per swarm is smaller than a threshold τ , the publisher incurs an aggregate

	not bundle	bundle
not bundle	A, W	B, X
bundle	C, Y	D, Z

TABLE III
GENERAL REWARDS FOR THE BUNDLING GAME

cost c for all swarms. Otherwise, the swarms provided by publisher s are self sustaining, and costs are zero (analogous to standard information goods in the economics literature).

$$G_s(\underline{\lambda}_s, K_s) = \begin{cases} c & \sum_{l=1}^N \lambda_s^{i_s(l)} K_s / N < \tau \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

This cost function resembles the active server cost function presented in Section III-B.2. However, the continuous function introduced in (12) for each swarm is now replaced by the above step function that captures costs aggregated across several swarms.

A. Duopoly

From here on we consider the special case where there are only two publishers in the market. Let $\mathcal{C}_1 = \{i : 1 \in \mathcal{O}^{(i)} \text{ and } 2 \notin \mathcal{O}^{(i)}\}$, $\mathcal{C}_2 = \{i : 2 \in \mathcal{O}^{(i)} \text{ and } 1 \notin \mathcal{O}^{(i)}\}$ and $\mathcal{C}_{1,2} = \{i : 1 \text{ and } 2 \in \mathcal{O}^{(i)}\}$ be the sets of contents offered solely by publisher 1, solely by publisher 2, and by both, respectively.

Each user seeks content that is provided either by publisher 1, publisher 2 or both. Requests arrive with rate $\Lambda_1 = \sum_{j \in \mathcal{C}_1} \Lambda^{(j)}$, $\Lambda_2 = \sum_{j \in \mathcal{C}_2} \Lambda^{(j)}$ and $\Lambda_B = \sum_{j \in \mathcal{C}_{1,2}} \Lambda^{(j)}$, respectively.

The normal form game is shown in Table III. We list publisher 1's strategies as rows and publisher 2's as columns. For each of the four cells, we give a pair of payoffs to the two publishers (in units of currency per second), first to publisher 1 and then to publisher 2, obtained using equations (5), (28) and (29). Given that $K_s = N$ if all the contents are bundled and $K_s = 1$ otherwise,

- $A, W = [p(\Lambda_i + \Lambda_B/2) - c \mathbf{1}_{(\Lambda_i + \Lambda_B/2)/N < \tau}]^+$
- $B, Y = [p(\Lambda_i + \Lambda_B) - c \mathbf{1}_{(\Lambda_i + \Lambda_B)/N < \tau}]^+$
- $C, X = [p\Lambda_i x - c \mathbf{1}_{\Lambda_i x < \tau}]^+$
- $D, Z = [p(\Lambda_i + \Lambda_B/2)x - c \mathbf{1}_{(\Lambda_i + \Lambda_B/2)x < \tau}]^+$

where $i = 1$ in cells A, B, C and D and $i = 2$ otherwise. $\mathbf{1}_C$ is an indicator variable equal to 1 if C is true and 0 otherwise. Note that all rewards are positive. Henceforth, this should be assumed even when the $^+$ symbol is omitted.

B. Large number of almost completely overlapping contents

We now consider the case where $\Lambda_i \approx 0$ ($i = 1, 2$) and $N \gg 0$. We assume that the following conditions hold:

	not bundle	bundle
not bundle	$p(\Lambda_1 + \frac{\Lambda_B}{2}) - c$, $p(\Lambda_2 + \frac{\Lambda_B}{2}) - c$	$p(\Lambda_1 + \Lambda_B) - c$, $p(\Lambda_2 x) - c$
bundle	$p(\Lambda_1 x) - c$, $p(\Lambda_2 + \Lambda_B) - c$	$p(\Lambda_1 + \frac{\Lambda_B}{2})x$, $p(\Lambda_2 + \frac{\Lambda_B}{2})x$

(a)

	not bundle	bundle
not bundle	4, 4	14, 0
bundle	0, 14	5, 5

(b)

TABLE IV

ALMOST COMPLETELY OVERLAPPING CONTENTS AND THE PUBLISHER'S DILEMMA.

- (C1) $(\Lambda_i + \Lambda_B)/N < \tau$: not bundling leads to publishing costs;
- (C2) $\Lambda_i x < \tau$: bundling but relying on a small population of users still leads to publishing costs;
- (C3) $(\Lambda_i + \frac{\Lambda_B}{2})x > \tau$: bundling and relying on a population of users with size greater than a critical threshold leads to negligible publishing costs.

Even though we write conditions (C2)-(C3) as a function of the variable x , they can be as well rewritten as a function of the parameter α (see comments following equation (27)). Given (C1)-(C3) the normal form game is depicted in Table IV(a).

For a numerical example, consider the following parameters: $\Lambda_B = 20$, $\Lambda_i = 0$, $x = 0.5$, $c = 6$, $p = 1$, $N = 10$, $\tau = 2.1$. The normal form game in this case is given in Table IV(b). Publishers are better off if both bundle. However, the Nash equilibrium of this game is (not bundle, not bundle), which is not Pareto optimal. A Nash equilibrium which is not Pareto optimal is referred to as a social trap.

We refer to a game in which bundling consists of a dominant strategy for both publishers and where the resulting equilibrium (not bundle, not bundle) is not Pareto optimal as a *publisher's dilemma*. A publisher's dilemma happens when $B > D > A > C$ and $Y > Z > W > X$ (Table III) and is an example of a prisoner's dilemma [21].

In what follows, we show that if in addition to (C1)-(C3) the arrival rate of requestors $\Lambda_B \gg 0$ and α , the reluctance to bundling is close to 0 (hence x is close to 1), then complete overlap ($\Lambda_i = 0, i = 1, 2$) favors a publisher's dilemma. For that purpose, we add a condition on the arrival rate of clients,

- (C4) $\frac{2c}{p} < \Lambda_B < \frac{2c}{p(1-x)}$

In particular, (C4) is valid for the example in Table IV.

Proposition 4.1: If conditions (C1)-(C4) hold then there exists an Λ_i^* such that if $0 \leq \Lambda_i \leq \Lambda_i^*$ ($i = 1, 2$) the game in Table IV is a publisher's dilemma.

Proof: Under the assumptions of the proposition, the condition for a publisher's dilemma $C < A < D < B$ maps into

$[p\Lambda_1 x - c]^+ < [p(\Lambda_1 + \frac{\Lambda_B}{2}) - c]^+ < p(\Lambda_1 + \frac{\Lambda_B}{2})x < [p(\Lambda_1 + \Lambda_B) - c]^+$. If $\Lambda_1^* = \max(0, \Lambda_1^{**})$, $\Lambda_1^{**} = \min(\frac{c}{px}, \frac{c - p(1-x)\frac{\Lambda_B}{2}}{p(1-x)})$ and $0 \leq \Lambda_1 \leq \Lambda_1^*$: (1) the first term $[p\Lambda_1 x - c]^+ = 0$ because $\Lambda_1 \leq \frac{c}{px}$; (2) the second term $[p(\Lambda_1 + \frac{\Lambda_B}{2}) - c]^+ > 0$ since (C4) implies $p\frac{\Lambda_B}{2} > c$; (3) the relation $[p(\Lambda_1 + \frac{\Lambda_B}{2}) - c]^+ < p(\Lambda_1 + \frac{\Lambda_B}{2})x$ is satisfied since $\Lambda_1 < \Lambda_1^{**}$ and (C4) implies $0 < p\frac{\Lambda_B}{2} - c < p(\Lambda_B/2)x$ so that even if $\Lambda_1 = 0$ still the relation holds; (4) (C4) also leads to $0 < p(\Lambda_B/2)x < p\Lambda_B - c$ hence $p(\Lambda_1 + \frac{\Lambda_B}{2})x < [p(\Lambda_1 + \Lambda_B) - c]^+$.

The condition $X < W < Z < Y$ follows similarly. ■

As a matter of fact, note that $\Lambda_i = 0$ favors (C1) and (C2) and $\Lambda_B \gg 0$ favors (C3). If $N \gg 0$, (C1) holds even when $\Lambda_B \gg 0$.

Key insight: If a publisher distinguishes itself by providing exclusive content ($\Lambda_i > 0$) it is less likely to end up in a social trap consisting of a non Pareto equilibrium where no publisher bundles.

C. Partially overlapping contents

We now consider the case where content is only partially overlapping and users are reluctant to download bundles ($x \approx 0$).² We replace (C1) by the following two conditions and add condition (C5):

- (C1') $(\Lambda_i + \frac{\Lambda_B}{2})/N < \tau$: not bundling leads to publishing costs if the population of requestors is small;
- (-C1) $(\Lambda_i + \Lambda_B)/N > \tau$: if the population of requestors is large, costs are zero;
- (C5) $c > p(\Lambda_1 + \Lambda_2 + \frac{\Lambda_B}{2})$: publishing costs, when incurred, are high.

Alternatively, (C1') and (-C1) can be interpreted as conditions over the number of files offered by each publisher: the number of files (N) offered by each publisher falls in the range $(\Lambda_i + \frac{\Lambda_B}{2})/\tau < N < (\Lambda_i + \Lambda_B)/\tau$.

Conditions (C1'), (-C1) and (C5) lead to the game shown in Table V(a). The reluctance of users to download bundles forces the publishers that bundle to go off the market. For $\Lambda_B = \Lambda_1 = \Lambda_2 = 10$, $x = 0$, $c = 16$, $p = 1$, $N = 6$, $\tau = 5$, we obtain the game shown in Table V(b). This game has three pure equilibria: (not bundle, not bundle), where none of the enterprises enters the market, and two others, (bundle, not bundle), (not bundle, bundle), where only one enterprise enters the market and does not bundle.

We refer to a game in which (1) there are at least two pure strategy equilibria, e_1 and e_2 ; (2) each publisher prefers a different pure strategy equilibrium and (3) both e_1 and e_2 are Pareto optimal as a *bundle off game*.

²In our technical report [13] we also analyze the case where x is not necessarily small and publishers may collude.

	not bundle	bundle
not bundle	0, 0	$p(\Lambda_1 + \Lambda_B), 0$
bundle	0, $p(\Lambda_2 + \Lambda_B)$	0, 0

(a)

	not bundle	bundle
not bundle	0, 0	4, 0
bundle	0, 4	0, 0

(b)

TABLE V

PARTIALLY OVERLAPPING CONTENTS AND THE BUNDLE OFF GAME.

Note that the bundle off game is an example of a degenerated hawk dove game [21]. The conditions for a bundle off game are (Table III) $B > D \geq C \geq A$ and $Y > Z \geq X \geq W$. Next, we show that reluctance of users to download bundles favors a bundle off game:

Proposition 4.2: If conditions [(C1'), (-C1) and (C5)] hold then there is a threshold Δ such that if $x < \Delta$ the game in Table III is a bundle off game.

Proof: If conditions (C1'), (-C1) and (C5) are met and $x < \tau/(\Lambda_i + \frac{\Lambda_B}{2})$ then the game in Table III reduces to the one shown in Table V(a), which is a bundle off game. ■

Key insight: If users are reluctant to download bundles and the costs to provide unbundled files are high, we may expect that in the equilibrium only one publisher offers an unbundled version of the contents it has.

V. RELATED WORK

In their seminal work on bundling, Adams and Yellen [1] showed that a firm that provides two products can take advantage of bundling by increasing the number of costumers that buy its products. Since then, there is a growing literature on bundling (for a survey, see [11]). More recently, certain authors have also considered strategic bundling [25], [15]. However, none of these works considered bundling of files in a peer-to-peer network.

Salinger [19] was the first to point out that bundling is a demand smoother in the case of homogeneous demands. This is a consequence of the law of large numbers: if all buyers draw their value for the goods from the same probability distribution, average valuation converges to the mean as the number of goods increases. Therefore, if the firm sets the price of the bundle to the mean times the number of elements in the bundle, virtually every customer will be willing to buy and profit is maximized. Later, Bakos and Brynjolfsson [3] reached the same conclusion in the context of information goods, i.e., goods that have almost zero cost to be replicated. This indicates that firms that cannot price discriminate may still use bundling to increase its sales. The same reasoning applies if demand is negatively correlated [17, Section 12.6].

Closely related to our work are those by Bakos and Brynjolfsson [4], Fay and Mackie-Mason [10] and Croson

and Sainders [8]. They were the first to consider bundling of information goods under competition. What distinguishes our work from these is the fact that in a peer-to-peer system the replication costs depend on the number of users that possess the content. In addition, we take into account performance factors related to network metrics such as download time and availability which were not considered in previous work.

There has been a vast amount of work on P2P systems focusing on their performance analysis, e.g., [27], [18], content search, e.g., [9], [23], neighbor selection, e.g., [22], [5], replication strategies, e.g., [7] and legislation, e.g., [26]. However none of these deal with the problems of when and how to bundle content.

VI. CONCLUSION

We considered publishers that rely on swarming systems to disseminate files and proposed a model to capture the tradeoffs in the choices of prices and bundling strategies. Using the proposed model, we showed conditions for existence and uniqueness of the equilibrium and how this equilibrium changes as a function of the bundling strategy. In the multi-firm case, we explained the effects of system parameters on the Nash equilibrium. Future work consists on analyzing the Bertrand-Nash equilibrium of the duopoly and using multi agent reinforcement learning to analyze the multi-firm case.

ACKNOWLEDGEMENTS

We thank fruitful discussions on the performance and availability implications of bundling with Arun Venkataramani, Bin Li and Antonio A. A. Rocha.

This work was supported in part by the National Science Foundation under award numbers CNS-0519922 and CNS-0721779. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation. Research of DSM also funded in part by a scholarship from CAPES/Fulbright (Brazil).

APPENDIX

In this section we verify that the hypothesis of Propositions 3.3 and 3.4 are satisfied respectively for functions (13) and (12). In order to check properties like convexity/concavity or monotonicity we can ignore multiplicative factors or scaling factors as long as they are positive. For (13) we consider $f_1(x) = 1 - e^{-x}$. The corresponding marginal cost is convex, because the the third derivative is $f_1'''(x) = e^{-x} > 0$.

For (12) we consider $f_2(x) = xe^{-x}$. The derivatives of interest are: $f_2''(x) = e^{-x}(x-2)$ and $f_2'''(x) = e^{-x}(3-x)$. f_2 is concave in $[0,2]$ and convex in $(2,\infty)$. Its derivative is convex in $(0,3)$.

REFERENCES

- [1] ADAMS, W., AND YELLEN, J. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics* 90 (1976), 475–498.
- [2] ANDERSON, C. *The Long Tail: Why the Future of Business is Selling Less of More*. Hyperion, 2006.
- [3] BAKOS, AND BRYNJOLFSSON. Bundling information goods: Pricing, profits and efficiency. *The Economics of Digital Info. Goods* (1998).
- [4] BAKOS, AND BRYNJOLFSSON. Bundling and competition on the internet. *Marketing Science* 19, 1 (2000), 63–82.
- [5] BERNSTEIN, D., FENG, Z., LEVINE, B., AND ZILBERSTEIN, S. Adaptive peer selection. In *IPTP* (2003).
- [6] COHEN, B. Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer Systems* (2003).
- [7] COHEN, E., AND SHENKER, S. Replication strategies in unstructured peer to peer. In *SIGCOMM* (2002).
- [8] CROSON, D., AND SAUNDERS, A. Competition and cooperation in the bundled software market. In *WISE* (2004).
- [9] FAST, A., LEVINE, B. N., AND JENSEN, D. Creating social networks to improve peer-to-peer networking. In *SIGKDD* (2005).
- [10] FAY, AND MACKIE-MASON. Competition between firms that bundle information goods. In *Proceedings of the 27th Annual Telecommunications Policy Research Conference* (1999).
- [11] FUERDERER, R., HERRMANN, A., AND WUEBKER, G. *Optimal Bundling*. Springer, 1999.
- [12] GALIDAKIS, I. N. On an application of lambert’s w function to infinite exponentials. *Complex Variables and Elliptic Equations* 49 (2004), 759–780.
- [13] MENASCHE, D. S., NEGLIA, G., TOWSLEY, D., AND ZILBERSTEIN, S. Strategic reasoning about bundling in swarming systems. *University of Massachusetts Technical Report UM-CS-2009-011* (2009).
- [14] MENASCHE, D. S., ROCHA, A. A., LI, B., TOWSLEY, D., AND VENKATARAMANI, A. Bundling builds availability in Bittorrent. *University of Massachusetts Technical Report UM-CS-2009-010* (2009).
- [15] NALEBUFF, B. Bundling. In *School of Management, Yale ICF Working Paper No. 99-14* (1999).
- [16] PAGE, W. More long tail debate: mobile music no, search yes, 2008. <http://longtail.typepad.com/>.
- [17] PERLOFF, J. *Microeconomic: Theory and Applications with Calculus*. Pearson International, 2008.
- [18] QIU, D., AND SRIKANT, R. Modeling and performance analysis of bittorrent-like p2p networks. In *SIGCOMM* (2004), pp. 367–378.
- [19] SALINGER, M. A. A graphical analysis of bundling. *Journal of Business* 68 (1995), 85–98.
- [20] SHERWOOD, R., BRAUD, R., AND BHATTACHARJEE, B. Slurpie: A cooperative bulk data transfer protocol. In *Infocom* (2004).
- [21] SHOR, M. Dictionary of game theory terms, game theory .net, 2009. <http://www.gametheory.net/dictionary/>.
- [22] TOKA, L., AND MICHIARDI, P. Analysis of user-driven peer selection in peer-to-peer backup and storage systems. In *GameNets* (2008).
- [23] VOULGARIS, S., KERMARREC, A.-M., AND MASSOULIE, L. Exploiting semantic proximity in peer-to-peer content searching. In *FTDCS* (2004), pp. 238–243.
- [24] WIKIPEDIA. Entry on digital rights management, 2009.
- [25] WILSON, R. Strategic models of entry deterrence. In *Handbooks in Economics* (1994).
- [26] WONG, S., ALTMAN, E., AND IBRAHIM, M. P2P networks: The interplay between legislation and information technology. In *INRIA Technical Report RR-6889* (2009).
- [27] YANG, X., AND DE VECIANA, G. Service capacity of peer to peer networks. In *INFOCOM* (2004).