

Coupon Replication Systems

Laurent Massoulié and Milan Vojnović

Abstract—Motivated by the study of peer-to-peer file swarming systems à la BitTorrent, we introduce a probabilistic model of coupon replication systems. These systems consist of users, aiming to complete a collection of distinct coupons. Users enter the system with an initial coupon provided by a bootstrap server, acquire other coupons from other users, and leave once they complete their coupon collection.

For open systems, with exogenous user arrivals, we derive stability condition for a layered scenario, where encounters are between users holding the same number of coupons. We also consider a system where encounters are between users chosen uniformly at random from the whole population. We show that sojourn time in both systems is asymptotically optimal as the number of coupon types becomes large.

We also consider closed systems with no exogenous user arrivals. In a special scenario where users have only one missing coupon, we evaluate the size of the population ultimately remaining in the system, as the initial number of users, N , goes to infinity. We show that this size decreases geometrically with the number of coupons, K . In particular, when the ratio $K/\log(N)$ is above a critical threshold, we prove that this number of leftovers is of order $\log(\log(N))$.

These results suggest that, under the assumption that the bootstrap server is not a bottleneck, the performance does not depend critically on either altruistic user behavior, or on load balancing strategies such as rarest first.

I. INTRODUCTION

ACCORDING to the file swarming paradigm, dissemination of a large file to a large population of interested users proceeds as follows. The file is chopped into smaller sub-files, or chunks; interested users not only download such chunks in order to retrieve the desired file, but also upload them when requested, thereby being both clients and servers of the whole system. The latter feature, which is characteristic of peer-to-peer systems, has the consequence that the service capacity of the system scales with the demand it is facing. File swarming is implemented in peer-to-peer file sharing networks such as eDonkey or KaZaA (see e.g. [1], [2]), and is the core functionality of BitTorrent [5], a popular file download system. Empirical studies of BitTorrent [11], [15] indicate that the download of a given file can be broken up into three distinct phases, which may last several days. The so-called *flash crowd* phase typically takes place shortly after the file has been made available. It is characterized by a sudden burst of user arrivals. The second phase is a *steady state* mode, where the rate of arrival of new users matches the rate of departures, and the number of active users fluctuates around some equilibrium value. Finally, the life of a file download is concluded by an *end* phase, where the arrival rate of new users gradually decreases.

Author Affiliations: L. Massoulié is with Thomson Research, Paris (e-mail: laurent.massoulie@thomson.net). M. Vojnović is with Microsoft Research, Cambridge (e-mail: milanv@microsoft.com). Work performed while L. Massoulié was with Microsoft Research.

The work of Yang and de Veciana [17] suggests that BitTorrent is efficient, in that its service capacity adapts to increases in demand, thereby providing fast downloads independently of the rate of demand. However, it is not completely understood which aspects of the system are critical to maintain this property. For instance, most BitTorrent users remain in the system for some time after having retrieved the desired file, then acting purely as servers; would performance deteriorate significantly in the absence of such cooperation? Another feature of BitTorrent is its implementation of the *rarest first policy*: that is, users download rare chunks preferentially. The rationale for this policy is that it balances the number of replicas of each chunk available within the system, and thus prevents starvation of some chunk types. It is not known however whether the system would perform well if one implemented simpler strategies instead.

Our aim in the present paper is to develop detailed models of file swarming systems, allowing to investigate these issues by analytical means. To this end, we introduce *coupon replication systems*. These stochastic models feature populations of users, each of which aims to complete a collection of coupons. With file swarming systems in mind, coupons should be interpreted as file chunks. We distinguish between *open* and *closed* systems, according to whether there are exogenous user arrivals or not. Open systems are appropriate to analyze the steady state mode in file downloads. Closed systems are appropriate to study the transient behavior of file swarming during flash crowds.

In both closed and open systems, each user contacts a randomly selected target user periodically. It then replicates a coupon from the collection of the target user to increase its own collection, if the collection of the target user is not fully contained in its own collection. Many distinct policies for random target selection, and choice of which coupon to replicate can be envisaged within this framework.

We focus mainly on two distinct systems, namely *layered* and *flat* systems. In layered systems, users choose their target user uniformly at random from the subpopulation of users having a collection of the same size—in number of chunks—as they have. Thus, partitioning users into *layers*, each layer corresponding to a collection size, in these systems user interactions are only within users of the same layer.

In flat systems, users select their targets uniformly at random from the total population present in the system. In both flat and layered systems, users choose which coupon to replicate uniformly at random from the set of coupons that their selected target has and that they don't have. This might be thought of as a blind selection policy, in contrast to the rarest first policy. We also assume that users leave the system as soon as they have completed their collection, i.e. they do not display the kind of altruism observed in actual BitTorrent users.

These systems are motivated as follows. In typical file swarming systems a user is provided upon joining with a fixed list of other users to sample its targets from. If this list contains users that recently joined, assuming users enlarge their collections at a similar rate, then interactions are largely among users holding the same number of coupons. This makes the layered system plausible. If the lists are instead drawn at random from the total user population, that is no preference is given to recently joined users, this corresponds more to the flat system.

We analyze such systems in a large population asymptotic regime. This is achieved for open systems by letting the exogenous user arrival rates increase to infinity; for closed systems, we simply let the initial population size increase to infinity. Taking such limits makes the analysis more tractable; in addition, there are practical reasons to focus on this limit regime, as empirical studies of BitTorrent [11], [15], show that popular downloads (e.g., Linux binaries) can involve hundreds of thousands of simultaneous users. The same studies also indicate that huge files can be chopped in several thousands of chunks, which provides motivation for looking at asymptotics in the number of chunks as well.

Our main findings for open systems are the following. In the layered case, we provide necessary and sufficient stability conditions, together with an explicit characterization of the system equilibrium points. These conditions are phrased in terms of the exogenous user arrival rates. For flat systems, we treat only systems where exogenous arrival rates satisfy a natural symmetry condition. Under this condition, we establish boundedness of the dynamics, and derive analytical upper bounds valid for any equilibrium point. These bounds show that the performance in flat systems, as measured by response time, approaches that of layered systems when the total number of coupons is large. For both systems, the performance is asymptotically optimal in this many coupons limit. These results imply that performance of file swarming systems such as BitTorrent is not critically dependent on either altruism of users, or on the implementation of refined balancing policies such as rarest first, in cases when bootstrap servers that grant initial coupon are sufficiently provisioned. Our treatment of open systems is similar in spirit to the work of Yang and de Veciana [17] and Qiu and Srikant [7], with the difference that their models do not capture precise description of coupon collections held by users.

For closed systems, we focus on a special case, where all users have only one missing coupon. They thus all belong to the last layer, which implies that both the layered and flat system display exactly the same behavior. Under this specific condition, we analyze the size of the population ultimately remaining in the system, in the limit where the initial population size, N , goes to infinity. We show that the size of this ultimate population decreases geometrically with the total number of coupons, K . In particular, when the ratio $K/\log(N)$ is above a critical threshold, we prove that this number of left-overs is of order $O(\log(\log(N)))$. This number of left-overs may be interpreted as the number of users who will eventually have to query a central server in order to complete their collection, at the end of a flash crowd phase initiated by a burst of N

clients. This result thus provides additional support to the claim, made in [17], that BitTorrent copes well with flash crowds. Related work on the performance of schemes for disseminating K items to a fixed population of N users can be found in Mundinger and Weber [16], where they propose an optimal centralized scheme, achieving this dissemination in $O(K + \log(N))$ time steps, and in Deb and Médard [6], where they propose a decentralized scheme using random network coding techniques.

The paper is organized as follows. The general model considered in this paper is defined in Section 2. Section 3 contains the analysis of layered open systems. Flat open systems are considered in Section 4, and closed systems in Section 5. Section 6 discusses special cases of our model and related work. Concluding remarks are given in Section 7.

II. MODEL DESCRIPTION

The model in this paper aims to capture efficiency of coupon exchange mechanisms in terms of global file download time. The model granularity is that of coupon exchanges. We abstract away the lower level details of transport layer that, for example, include packet transfer and network delay. It is assumed that each user peers with other users by random sampling, and the time interval between such peerings is interpreted as the time for a user to download a coupon. The rate of peering attempts is thus equal to the upload capacity of a contacted user. Each peering attempt either results in a successful coupon exchange or no exchange at all. The latter case occurs when the contacted user has no useful coupon to offer to the instigator user. Thus, we model the dynamics of collection of coupons held by users, which depend on the rate of successful peerings. Details are as follows.

Consider a collection of K distinct coupons, indexed by $k \in \{1, \dots, K\}$. Denote by \mathcal{C} the family of strict subsets c of $\{1, \dots, K\}$. Let $X_c(t)$ denote the number of users holding the sub-collection c , for any $c \in \mathcal{C}$, present in the system at time t . A user holding sub-collection c is also called a type c -user. The system dynamics are then specified as follows.

For each c , exogenous arrivals of type c -users take place at the instants of a Poisson process with rate λ_c . Of particular interest is the case where $\lambda_c = 0$ for all c that are not singletons. The non-zero arrival rates then correspond to sets $\{k\}$, for some $k \in \{1, \dots, K\}$. We interchangeably denote the corresponding arrival rate $\lambda_{\{k\}}$ or λ_k . In the context of file swarming systems, this reflects the situation where users initially obtain one coupon from a central server, before entering the system, this being coupon k with probability $\lambda_k / \sum_{i=1}^K \lambda_i$.

Given the system state $X(t) := \{X_c(t)\}_{c \in \mathcal{C}}$, and $c \in \mathcal{C}$, each type c -user initiates encounters at the instants of a Poisson process of rate 1^1 , and the encountered user is of type $s \in \mathcal{C}$ with probability $p_{cs}(X(t))$, where for each c and X , the family $(p_{cs}(X))_{s \in \mathcal{C}}$ is a probability distribution on \mathcal{C} . As an example, for the layered systems we shall consider, we assume that

¹We make this assumption for tractability. In a more general model, one might instead consider general, independent and identically distributed inter-encounter times.

each user, when initiating an encounter, selects the target user uniformly from the set of users holding the same number of coupons. The corresponding probability distribution then reads

$$p_{cs}^{layered}(X) = \begin{cases} 0 & \text{if } |c| \neq |s|, \\ \frac{X_s}{\sum_{s':|s'|=|c|} X_{s'}} & \text{if } |c| = |s|. \end{cases} \quad (1)$$

In contrast, in flat systems, users are assumed to sample uniformly from the total population of users when selecting a target user. The corresponding probability distribution then reads

$$p_{cs}^{flat}(X) = \frac{X_s}{\sum_{s' \in \mathcal{C}} X_{s'}}, \quad s \in \mathcal{C}. \quad (2)$$

Given that a type c -user initiates an encounter with a type s -user, the following then takes place. Provided some coupons of the target user are of interest to the user initiating the encounter, that is if $s \setminus c$ is not empty, then the target user selects a coupon uniformly at random from $s \setminus c$, say coupon i , and replicates it at the initiator, whose type then becomes $c \cup \{i\}$. In case $c \cup \{i\} = \{1, \dots, K\}$, this user has completed its collection, and leaves the system. If $s \setminus c = \emptyset$, then nothing happens.

Under these assumptions, it is easily verified that $\{X(t)\}_{t \geq 0}$ is a Markov process taking its values in $\mathbb{Z}_+^{\mathcal{C}}$. Denoting by e_c the unit vector of $\mathbb{Z}_+^{\mathcal{C}}$ whose c -coordinate equals 1, and with all other coordinates equal to zero, the non-zero transition rates of the Markov process are, for all $c \in \mathcal{C}$, all $i \notin c$,

$$\begin{aligned} X &\rightarrow X + e_c && \text{with rate } \lambda_c, \\ X &\rightarrow X - e_c + e_{c \cup i} && \text{with rate } X_c \sum_{s:i \in s} p_{cs}(X) \frac{1}{|s \setminus c|}. \end{aligned}$$

Note that λ_c is the rate of exogenous user arrivals holding coupon collection c , while the second line of the above display is the rate at which a user with collection c acquires coupon i .

In the sequel we will be interested in analyzing these systems under a *large population* asymptotic regime. We define this regime as follows. For any integer $N > 0$, we consider an associated Markov process, with the same parameters as above, except that the arrival rates scale linearly with N , i.e. in the N -th system we take $\lambda_c^N = N\lambda_c$, and the initial condition $X^N(0)$ is also dependent on N . We assume that it is such that $N^{-1}X^N(0)$ converges to $x(0) \in \mathbb{R}_+^{|\mathcal{C}|}$.

We note that for both flat and layered systems, the probability $p_{cs}(X)$ is invariant by joint rescaling of the coordinates X_c by a common factor. This implies that in both cases, the Markov process X^N has non-zero transition rates of the form

$$q^N(X, X + v) = Nq_v(N^{-1}X),$$

for fixed functions $q_v(\cdot)$ of the jump vector v , and of the rescaled state $N^{-1}X$. This is precisely the definition of *density dependent jump Markov processes*, given in Kurtz [13], p. 51 Eq. 8.5. Under technical assumptions² which are trivially verified in our present setting, then by Theorem 8.1, p. 52 in [13], we obtain that the rescaled process $N^{-1}X^N(t)$ converges almost surely, uniformly on all finite intervals $[0, T]$,

²Namely, boundedness, and Lipschitz continuity of the associated functions $q_v(\cdot)$.

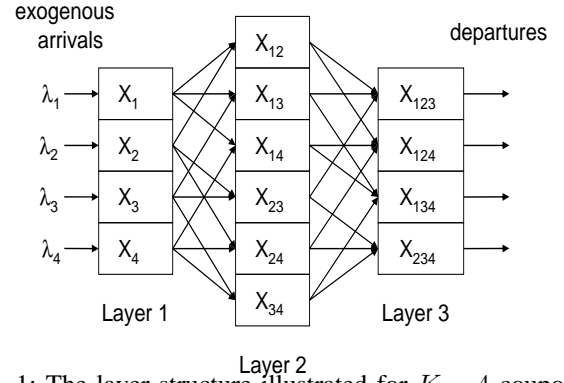


Fig. 1: The layer structure illustrated for $K = 4$ coupons.

i.e.

$$\lim_{N \rightarrow \infty} \sup_{0 \leq t \leq T} \left| \frac{1}{N} X^{(N)}(t) - x(t) \right| = 0, \quad (3)$$

where $x(t)$ is the solution of the system of differential equations

$$\begin{aligned} \frac{d}{dt} x_c(t) &= \lambda_c + \sum_{i \in c} x_{c \setminus i} \sum_{s: i \in s} p_{c \setminus i, s}(x) \frac{1}{|s \setminus c| + 1} \\ &\quad - x_c \sum_{s: |s \setminus c| \neq 0} p_{cs}(x), \end{aligned} \quad (4)$$

with initial condition $x(0)$. It will be convenient to rewrite this equation as

$$\frac{d}{dt} x_c(t) = \beta_c(x) - \mu_c(x), \quad (5)$$

where

$$\begin{aligned} \beta_c(x) &= \lambda_c + \sum_{i \in c} x_{c \setminus i} \sum_{s: i \in s} p_{c \setminus i, s}(x) \frac{1}{|s \setminus c| + 1}, \\ \mu_c(x) &= x_c \sum_{s: |s \setminus c| \neq 0} p_{cs}(x). \end{aligned} \quad (6)$$

We shall refer to $\beta_c(x)$ as the *arrival rate* of type c -users, and $\mu_c(x)$ as the *departure rate* of type c -users. This is precisely the system of differential equations (4) we study in the next two sections, for layered and flat systems respectively.

III. OPEN SYSTEMS: THE LAYERED CASE

We let \mathcal{C}_k denote the set of those $c \in \mathcal{C}$ with cardinality k , $k = 1, \dots, K - 1$. We shall say that $\{x_c\}_{c \in \mathcal{C}_k}$ constitutes the state of the k -th layer of the system (see Figure 1). The following property of layered systems, i.e. with $p_{sc} = p_{sc}^{layered}$, will allow us to analyze their stability by decomposing them in subsystems. For any $c \in \mathcal{C}_k$, $k = 1, \dots, K - 1$, the arrival rate $\beta_c(x)$ of type c -users depends only on the state of the $k - 1$ -th layer, and the departure rate $\mu_c(x)$ of type c -users depends only on the state of the k -th layer. Thus the dynamics of the first k layers do not rely on the state variables of the subsequent layers, and can be studied in isolation.

In order to study the limiting behavior of the whole system for large t we can thus proceed by induction, and study the asymptotic behavior of each of the subsystems consisting of the first k layers, for $k = 1, \dots, K - 1$. Provided the $k - 1$ -th subsystem reaches an equilibrium point as $t \rightarrow \infty$, then the limiting behavior of the k -th subsystem is characterized by the dynamics of the k -th layer, which is similar to those of the first layer.

In the remainder of this section we first study single layers

in isolation, then consider the global system, and finally look at the corresponding response times.

A. Single layer: stability analysis

Without loss of generality, we consider the first layer, with arrival rates $\beta_c = \lambda_c$, where λ_i is the arrival rate of users holding initial coupon i with convention $0 \leq \lambda_K \leq \lambda_{K-1} \leq \dots \leq \lambda_1 < \infty$. The differential system of interest can be rewritten as:

$$\frac{d}{dt}x_i(t) = \lambda_i - x_i(t) \left(1 - \frac{x_i(t)}{s(t)}\right), \quad i = 1, 2, \dots, K, \quad (7)$$

where $s(t) := \sum_{k=1}^K x_k(t)$ and by convention $0/0 := 1$. We say a point $x^0 \in \mathbb{R}_+^K$ is a rest point if $dx(t)/dt|_{x(t)=x^0} = 0$. It is easily seen that for $K = 2$, there exists a rest point if and only if $\lambda_1 = \lambda_2$, in which case there is an infinity of rest points, given by $(x_1, x_2) \in \mathbb{R}_+^2$ such that $1/x_1 + 1/x_2 = 1/\lambda_1$. Also, one can establish that, for $K > 2$, if $\lambda_3 = \dots = \lambda_K = 0$, then there exist rest points if and only if $\lambda_1 = \lambda_2$, in which case they are again given by $(x_1, x_2, 0, \dots, 0)$ such that $1/x_1 + 1/x_2 = 1/\lambda_1$. Thus in the sequel we focus on the case where $K > 2$ and $\lambda_3 > 0$. Existence and uniqueness of rest points is settled by the following result, the proof of which is given in Appendix A.

Theorem 1. *For any $K > 2$, and $\lambda_3 > 0$, there exists a rest point for the system of ODEs (7) if and only if (D) $\lambda_1 < \sum_{k=2}^K \lambda_k$. When (D) holds, the rest point is unique, and is characterized as follows. If (C) $\sum_{k=2}^K \sqrt{1 - \frac{\lambda_k}{\lambda_1}} < K - 2$, then x is the rest point given by*

$$x_i = \frac{s}{2} \left(1 - \sqrt{1 - \frac{4\lambda_i}{s}}\right), \quad i = 1, 2, \dots, K, \quad (8)$$

where $s \in [4\lambda_1, \infty)$ is the unique solution of

$$\sum_{i=1}^K \sqrt{1 - \frac{4\lambda_i}{s}} = K - 2. \quad (9)$$

If (C) fails, the rest point is given by

$$x_i = \begin{cases} \frac{s}{2} \left(1 + \sqrt{1 - \frac{4\lambda_1}{s}}\right) & i = 1 \\ \frac{s}{2} \left(1 - \sqrt{1 - \frac{4\lambda_i}{s}}\right) & i = 2, 3, \dots, K, \end{cases} \quad (10)$$

where $s \in [4\lambda_1, \infty)$ is the unique solution of

$$-\sqrt{1 - \frac{4\lambda_1}{s}} + \sum_{i=2}^K \sqrt{1 - \frac{4\lambda_i}{s}} = K - 2. \quad (11)$$

Remark 1. *It is convenient to rewrite the conditions (C) and (D) by using the substitution $u_i = 1 - \lambda_i/\lambda_1$, for $i = 1, 2, \dots, K$. They now read (D) $\sum_{i=2}^K u_i < K - 2$ and (C) $\sum_{i=2}^K \sqrt{u_i} < K - 2$. By definition $0 \leq u_i \leq 1$, for all $i = 1, 2, \dots, K$, and hence $\sum_{i=2}^K u_i \leq \sum_{i=2}^K \sqrt{u_i}$, with equality if and only if $u_i \in \{0, 1\}$ for all i . Thus (C) implies (D). Also both conditions hold whenever $\lambda_1 = \lambda_2$, i.e. no single sub-population dominates.*

The symmetric case $\lambda_1 = \lambda_2 = \dots = \lambda_K$ is extremal in the following sense.

Proposition 1. *Let $\lambda_1, \lambda_2, \dots, \lambda_K > 0$ be any such that $\sum_{i=1}^K \lambda_i = \lambda$, for a fixed $\lambda > 0$. The rest point x of (7) satisfies $\sum_{i=1}^K x_i \geq K\lambda/(K-1)$, where equality is achieved only if $\lambda_1 = \lambda_2 = \dots = \lambda_K = \lambda/K$.*

Proof: Case 1: x is given by (8)–(9). From (9),

$$K - 2 = \sum_{i=1}^K \sqrt{1 - \frac{4\lambda_i}{s}} \leq K \sqrt{1 - \frac{4\sum_{i=1}^K \lambda_i}{Ks}}.$$

The last inequality follows from Jensen's inequality, for which equality holds only if $1 - 4\lambda_i/s = 1 - 4\lambda_1/s$, for all $i = 2, 3, \dots, K$, or equivalently, $\lambda_1 = \lambda_2 = \dots = \lambda_K$. The above inequality is equivalent to $s \geq \lambda K/(K-1)$. As $s = \sum_{i=1}^K x_i$, the result of case 1 follows.

Case 2: x is given by (10)–(11). Equation (11) implies $\sum_{i=1}^K \sqrt{1 - \frac{4\lambda_i}{s}} \geq K - 2$. Proof follows that of Case 1.

We next establish that (D) is a necessary and sufficient condition for (7) to have a bounded solution. We say $x(t) = (x_1(t), x_2(t), \dots, x_K(t))$ is uniformly bounded if there exists a finite $D \geq 0$ such that $\limsup_{t \rightarrow +\infty} x_i(t) \leq D$, for all $i = 1, 2, \dots, K$.

Theorem 2 (Boundedness). *Any solution $x(t)$ of (7) is uniformly bounded if (D) holds, while when (D) fails, $\lim_{t \rightarrow \infty} x_1(t) = +\infty$.*

The proof is given in Appendix B. Condition (D) admits a natural interpretation in terms of stability of service systems. Indeed, in the system under consideration, individuals are both clients and servers, and they can serve clients of all types except their own. In this light, condition (D) can be read as: for any type i , the arrival rate of corresponding clients, λ_i , is strictly less than the total arrival rate of servers that can serve such clients, that is $\sum_{j \neq i} \lambda_j$. Thus, formally, condition (D) is related to the classical stability condition of single server queues, which reads $\rho < 1$, where the load ρ is defined as the ratio of work arrival rate per time unit to service rate per time unit (see Loynes [14]).

The main result of this section is the following

Theorem 3 (Global stability). *Under (D), the system of the ODEs (7) is globally asymptotically stable, i.e. there exists a unique limit point to which the system converges from any initial point.*

Proof: Consider $V : (0, +\infty)^K \rightarrow \mathbb{R}$ defined by

$$V(x) = \sum_{i=1}^K \log(x_i) - \log \left(\sum_{i=1}^K x_i \right) + \sum_{i=1}^K \frac{\lambda_i}{x_i}. \quad (12)$$

This is a Lyapunov function for the dynamics (7). Indeed, it is continuously differentiable on $(0, \infty)^K$, and such that

$$\begin{aligned} \frac{d}{dt}V(x(t)) &= \sum_{i=1}^K \frac{\partial}{\partial x_i} V(x(t)) \frac{d}{dt}x_i(t) \\ &= \sum_{i=1}^K [\lambda_i - x_i (1 - \frac{x_i}{s})] \left[\frac{1}{x_i} - \frac{1}{s} - \frac{\lambda_i}{x_i^2} \right] \\ &= - \sum_{i=1}^K \frac{1}{x_i^2} [\lambda_i - x_i (1 - \frac{x_i}{s})]^2. \end{aligned}$$

Thus, the time derivative $(d/dt)V(x(t))$ is strictly negative,

except at the unique rest point of the dynamics (7), where it equals zero. As we know from Theorem 2 that for any initial condition $x(0)$, the trajectory $x(t)$, $t \in \mathbb{R}_+$ is contained in a compact set, LaSalle's invariance theorem (see e.g. Khalil [12], p. 128, Theorem 4.4) can be applied to conclude that for any initial condition $x(0)$, $x(t)$ converges to the unique rest point identified in Theorem 1.

B. Global system: stability analysis

Assuming that the state of the $k-1$ st layer converges to a unique rest point, $(x_c^*)_{c \in \mathcal{C}_k}$, we denote by $\tilde{\lambda}_c$ the corresponding limiting arrival rate of type c -users for any $c \in \mathcal{C}_k$, which in view of (1) and (6) is given by

$$\tilde{\lambda}_c = \beta_c(x^*) = \lambda_c + \sum_{i \in c} x_{c^*i}^* \sum_{s \in \mathcal{C}_{k-1}, i \in s} \frac{x_s^*}{x_{s^*}^*} \frac{1}{|s \setminus c| + 1}.$$

We also define the corresponding condition

$$(D_k) \quad 2 \max_{c \in \mathcal{C}_k} \tilde{\lambda}_c < \sum_{c \in \mathcal{C}_k} \tilde{\lambda}_c.$$

Note that, in view of the preceding subsection, provided conditions $(D_1), \dots, (D_{k-1})$ hold, the rest points of layer $k-1$ exist and are unique, being characterized by the relations

$$\tilde{\lambda}_c = x_c^* \left(1 - \frac{x_c^*}{\sum_{c' \in \mathcal{C}_{k-1} x_{c'}^*} x_{c'}^*} \right), \quad c \in \mathcal{C}_{k-1},$$

and thus the rates $\tilde{\lambda}_c$, $c \in \mathcal{C}_k$ are indeed well defined. We then have the following

Theorem 4. *Assume that for all $k = 1, \dots, K-1$, condition (D_k) holds. Assume also that for all $c \in \mathcal{C}$, $\tilde{\lambda}_c > 0$. Then the dynamics (4) are globally asymptotically stable.*

Proof: We show by induction that the state of each layer converges asymptotically to its unique rest point. Assuming that this holds for layer $k-1$, then the arrival rates of type c -users, for all $c \in \mathcal{C}_k$, converge to $\tilde{\lambda}_c$. The state of layer k thus evolves according to

$$\frac{d}{dt} x_c(t) = \tilde{\lambda}_c + \epsilon_c(t) - x_c(t) \left(1 - \frac{x_c(t)}{s_k(t)} \right), \quad c \in \mathcal{C}_k, \quad (13)$$

where $s_k(t) := \sum_{c \in \mathcal{C}_k} x_c(t)$, and $\epsilon_c(t) \rightarrow 0$ as $t \rightarrow \infty$. These dynamics are the same as those of the single layer in isolation, described by (7), except for the additional vanishing inputs $\epsilon_c(t)$. The proof of Theorem 2 is easily adapted to show that, under condition (D_k) , any solution of (13) is uniformly bounded, i.e. there is a constant $A > 0$ such that $x_c(t) \leq A$ for large enough t . Also, since by assumption $\tilde{\lambda}_c > 0$ for all c , there exists a positive constant $a > 0$ such that, for large enough t , $x_c(t) \geq a$.

As in the proof of Theorem 3, introduce the function

$$V(x_c, c \in \mathcal{C}_k) = \sum_{c \in \mathcal{C}_k} \frac{\tilde{\lambda}_c}{x_c} + \log(x_c) - \log(s_k),$$

where $s_k = \sum_{c \in \mathcal{C}_k} x_c$. We then have that

$$\begin{aligned} \frac{d}{dt} V(x_c(t), c \in \mathcal{C}_k) = & - \sum_{c \in \mathcal{C}_k} \frac{1}{x_c^2} \left[\tilde{\lambda}_c - x_c(1 - x_c/s_k) \right]^2 \\ & + \sum_{c \in \mathcal{C}_k} \epsilon_c \left[\frac{-\tilde{\lambda}_c}{x_c^2} + \frac{1}{x_c} - \frac{1}{s_k} \right]. \end{aligned}$$

Since $x_c(t) \in [a, A]$ for large enough t , the second summation in the right-hand side of this equation goes to zero as t goes to infinity; denote it by $\epsilon(t)$. Let $\delta_{min} := V(x_c^*, c \in \mathcal{C}_k)$. Fix some $\delta > \delta_{min}$, and define

$$\phi(\delta) := \inf \left\{ \sum_{c \in \mathcal{C}_k} \frac{1}{x_c^2} \left[\tilde{\lambda}_c - x_c(1 - x_c/s_k) \right]^2 \right\}$$

where the infimum is taken over the $x_c, c \in \mathcal{C}_k$ such that $x_c \in [a, A]$ and $V(x) \geq \delta$. It thus holds that, for t large enough, if $V(x_c, c \in \mathcal{C}_k) \geq \delta$,

$$\frac{d}{dt} V(x_c(t), c \in \mathcal{C}_k) \leq -\phi(\delta) + \epsilon(t).$$

As $\phi(\delta) > 0$ for $\delta > \delta_{min}$, it follows that necessarily $V(x(t)) \leq \delta$ for t large enough. Since $\delta > \delta_{min}$ was chosen arbitrarily, this guarantees that $x(t)$ converges to a minimizer of V . As $(x_c^*, c \in \mathcal{C}_k)$ is the unique such minimizer, the announced convergence holds.

With some more work it can be shown that, when some condition (D_k) fails, given that the conditions $(D_1), \dots, (D_{k-1})$ hold, then for $c \in \mathcal{C}_k$ such that $\tilde{\lambda}_c = \sup_{c' \in \mathcal{C}_k} \tilde{\lambda}_{c'}$, the component $x_c(t)$ goes to infinity as $t \rightarrow \infty$. We omit the argument for lack of space. The main application of the previous theorem is the following

Corollary 1. *When there are no exogenous inputs to layers $2, \dots, K-1$, i.e. $\lambda_c = 0$ for $|c| \geq 2$, and the exogenous inputs to layer 1 are symmetrical, i.e. $\lambda_1 = \lambda_2 = \dots = \lambda_K > 0$, all solutions $x(t)$ of (4) satisfy:*

$$\lim_{t \rightarrow +\infty} x_c(t) = \frac{K\lambda_1}{\binom{K}{i} - 1}, \quad c \in \mathcal{C}_i, \quad i = 1, \dots, K-1. \quad (14)$$

C. Sojourn times

We define the sojourn time T_i within layer i as the reciprocal of the rate of successful encounters within layer i , evaluated at the rest point of the differential systems (4). We expect such sojourn times to be close to the expected sojourn time of users within layer i in the original stochastic system, when in equilibrium³.

In this section, we consider only the symmetrical setup of the previous corollary. By Proposition 1, this achieves the smallest sojourn time over all systems with given total arrival rate $\sum_{i=1}^K \lambda_i$. With the above definition of sojourn times T_i , it follows from Little's theorem that

$$(\lambda_1 K) T_i = \sum_{c \in \mathcal{C}_i} x_c^*, \quad i = 1, 2, \dots, K-1,$$

³Note that convergence of expected sojourn times in the stochastic system with arrival rates $\lambda_c^N = N\lambda_c$ to those as defined above does not follow directly from the results of Kurtz [13]. Indeed, these establish convergence of processes over finite intervals, and not convergence of equilibrium quantities. Proving such convergence is beyond the scope of the present work.

and thus, in view of (14),

$$T_i = 1 + \frac{1}{\binom{K}{i} - 1}, \quad i = 1, 2, \dots, K-1. \quad (15)$$

It follows that $1 < T_i \leq K/(K-1)$ for all $i = 1, 2, \dots, K-1$, and the normalized total sojourn time $\sum_{i=1}^{K-1} T_i/(K-1)$ also lies in the interval $[1, K/(K-1)]$. Thus in the large K limit,

$$\lim_{K \rightarrow +\infty} \frac{\sum_{i=1}^{K-1} T_i}{K-1} = 1.$$

As the per layer sojourn is necessarily larger than 1, this shows that the time to complete the whole collection is asymptotically optimal in the large K limit.

IV. OPEN SYSTEMS: THE FLAT CASE

We now focus on the flat system, i.e. the differential system (4), with encounter probabilities p_{cs} given by (2), corresponding to users sampling uniformly from the total user population to select their targets.

We confine our analysis to a particular symmetrical case where λ_c depends only on the cardinality of c , i.e. there exist λ'_k , $k = 1, \dots, K-1$, such that $\lambda_c = \lambda'_{|c|}$. We further assume symmetry of the initial conditions, i.e. $x_c(0) = x_{c'}(0)$ whenever $|c| = |c'|$. The dynamics (4) then simplify as follows (proof omitted for the benefit of space):

Lemma 1. *Assume $\lambda_c = \lambda'_i$ and $x_c(0) = y_i(0)$, for all $c \in \mathcal{C}_i$, $i = 1, 2, \dots, K-1$. Then, $x_c(t) = y_i(t)$, $t \geq 0$, for all $c \subseteq \mathcal{C}_i$, where for $i = 1, 2, \dots, K-1$,*

$$\begin{aligned} \frac{d}{dt} y_i(t) = & \lambda'_i + \frac{i}{K-(i-1)} \left(1 - \frac{s_{i-2}(t)+y_{i-1}(t)}{s_{K-1}(t)} \right) y_{i-1}(t) \\ & - \left(1 - \frac{s_{i-1}(t)+y_i(t)}{s_{K-1}(t)} \right) y_i(t), \end{aligned} \quad (16)$$

with $s_i(t) := \sum_{k=1}^i \binom{i+1}{k} y_k(t)$, $y_0 \equiv 0$, and $s_0 \equiv 0$.

We then have the following

Lemma 2. *Let $y(t)$ be a solution of (16). Let $y^+(0) = y(0)$, and for $i = 1, 2, \dots, K-1$,*

$$\frac{d}{dt} y_i^+(t) = \lambda'_i + \frac{i}{K-(i-1)} y_{i-1}^+(t) - \frac{K-i}{K} y_i^+(t). \quad (17)$$

Then, $y(t) \leq y^+(t)$, for all $t > 0$.

The proof is given in Appendix C. An important consequence of this lemma is the following

Proposition 2. *The solution of (16) is uniformly bounded. More precisely, for any initial condition $y(0)$, for all $i = 1, 2, \dots, K-1$, we have*

$$\limsup_{t \rightarrow +\infty} y_i(t) \leq i! \sum_{j=1}^i \frac{K^{i-j+1}}{j!(K-j)} \lambda'_j < \infty.$$

The proof is by induction over i ; we omit for lack of space.

A. Rest points

Let $y = (y_1, y_2, \dots, y_{K-1})$ be a rest point of (16). The flow of users arriving to layers $1, 2, \dots, i$ equals the flow of users

leaving layer i , hence

$$\sum_{j=1}^i \binom{K}{j} \lambda'_j = \binom{K}{i} \left(1 - \frac{s_{i-1} + y_i}{s_{K-1}} \right) y_i.$$

We introduce the notations $\tilde{y}_i = y_i/s_{K-1}$, $\tilde{s}_i = s_i/s_{K-1}$, and $\tilde{\lambda}_i = \lambda'_i/s_{K-1}$, $i = 1, 2, \dots, K-1$. The above equation may thus be written

$$\tilde{y}_i^2 - (1 - \tilde{s}_{i-1}) \tilde{y}_i + \sum_{j=1}^i \frac{\binom{K}{j}}{\binom{K}{i}} \tilde{\lambda}_j = 0. \quad (18)$$

This entails the following

Proposition 3 (Existence). *Consider (16) with $\lambda'_2 = \lambda'_3 = \dots = \lambda'_{K-1} = 0$. There exists a rest point given by*

$$\tilde{y}_i = \frac{1}{2} \left((1 - \tilde{s}_{i-1}) - \sqrt{(1 - \tilde{s}_{i-1})^2 - \frac{4}{\binom{K}{i}} \frac{1}{T}} \right).$$

where the parameter $T := s_{K-1}/(K\lambda'_1)$ is defined by the relation $\tilde{s}_{K-1} = 1$.

It should be noted that T , being the ratio of the number of users in the system to the arrival rate, can be interpreted as the sojourn time of a user in the system.

Although we suspect that, under the assumptions of the proposition, the identified rest point is the unique one for the dynamics (16), and that these are globally asymptotically stable, we have not been able to prove these properties yet.

Proof: The asserted rest point satisfies the quadratic equation (18). Now, note that $\tilde{s}_{K-1}(T)$ is an increasing function of $1/T$. Indeed, the sequence $\{\tilde{s}_i\}_{1 \leq i < K}$ is defined by $\tilde{s}_0 = 0$, and for $i = 1, 2, \dots, K-1$,

$$\tilde{s}_i = \frac{1}{2} \sum_{j=1}^{i-1} \binom{i}{j} \left(1 - \tilde{s}_{i-1} - \sqrt{(1 - \tilde{s}_{i-1})^2 - \frac{4}{\binom{K}{i}} \frac{1}{T}} \right).$$

For $i = 1$, we have $\tilde{s}_1 = (1 - \sqrt{1 - 4/(KT)})/2$, which is clearly increasing with $1/T$, for $T \geq 4/K$. Suppose now that \tilde{s}_n is increasing with $1/T$, for all $1 < n < K-1$. Combine this with the fact that \tilde{s}_{n+1} is increasing with both \tilde{s}_n and $1/T$, hence, the result follows by induction.

B. Sojourn times

In the remainder of this section we place ourselves under the symmetry assumptions of Proposition 3, and assume $\lambda'_1 > 0$ and $\lambda'_2 = \dots = \lambda'_{K-1} = 0$. We consider the evaluation of sojourn times, as previously defined, at some arbitrary rest point of dynamics (16). In particular we do not rely on our conjecture that the rest point identified in Proposition 3 is unique.

By definition of sojourn times, and from (16), it follows:

$$\frac{1}{T_i} = 1 - \sum_{j=1}^i \binom{i}{j} \tilde{y}_j, \quad i = 1, 2, \dots, K-1. \quad (19)$$

Indeed, the right-hand side is the rate of successful encounters, for an individual holding a collection of i coupons: $\binom{K}{j} \tilde{y}_j$ is the probability that the encountered individual has a collection of j coupons and the remaining factor $\binom{i}{j}/\binom{K}{j}$ is the

probability that the encounter's collection of j coupons is a subset of the instigator's collection, given that the instigator has i coupons. Thus, the sum in (19) is the probability that the target user's collection is a subset of the instigator's collection, given that the instigator has i coupons. The last event is by definition an unsuccessful encounter.

Note now, by Little's formula (or by direct verification from the fixed point relations) that the sojourn time T_i is related to the total sojourn time $T = \sum_{j=1}^{K-1} T_j$ via

$$T_i = \tilde{y}_i \binom{K}{i} T, \quad i = 1, 2, \dots, K-1.$$

Plugging this relation into (19) yields

$$\frac{1}{T_i} = 1 - \sum_{j=1}^i \frac{\binom{i}{j} T_j}{\binom{K}{j} T}, \quad i = 1, 2, \dots, K-1. \quad (20)$$

This relation is the key to the following proposition, providing evaluations of the sojourn times T_i .

Proposition 4. *The sojourn times are such that $1 < T_1 < T_2 < \dots < T_{K-1} < 2$. Thus the success rate of attempted encounters is at least one half within each layer. Furthermore, the total sojourn time T satisfies*

$$T \leq K + O(\sqrt{K}) \quad (21)$$

as $K \rightarrow \infty$. It is in this sense asymptotically optimal.

Proof: It is immediate from (20) that the sequence $\{T_i\}$ is increasing. It also follows from (20) that $T_i > 1$, for all $i = 1, 2, \dots, K-1$. We next prove the asserted upper bound. For $i = K-1$, (20) can be rewritten as

$$\frac{1}{T_{K-1}} = 1 - \sum_{i=1}^{K-1} a_i t_i,$$

where $a_i = 1 - i/K$, and $t_i = T_i/T$. As the sequences $\{a_i\}$ and $\{t_i\}$ are both non-decreasing, Theorem 43 in Hardy, Littlewood and Pólya [8] guarantees the inequality in the following derivation:

$$\begin{aligned} \frac{1}{T_{K-1}} &> 1 - (K-1) \left(\frac{1}{K-1} \sum_{i=1}^{K-1} a_i \right) \left(\frac{1}{K-1} \sum_{i=1}^{K-1} t_i \right) \\ &= \frac{1}{2}. \end{aligned}$$

The asserted bound of 2 on the T_i 's follows.

The first part of the proposition ensures that $(K-1)T_j/T \in [1/2, 2]$ for all j . This together with (20) implies the inequality in the following:

$$\frac{1}{T_i} \geq 1 - \frac{2}{K-1} \sum_{j=1}^i \frac{\binom{i}{j}}{\binom{K}{j}} = 1 - \frac{2}{K-1} \sum_{j=1}^i \frac{\binom{K-j}{i-j}}{\binom{K}{i}},$$

where the equality follows by elementary manipulations. Applying the parallel summation formula (see e.g. Graham et al. [9], p. 174), this simplifies to

$$\frac{1}{T_i} \geq 1 - \frac{2}{K-1} \frac{\binom{K}{i-1}}{\binom{K}{i}} = 1 - \frac{2}{K-1} \frac{i}{K-i+1}.$$

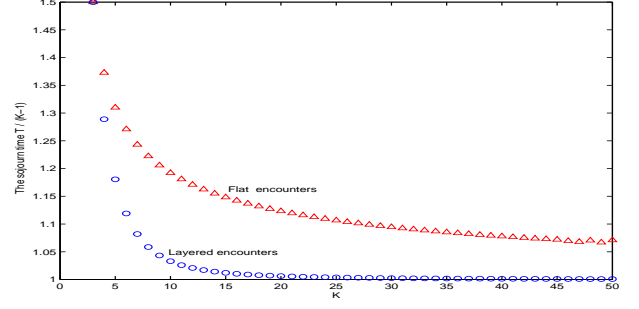


Fig. 2: Normalized sojourn time $\sum_{i=1}^{K-1} T_i / (K-1)$ for the open system with flat and layered encounters.




Fig. 3: Numerical values for the sojourn times T_i , $i = 1, 2, \dots, K-1$, $K = 10, 20, 30, 40$, for the open system with flat and layered encounters.

Select $k \leq K-1$, such that $K-1-k \sim A\sqrt{K}$ for some constant $A > 0$. We then have

$$\begin{aligned} T &\leq \sum_{i=1}^k \left[1 - \frac{2}{K-1} \frac{i}{K-i+1} \right]^{-1} + 2(K-1-k) \\ &\leq 2(K-1-k) + k \left\{ 1 - \frac{2k}{[(K-1)(K-k+1)]} \right\}^{-1} \\ &= 2A\sqrt{K} + (K-1-A\sqrt{K}) \left(1 + \frac{1}{A\sqrt{K}} \right) + o(\sqrt{K}) \\ &= K-1 + (2A + \frac{1}{A})\sqrt{K} + o(\sqrt{K}). \end{aligned}$$

The announced result follows.

Numerical example. We provide numerical values for the mean sojourn times in each layer for both layered and flat system using Equation (15) and Equation (20), respectively. In Figure 2, we show total mean sojourn time normalized by the number of layers for varying K . The results indicate better performance for layered system, but this vanishes for large K . In Figure 3, we show mean sojourn time per layer for both layered and flat system. The first few coupons are obtained faster in the flat system while the opposite holds for the subsequent coupons. The latter effect dominates the overall mean sojourn time. The reason for this behaviour is that many encounters are wasted in flat systems due to users with many coupons trying to pull from users with a few coupons.

Our present analysis of flat systems is not as complete as that of layered systems. However it gives support to the belief that, as layered systems, flat systems settle to an equilibrium in the large arrival rate limit, and the sojourn times under this equilibrium are near optimal when the collection size K is large.

V. ULTIMATE LEFT-OVER IN A CLOSED SYSTEM

In the present section, we consider the same stochastic system as previously, but without exogenous user arrivals. It is easily seen that, when the system contains users of only one type c , no further transitions can take place. The performance measure of interest is then the size of the population ultimately remaining in the system, or put differently, the left-over.

At this stage we are able to analyze only a specific type of initial population distributions, namely with users only in layer $K - 1$. Under this assumption, denoting by $X_k(t)$ the number of users in the system at time t who have all coupons except for coupon k , then the $X_k(t)$ constitute a \mathbb{Z}_+^K -valued Markov process, with non-zero transition rates

$$x \rightarrow x - e_i \quad \text{with rate } x_i \left(1 - \frac{x_i}{\sum_{j=1}^K x_j}\right), \quad (22)$$

under both the flat and layered encounter mechanisms, where $x = (x_1, \dots, x_K)$ and e_i denotes the i -th unit vector. In the sequel, the initial population sizes are denoted by X_i . Without loss of generality we assume that $X_1 \geq \dots \geq X_K$.

As previously, we first let the initial condition scale to infinity, with K fixed, by setting $X_i \sim Nx_i$, where $x = (x_1, \dots, x_K)$ is fixed, and N tends to infinity. We denote by $X^{(N)}(t)$ the corresponding state at time t . We are still in the framework of density dependent Markov jump processes of Kurtz [13], the results of which imply that for all finite $T > 0$, Equation (3) holds, where convergence in the above equation is almost surely, and $x(t)$ denotes the solution of the system of ordinary differential equations (ODE)

$$\frac{d}{dt} x_i(t) = -x_i(t) \left(1 - \frac{x_i(t)}{\sum_j x_j(t)}\right), \quad (23)$$

with initial condition $x(0) = x$. From this, we can establish the following result (proof in Appendix D).

Theorem 5. *For initial population sizes $X_i = Nx_i$, $i = 1, \dots, K$, arranged in decreasing order, denoting by $X_i^{(N)}(\infty)$ the final number of type i -individuals remaining in the system, we have*

$$\lim_{N \rightarrow \infty} \frac{1}{N} X_i^{(N)}(\infty) = \begin{cases} x_1 \prod_{i=2}^K \left(1 - \frac{x_i}{x_1}\right) & \text{if } i = 1, \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

where convergence is almost sure.

The expression in (24) captures the following interesting phenomenon: increasing the number of populations, K , has a dramatic impact on the remaining number of individuals in the system, in that the latter decreases geometrically in K . As an example, assuming $x_1 > x_2 = \dots = x_K$, the right-hand side in (24) reads $x_1(1 - x_1/x_2)^K$. For fixed x_1 and x_2 , this expression goes to zero exponentially fast with K . In fact, it suggests (heuristically), by writing $K = \alpha \log(N)$, and $x_1/x_2 = b$, that the total left-over in the system is of order

$$Nx_1(1 - x_1/x_2)^K = x_1 N^{1+\alpha \log(1-b)}.$$

This is only a heuristic reasoning, as the above theorem assumes K is fixed. It nevertheless suggests that some qualitative

change in the left-over occurs when K is of order $\alpha \log(N)$, and when α crosses the value $1/|\log(1-b)|$. The next result provides bounds on the left-over, which capture this intuition.

Theorem 6. *Given initial conditions $X_1 \geq \dots \geq X_K$, introduce the notations $X_1 = N$, $X_2 + \dots + X_K = \beta(K-1)N$. Assume that $K \sim \alpha \log(N)$, and that $\beta \rightarrow b \in (0, 1)$ as $N \rightarrow \infty$. Assume further that*

$$\alpha \log \left(\frac{1}{1-b} \right) > 1. \quad (25)$$

Then

$$\lim_{N \rightarrow \infty} \mathbf{P} \left(\sum_{i=1}^K X_i(\infty) \geq \theta \log(K) \right) = 0 \quad (26)$$

for $\theta > 0$ large enough.

In words, when $K/\log(N)$ is above the critical threshold $1/|\log(1-b)|$, the total left-over is of order $\log(\log(N))$. Before stating the proof, we need some auxiliary results.

In the sequel, population 1 plays a special role, as we aim at bounding the remainder in this population specifically. It is convenient to modify the original transition rates, as in (22), by multiplying them by $y := \sum_{k=1}^K x_k / \sum_{k=2}^K x_k$. This amounts to make a state-dependent time rescaling, and thus does not affect the limiting state of the process. The new transition rates are then given by

$$\begin{aligned} x &\rightarrow x - e_1 && \text{with rate } x_1 \mathbf{1}_{\sum_{j=2}^K x_j > 0}, \\ x &\rightarrow x - e_i && \text{with rate } x_i \left(1 + \frac{x_1 - x_i}{\sum_{k=2}^K x_k}\right), \quad 1 < i \leq K. \end{aligned} \quad (27)$$

We now introduce an auxiliary, two dimensional Markov process, denoted by $X'(t) := (X'_1(t), Y'(t))$, defined by the following transition rates:

$$\begin{aligned} (x'_1, y') &\rightarrow (x'_1 - 1, y') && \text{with rate } x'_1 \mathbf{1}_{y' > 0}, \\ (x'_1, y') &\rightarrow (x'_1, y' - 1) && \text{with rate } \frac{K-2}{K-1} y' + x'_1. \end{aligned} \quad (28)$$

The interpretation of this process is as follows: it is meant to represent dynamics similar to the previous ones, X'_1 playing the role of X_1 , and Y' playing the role of the aggregate $X_2 + \dots + X_K$, but with transition rates chosen to be as if the population sizes X_2, \dots, X_K were perfectly balanced, and all equal. It turns out that forcing this balance property makes it harder for type 1-individuals to leave the system, as is captured by the following lemma, whose proof is deferred to Appendix E.

Lemma 3. *Consider the two Markov processes $X = (X_1, \dots, X_K)$ and $X' = (X_1, Y)$, with initial conditions $X'_1 = X_1$, and $Y' = \sum_{k=2}^K X_k$, and respective transition rates (27) and (28). It is then possible to construct the two processes on the same probability space so that for all $t \geq 0$, the following orderings hold:*

$$X'_1(t) \geq X_1(t), \quad Y'(t) \leq Y(t) := \sum_{k=2}^K X_k(t). \quad (29)$$

We may thus obtain upper bounds on $X_1(\infty)$ by studying the value of $X'_1(T)$ at the time T when Y' reaches 0. The next lemma provides the main results we shall need about process

$X' = (X'_1, Y')$. Its proof is in Appendix F.

Lemma 4. Consider the Markov process $X' = (X'_1, Y')$ with initial conditions $X'_1(0) = N$, and $Y'(0) = Y$. Introduce the notation

$$r := \frac{K-2}{K-1}. \quad (30)$$

Then its state at a given time $t > 0$ can be represented as

$$X'_1(t) = U, \quad Y'(t) = V - W, \quad (31)$$

where the random variables U, V are binomially distributed, with respective distributions:

$$U \sim \text{Bin}(N, \exp(-t)), \quad V \sim \text{Bin}(Y, \exp(-rt)), \quad (32)$$

and the random variable W has mean given by:

$$\mathbf{E}[W] = N \frac{e^{-rt} - e^{-t}}{1-r}. \quad (33)$$

Furthermore, for $K \geq 3$, all three random variables U, V, W satisfy large deviations bounds of the type

$$\mathbf{P}(|Z - \mathbf{E}(Z)| \geq \epsilon \mathbf{E}(Z)) \leq 4e^{-c_\epsilon \mathbf{E}(Z)}, \quad (34)$$

where $c_\epsilon > 0$ for all $\epsilon > 0$.

We are now ready to prove Theorem 6. *Proof:* [Theorem 6] Since only one component $X_i(\infty)$ is non zero, one has

$$\mathbf{P}\left(\sum_{i=1}^K X_i(\infty) > A\right) = \sum_{i=1}^K \mathbf{P}(X_i(\infty) > A).$$

Let us consider the first term in this sum. In view of the preceding lemmas, one has

$$\mathbf{P}(X_1(\infty) \leq A) \geq \mathbf{P}(U \leq A, V - W > 0),$$

where the r.v. U, V, W correspond to any suitable choice of t . Thus,

$$\mathbf{P}(X_1(\infty) > A) \leq \mathbf{P}(U > A) + \mathbf{P}(V - W \leq 0).$$

Take $t = \log(N/\theta \log(K))$, for some $\theta > 0$. One then has:

$$\begin{aligned} \mathbf{E}(U) &= Ne^{-t} = \theta \log(K), \\ \mathbf{E}(V) &= \beta(K-1)\theta \log(K)e^{t/(K-1)} \sim \beta\theta K \log(K)e^{1/\alpha} \\ \mathbf{E}(W) &= (K-1)\theta \log(K) \left(e^{t/(K-1)} - 1\right) \\ &\sim \theta K \log(K) \left(e^{1/\alpha} - 1\right). \end{aligned}$$

For any $\epsilon > 0$, by (34) it holds that

$$\begin{aligned} \mathbf{P}(U > (1+\epsilon)\mathbf{E}(U)) &\leq 2e^{-c_\epsilon \mathbf{E}(U)} = 2K^{-\theta c_\epsilon}, \\ \mathbf{P}(V < (1-\epsilon)\mathbf{E}(V)) &\leq 2e^{-c_\epsilon \mathbf{E}(V)} = o(K^{-1}), \\ \mathbf{P}(W > (1+\epsilon)\mathbf{E}(W)) &\leq 4e^{-c_\epsilon \mathbf{E}(W)} = o(K^{-1}). \end{aligned}$$

These estimations, together with assumption (25), entail that

$$\mathbf{P}(X_1(\infty) \geq \theta \log(K)) = o(K^{-1})$$

for $\theta > 0$ sufficiently large.

One can apply exactly the same method so as to obtain upper bounds on the probability that $X_i(\infty)$ exceeds $\theta \log(K)$. In fact, these upper bounds can be made tighter as $X_1(0)$ is the

largest component. By summing such evaluations, one arrives at

$$\mathbf{P}\left(\sum_{i=1}^K X_i(\infty) > \theta \log(K)\right) = o(1)$$

for $\theta > 0$ large enough. The conclusion (26) of the theorem follows.

VI. DISCUSSION

Simple models under symmetry. The system (4) and its particular instances for layer and flat open systems, studied respectively in Section III and Section IV, have dimensionality $2(2^{K-1} - 1)$. In practise, for peer-to-peer file sharing applications, some data files are broken down into many pieces, for example, in the order of thousand.⁴ This dimensionality is brought down to $K-1$ under symmetry assumptions discussed in this section. The special cases that we are about to consider enable us to compare our systems with that of related work by Yang and Veciana [17] and Qiu and Srikant [7], which we do later in this section.

We first consider open system with layered peering, under the symmetry assumptions: (i) initial user population is such that an equal number of users hold a coupon collection of a given size (i.e. $x_c(0) = x_{c'}(0)$, for any c, c' for which $|c| = |c'|$), and (ii) users enter the system holding a coupon i with rate λ_1 that is independent of i . The former, in particular, accommodates the special case of an initially empty system, e.g. beginning of a file dissemination. The latter is non restrictive in practise, as servers would offer initial coupon to users with no bias towards a specific coupon.

We denote with n_k the number of users holding a collection of k coupons, which relates to the existing notation as $n_k = \binom{K}{k} x_c$, $|c| = k$. The layer system dynamics is described by the system of linear differential equations⁵, for all $k = 1, 2, \dots, K-1$:

$$\frac{d}{dt} n_k = \lambda_k + a_{k-1} n_{k-1} - a_k n_k \quad (35)$$

where λ_k may depend on time t and

$$a_k = 1 - \frac{1}{\binom{K}{k}}.$$

Note that $1/a_k$ is the sojourn time in layer k .

It is noteworthy that, under the prevailing symmetry assumptions, the sojourn time in each layer is time invariant. This is intuitive as under our symmetry assumptions, the perfect balance of the number of users holding distinct collections within a layer is preserved over time for each layer. Note that this holds even for time-dependent rate of new user arrivals.

It may be of interest to note the following explicit characterization for the total number of users in an open layer system, under the large population limit and prevailing symmetry assumptions of this section. This characterization may be of

⁴For concreteness, consider a movie file of 2 GB and typical piece size of 0.5 MB, which amounts to 4000 distinct pieces.

⁵We phrase this slightly more generally by allowing new users to join with any coupon collection subject only to the symmetry condition that the rate of users that enter with a collection c depends on c only through $|c|$.

interest for numerical computations. Denote with $n(t)$, the total number of users at time t , i.e. $n(t) := \sum_{k=1}^{K-1} n_k(t)$, $t \geq 0$. Note that the sojourn time for a user is a sum of exponentially distributed random variables with means $1/a_k$, $k = 1, 2, \dots, K-1$. Denote the distribution of a user sojourn time with F . Let further $\lambda(t)$ be the rate of user arrivals at time $t \geq 0$. We then have,

$$n(t) = n(0)\bar{G}(t) + \int_0^t \bar{F}(t-s)\lambda(s)ds, \quad t \geq 0,$$

where G is the residual sojourn time distribution for users that are in the system at time $t = 0$, and $\bar{G}(t) := 1 - G(t)$ with analogous definition for \bar{F} . Note that F is a convolution of exponential distributions with parameters a_k , $k = 1, 2, \dots, K-1$. Recall that $a_k = 1 - 1/\binom{K}{k}$ and thus $a_k \sim 1$, for large K . This tells us that for large K , F is well approximated by a normal distribution with mean and variance equal to $K-1$.

We now similarly consider open flat systems under our symmetry assumptions. The dynamics is described by the system of ordinary differential equations, for $k = 1, 2, \dots, K-1$:

$$\frac{d}{dt}n_k = \lambda_k + a_{k-1}(n)n_{k-1} - a_k(n)n_k \quad (36)$$

where $a_0(0) = 0$ and

$$a_k(n) = 1 - \frac{\sum_{i=1}^k \binom{k}{i} n_i}{\sum_{i=1}^{K-1} n_i}, \quad k = 1, 2, \dots, K-1.$$

Note that, unlike for the layer case, the system is non-linear and sojourn times are in general time variant.

Related work. We discuss the model of Yang and de Veciana [17] and its extension by Qiu and Srikant [7]. These models appears to be *a priori* postulated. In contrast, our systems derive from the underlying replication strategies.

Yang and de Veciana [17] introduced the following model for BitTorrent-like systems. Denote with $N_d(t)$ the total number of downloaders in the system at time t and with N_s , the number of users that have finished their downloads, and then remained in the system uploading to other users (seeds). The state (N_d, N_s) is assumed to evolve as a Markov process with the following transition rates:

$$(N_d, N_s) \rightarrow \begin{cases} (N_d + 1, N_s) & \text{rate } \lambda \\ (N_d - 1, N_s + 1) & \text{rate } \mu(\eta N_d + N_s) \\ (N_d, N_s - 1) & \text{rate } \sigma N_s \end{cases} \quad (37)$$

The parameter μ is interpreted as the per user upload capacity and η specifies the utilization of the user upload capacity achieved by the downloaders, so that $\mu(\eta N_d(t) + N_s(t))$ is the total effective capacity at time t . The utilization parameter η is fixed and related to the number of coupons, K , in some way to reflect an underlying replication strategy. Each user remains in the system as a seed for a random time, assumed to have exponential distribution with mean $1/\sigma$.

Qiu and Srikant [7] consider the evolution of the number of downloaders and seeds, (n_d, n_s) , specified by the following

differential equations:

$$\frac{d}{dt}n_d = \lambda - \theta n_d - \min(cn_d, \mu(\eta n_d + n_s)) \quad (38)$$

$$\frac{d}{dt}n_s = \min(cn_d, \mu(\eta n_d + n_s)) - \sigma n_s \quad (39)$$

where new parameters are the per user download abort rate θ and the per user download capacity c . The appeal of the model is its low dimensionality.

The model is an extension of [17], as for the choice of parameters $(\theta, c) = (0, +\infty)$, it is precisely the large population limit of (37); in the precise limit sense as described in Section II and with appropriately scaled original rate λ . In [7], the utilization parameter η is approximated by $1 - \log(K)/K$ for a collection of K coupons, which is motivated in [7] under assumption (H) that the number of users holding a collection of k coupons is balanced with respect to k .

In order to compare with our model, suppose now that users instantly leave upon finishing their downloads (i.e. $\sigma = +\infty$), then (38)–(39) boils down to

$$\frac{d}{dt}n = \lambda - \eta n \quad (40)$$

where, with no loss of generality, we consider users with unit upload capacity ($\mu = 1$) and omit the subscript for the counter of the downloaders. The question is whether there exist cases and what these cases are for which (35) and (36) conform to or are well approximated by (40). We now consider (35) and with a slight abuse of notation denote $n = \sum_{k=1}^{K-1} n_k$. It follows from (35)

$$\frac{d}{dt}n = \lambda - \left(1 - \frac{1}{K}\right)n_{K-1}$$

where $\lambda := \sum_{k=1}^{K-1} \lambda_k$. In general, the last dynamics conforms to (40) only for the special case $K = 2$, as otherwise, η depends on the state as $(1 - 1/K)n_{K-1}/n$. In view of the assumption H, note that we found that at equilibrium, n_k/n is *not* independent on k , but increases from the center towards the end points 1 and $K-1$. Moreover, in the equilibrium, $n_k/n \sim 1/(K-1)$, for large K . This suggests that in (40) η would be of the order $1/K$. Note, however, that this does not conform to the suggestion in [7] to approximate η with $1 - \log(K)/K$. In summary, we find that (40) describes the dynamics of the open layer system only for the special case of two distinct coupons, i.e. for $K = 2$.

We now similarly consider the open flat dynamics (36), from which it follows:

$$\frac{d}{dt}n = \lambda - \left(\frac{1}{K} \sum_{k=1}^{K-1} k \frac{n_k}{n}\right)n_{K-1}.$$

Again, the dynamics conforms to (40) only for $K = 2$ and otherwise the utilization parameter η is state dependent, equal to $\left((1/K) \sum_{k=1}^{K-1} k(n_k/n)\right)(n_{K-1}/n)$. The assumption H is again not verified, as we found that in the equilibrium of open flat systems n_k/n increases with k .

Extensions. The model (4) can be extended in a straightforward way to accommodate both time-dependent user arrival rate and users remaining in the system as seeds.

VII. CONCLUSION

In this paper, we introduced coupon replication systems as models of file swarming. We analyzed their stability properties in the presence of exogenous arrivals, for both flat and layered encounter mechanisms, in a large population limit. The results suggest that file swarming does stabilize around a finite equilibrium point, no matter how large the arrival rate, thereby giving additional support to the claims of Yang and de Veciana [17]. Future work could aim to strengthen our conclusions for flat encounters. Another promising direction is to relax the assumption that users always enter the system with at least one coupon, which implicitly assumes that some centralized server has a capacity of ρ/K , where ρ is the work arrival rate, and K the size of the coupon collection. This extension fits into our framework, by explicitly accounting for the number of users in the system holding no coupon at all.

We also analyzed the left-over in a closed system, under the assumption that all users initially have only one missing coupon. In doing so, we developed explicit solutions of the ODE (23), which is a special instance of Lotka-Volterra system (e.g. see Hofbauer and Sigmund [10]), which may be of independent interest. We then used stochastic comparison results, combined with large deviations estimations, to have finer bounds on the order of the left-over when K is allowed to scale as $\log(N)$, and N denotes the largest initial population, thereby confirming that for $K/\log(N)$ above a critical threshold, the left-over is small. These results provide analytical support to the belief that file swarming performs well under flash crowds, even when rarest first is replaced by a random coupon selection strategy. More work is needed to analyze the left-over for more general initial conditions.

REFERENCES

- [1] eDonkey. <http://www.edonkey2000.com/index.html>.
- [2] KaZaA. <http://www.kazaa.com/>.
- [3] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. Wiley Series in Probability and Statistics, 2002.
- [4] N. Alon and J. Spencer. *The probabilistic method*. Wiley Interscience Series in Discrete Mathematics and Optimization, 2 edition, 2000.
- [5] Bram Cohen. Incentives build robustness in bittorrent, May 2003.
- [6] Supratim Deb and Muriel Médard. Algebraic Gossip: A Network Coding Approach to Optimal Multiple Rumor Mongering, April 2004. preprint.
- [7] Dongyu Qiu and R. Srikant. Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks. In *Proc. of ACM Sigcomm 2004*, Portland, Oregon, 2004.
- [8] G. Hardy and J. E. Littlewood and G. Pólya. *Inequalities*. Cambridge Mathematical Library, 2 edition, 1952.
- [9] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Addison Wesley, 1998.
- [10] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [11] J. A. Pouwelse and P. Garbacki and D. H. J. Epema and H. J. Sips. A Measurement Study of the BitTorrent Peer-to-Peer File-Sharing System, April 2004. preprint.
- [12] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2002.
- [13] Thomas Kurtz. *Approximation of Population Processes*, volume 36. CBMS-NSF Regional Conference Series in Applied Mathematics, 1981.
- [14] R. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. of Cambr. Phil. Soc.*, 58(3):497–520, 1962.
- [15] M. Izal and G. Urvoy-Keller and E. W. Biersack and P. A. Felber and A. Al Hamra and L. Garcés-Erice. Dissecting BitTorrent: Five Months in a Torrent's Life. In *Proc. of PAM 2004*, Antibes, France, 2004.
- [16] Jochen Munding and Richard Weber. Efficient File Dissemination using Peer-to-Peer Technology. Technical Report 2004-01, University of Cambridge, UK, 2004.

- [17] Xiangying Yang and Gustavo de Veciana. Service Capacity of Peer to Peer Networks. In *Proc. of Infocom 2004*, Hong-Kong, China, 2004.

APPENDIX

A. Proof of Theorem 1

Let us first establish that any rest point x of (7), if it exists, is given either by (8) or by (10). First, considering $s = \sum_{i=1}^K x_i$ as a parameter, the equation $dx_i/dt = 0$ is quadratic in x_i , and its solution must satisfy

$$x_i = \frac{s}{2} \left(1 \pm \sqrt{1 - \frac{4\lambda_i}{s}} \right), \quad i = 1, 2, \dots, K. \quad (41)$$

This also shows that necessarily, $s \geq 4\lambda_1$. Next, since $x_i \geq 0$ for all i , if for some m , $x_m > s/2$, then $x_j < s/2$, for all $j \neq m$. Hence, we have either

$$x_i = \frac{s}{2} \left(1 - \sqrt{1 - \frac{4\lambda_i}{s}} \right), \quad i = 1, 2, \dots, K, \quad (42)$$

or, for some $m \in \{1, 2, \dots, K\}$,

$$x_m = \begin{cases} \frac{s}{2} \left(1 + \sqrt{1 - \frac{4\lambda_i}{s}} \right) & i = m \\ \frac{s}{2} \left(1 - \sqrt{1 - \frac{4\lambda_i}{s}} \right) & i \neq m. \end{cases} \quad (43)$$

It remains to show that in this latter case, necessarily one must have $m = 1$. To this end, assume that there exists x given by (43) for some $1 < m \leq K$. The definition $s = \sum_{i=1}^K x_i$, with x given by (43) and the change of variable $u = 1/s$ can be rewritten as $f_m(u) = K - 2$, where

$$f_m(u) := \sum_{i \neq m} \sqrt{1 - 4\lambda_i u} - \sqrt{1 - 4\lambda_m u}, \quad 0 \leq u \leq 1/(4\lambda_1). \quad (44)$$

The function $f_m(\cdot)$ is such that $f_m(0) = K - 2$ and $f'_m(u) < 0$, for all $u \in [0, 1/(4\lambda_1)]$, hence, there exists no solution for $f_m(u) = K - 2$ on $u \in (0, 1/(4\lambda_1)]$ (i.e. $s \in [4\lambda_1, \infty)$), a contradiction. It remains to show that indeed $f'_m(u) < 0$ on $(u, 1/(4\lambda_1)]$. To that end, consider

$$f'_m(u) = - \sum_{i \neq m} \frac{2\lambda_i}{\sqrt{1 - 4\lambda_i u}} + \frac{2\lambda_m}{\sqrt{1 - 4\lambda_m u}}, \quad 0 \leq u < 1/(4\lambda_1). \quad (45)$$

Under $K > 2$, for $f'_m(u) < 0$ to hold, it suffices that

$$\frac{\lambda_1}{\sqrt{1 - 4\lambda_1 u}} \geq \frac{\lambda_m}{\sqrt{1 - 4\lambda_m u}}. \quad (46)$$

The function $f(\cdot)$ defined as, $a > 0$,

$$f(x) = \frac{x}{\sqrt{1 - ax}}, \quad x \in [0, 1/a),$$

is increasing with x . Indeed, $f'(x) = (1 - ax/2)/(1 - ax)^{3/2} > 0$. Hence, (46) is true whenever $\lambda_1 \geq \lambda_m$, which is true by assumption. This shows the asserted monotonicity of $f'_m(\cdot)$.

To prove that (D) is necessary for the existence of rest points, suppose there exists a rest point x and (D) is not true. Assume first that x is given by (10). Plugging (10) into $s = \sum_{i=1}^K x_i$ and making the change of variable $u = 1/s$, we obtain $f_1(u) = K - 2$, where $f_1(\cdot)$ is defined by (44).

Now, note that $f_1(0) = K - 2$. Also, elementary calculus shows that the function $g : x \rightarrow x/\sqrt{1-x}$ is such that $g(x+y) \geq g(x) + g(y)$ for all $x, y \geq 0$, $x+y \leq 1$, the inequality being strict unless either x or y equals zero. Since by assumption $\lambda_1 \geq \sum_{i=2}^K \lambda_i$, this implies that $f_1'(u) > 0$, for $u \in [0, 1/(4\lambda_1)]$, except in the degenerate case where $\lambda_1 = \lambda_2$, and $\lambda_i = 0$ for $i > 2$. As the latter is ruled out, there exists no $u \in (0, 1/(4\lambda_1)]$ such that $f_1(u) = K - 2$, which contradicts the assumed existence of x .

Assume now that x exists and is given by (8). By the variable change $u = 1/s$, the relation $\sum_{i=1}^K x_i = s$ can again be rewritten as $\phi(u) = K - 2$, where $\phi(u) = \sum_{i=1}^K \sqrt{1 - 4\lambda_i u}$. As ϕ is decreasing, necessarily one must have that $\phi(1/4\lambda_1) \leq K - 2$. If the inequality is strict, then (C) holds, and a fortiori, (D) holds. If we have equality here, then (D) holds, except in the degenerate case $\lambda_1 = \lambda_2$, $\lambda_i = 0$ otherwise.

Assume that (C) holds. That is to say, $\phi(1/4\lambda_1) < K - 2$. Thus there exists a fixed point satisfying (8). Moreover, it is unique as ϕ is strictly decreasing.

Assume now that (C) fails, and (D) holds. Let us establish existence of a unique rest point given by (10). We again consider $f_1(u) = K - 2$ on $u \in [0, 1/(4\lambda_1)]$ where $f_1(\cdot)$ is defined by (44). Note that $f_1(0) = K - 2$ and $f_1(1/(4\lambda_1)) = \sum_{i=2}^K \sqrt{1 - \lambda_i/\lambda_1} \geq K - 2$. Moreover, $f_1'(0) = 2(\lambda_1 - \sum_{i=2}^K \lambda_i) < 0$ and $\lim_{u \uparrow 1/(4\lambda_1)} f_1'(u) = +\infty$. Hence, there must exist a $u \in (0, 1/(4\lambda_1)]$ such that $f_1(u) = K - 2$, which establishes the existence of the rest point satisfying (10). Moreover, we can show that there exists a unique $u \in [0, 1/(4\lambda_1)]$ such that $f_1'(u) = 0$. This implies uniqueness of such a rest point.

Finally, in order to conclude the proof, we have to establish that rest points of type (10) do not exist when condition (C) holds. This is easily done by examining the properties of f_1 and f_1' on the interval $[0, 1/4\lambda_1]$.

B. Proof of Theorem 2

We first note that, as $(d/dt)x_i(t) \geq \lambda_i - x_i(t)$, after some time t_0 , $x_i(t) \geq \lambda_i/2$. In order to establish boundedness (or lack of it), we introduce an auxiliary 2-dimensional system as follows. Fix $m \in \{1, 2, \dots, K\}$ and define $u = x_m$ and $v = \sum_{i \neq m} x_i$. We also use the notation $a := \lambda_m$, $b := \sum_{i \neq m} \lambda_i$ and $z := \sum_{i \neq m} x_i^2$. From (7), we have

$$\frac{d}{dt}u(t) = a - \frac{u(t)v(t)}{u(t) + v(t)} \quad (47)$$

$$\frac{d}{dt}v(t) = b - \frac{u(t)v(t) - v(t)^2 + z(t)}{u(t) + v(t)}. \quad (48)$$

Let us first establish that $x_1(t)$ goes to infinity when (D) does not hold, by taking $m = 1$. Note that

$$\frac{d}{dt}(u(t) - v(t)) = a - b + \frac{v(t)^2 - z(t)}{u(t) + v(t)}.$$

Since (D) is violated, $a - b \geq 0$. If the inequality is strict, as the right-hand side is larger than $a - b$, we see at once that $u(t) - v(t)$ increases linearly to infinity, and a fortiori $x_1(t)$ also increases linearly to infinity. If $a = b$, note that $v(t)^2 - z(t) =$

$2 \sum_{1 < i < j \leq K} x_i x_j \geq \lambda_2 \lambda_3 / 2$, because we placed ourselves after time t_0 , so that $x_i \geq \lambda_i/2$ for all i . We thus have

$$\frac{d}{dt}(u(t) - v(t)) \geq \frac{\lambda_2 \lambda_3 / 2}{u(t) + v(t)}.$$

Integrating, this yields

$$u(t) + v(t) \geq u(t) - v(t) \geq u(t_0) - v(t_0) + \frac{\lambda_2 \lambda_3}{2} \int_{t_0}^t \frac{d\tau}{u(\tau) + v(\tau)}. \quad (49)$$

Provided $u(t_0) \geq v(t_0)$, this implies that $u(t) + v(t)$ is at least of order \sqrt{t} . Since we also have that

$$u(t) \geq v(t) + u(t_0) - v(t_0),$$

necessarily $u(t)$ must also be at least of order \sqrt{t} . However, we can always choose t_0 such that $u(t_0) - v(t_0) \geq 0$, for otherwise, it would follow from (47) that $u(t)$ is bounded, and in turn from the previous display that $v(t)$ is bounded. However, this cannot hold in view of (49).

Assume now that (D) holds. By the same argument as above, we may place ourselves after some time t_0 such that for all $t \geq t_0$, and each $i \neq m$, $x_i(t) \geq \lambda_i - \epsilon$, for any given $\epsilon > 0$. In particular, for a suitably chosen $\epsilon > 0$, this ensures that for all $t \geq t_0$, $v(t) := \sum_{i \neq m} x_i(t) \geq b - (K - 1)\epsilon \geq a + \epsilon$. This is indeed feasible as, by assumption (D), $b > a$. It thus holds that

$$\frac{d}{dt}u(t) \leq a - \frac{(a + \epsilon)u(t)}{a + \epsilon + u(t)} = \frac{a(a + \epsilon) - \epsilon u(t)}{u(t) + a + \epsilon}.$$

This guarantees that $\limsup_{t \rightarrow \infty} u(t) \leq \epsilon^{-1}a(a + \epsilon)$. This argument holds for all $m \in \{1, 2, \dots, K\}$, and thus uniform boundedness of the dynamical system under consideration holds as claimed.

C. Proof of Lemma 2

The result follows by induction. Base step. $m = 1$. We have

$$\frac{d}{dt}y_1(t) = \lambda_1 - \left(1 - \frac{y_1(t)}{s_{K-1}(t)}\right) y_1(t).$$

Note that

$$\frac{y_1(t)}{s_{K-1}(t)} = \frac{y_1(t)}{\sum_{k=1}^{K-1} \binom{K}{k} y_k(t)} < \frac{y_1(t)}{\binom{K}{1} y_1(t)} = \frac{1}{K}.$$

From the last two relations, $(d/dt)y_1(t) < \lambda_1 - (K - 1)/K y_1(t)$. It follows that (17) for $m = 1$ is a super-solution for $y_1(t)$.

Induction step. Suppose (17) is a super-solution for the ODEs (16) for $m = 1, 2, \dots, n < K - 1$. We have

$$\begin{aligned} \frac{s_n(t) + y_{n+1}(t)}{s_{K-1}(t)} &= \frac{\sum_{k=1}^{n+1} \binom{n+1}{k} y_k(t)}{\sum_{k=1}^{K-1} \binom{K}{k} y_k(t)} \\ &\leq \frac{\sum_{k=1}^{n+1} \binom{n+1}{k} y_k(t)}{\sum_{k=1}^{n+1} \binom{K}{k} y_k(t)} < \max_{k=1}^{n+1} \frac{\binom{n+1}{k}}{\binom{K}{k}} = \frac{n+1}{K}. \end{aligned}$$

Combine this with (16) to write

$$\frac{d}{dt}y_{n+1}(t) < \lambda_{n+1} + \frac{n+1}{K - (n+1)} y_n^+(t) - \frac{K - (n+1)}{K} y_{n+1}(t).$$

This, together with $y(0)^+ = y(0)$, shows that $y_{n+1}(t) < y_{n+1}^+(t)$, for all $t > 0$. The proof is completed.

D. Proof of Theorem 5

The result is an easy consequence of the following lemma and the uniform convergence result of Kurtz.

Lemma 5. *The ODE (23) admits the limit point*

$$x_i(+\infty) = \begin{cases} x_1 \prod_{i=2}^K \left(1 - \frac{x_i}{x_1}\right) & \text{if } i = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (50)$$

The lemma follows from two auxiliary results that we display first. We introduce the transform:

$$y_i(t) := \exp\left(-\int_0^t x_i(u) du\right), \quad i = 1, 2, \dots, K.$$

The next lemma shows that a particular ‘‘constant of motion’’ is enjoyed by the ODE (23).

Lemma 6. *For the ODE (23), for all $i, j \in \{1, 2, \dots, K\}$,*

$$x_i(0)y_j(t) - x_j(0)y_i(t) = x_i(0) - x_j(0).$$

Proof: From (23), we have that

$$x_i(t) = x_i(0) \exp\left(-\int_0^t (s(u) - x_i(u)) du\right), \quad i = 1, 2, \dots, K.$$

Hence, for any $i, j \in \{1, 2, \dots, K\}$,

$$x_i(t) - x_j(t) = \left(\frac{x_i(0)}{y_i(t)} - \frac{x_j(0)}{y_j(t)}\right) \exp\left(-\int_0^t s(u) du\right). \quad (51)$$

On the other hand, by subtracting the right- and left-hand sides of the i -th and the j -th ODE in (23), we have that

$$\frac{d}{dt}(x_i(t) - x_j(t)) = -(x_i(t) - x_j(t))(s(t) - x_i(t) - x_j(t)).$$

Integrating the last ODE, we obtain

$$x_i(t) - x_j(t) = (x_i(0) - x_j(0)) \frac{1}{y_i(t)y_j(t)} e^{-\int_0^t s(u) du}. \quad (52)$$

By equating the right-hand sides of (51) and (52), the asserted relation follows.

We next give a lemma that characterizes a solution of the ODE (23). We use the notation $q_i := (x_1(0) - x_i(0))/x_1(0)$ and denote the geometric mean $\rho = \left(\prod_{i=2}^K q_i\right)^{\frac{1}{K-1}}$.

Lemma 7. *Let $x(t) = (x_1(t), x_2(t), \dots, x_K(t))$ be the solution of the initial value problem (23). Let the function $y_1(\cdot)$ be given by $y_1(0) = 1$ and*

$$\frac{d}{dt}y_1(t) = -x_1(0)y_1(t) \prod_{j=2}^K (q_j + (1 - q_j)y_1(t)). \quad (53)$$

Then, for all $i = 1, 2, \dots, K$,

$$x_i(t) = x_i(0) \prod_{j \neq i} (q_j + (1 - q_j)y_1(t)), \quad t \geq 0. \quad (54)$$

Proof: From (51), we have

$$x_i(t) = x_i(0) \prod_{j \neq i} y_j(t), \quad i = 1, 2, \dots, K. \quad (55)$$

Now, in the last display, substitute $y_j(t)$ by using Lemma 6 with $i = 1$ therein. This yields the asserted expression for $x_i(t)$. It remains only to show that y_1 is defined as asserted. By definition of y_i , we have that $(d/dt)y_i(t) = -x_i(t)y_i(t)$, $i = 1, 2, \dots, K$. Combining the last relation with (55) and Lemma 6, we obtain (53).

Proof: [Theorem 5] We consider two cases:

Case 1: $x_1(0) = x_2(0)$. We show that $x_i(+\infty) = 0$, for all $i = 1, 2, \dots, K$, as asserted in the theorem. By symmetry of the right-hand sides in (23), $x_i(t) > x_j(t)$ (resp. $=$), for all $t \geq 0$, whenever $x_i(0) > x_j(0)$ (resp. $x_i(0) = x_j(0)$). To contradict, suppose that $x_1(+\infty) = c > 0$. We now reuse the representation of the solution given in (55) to write

$$x_1(t) = x_1(0) \prod_{j=3}^K y_j(t)y_1(t).$$

Taking the limit $t \rightarrow \infty$ on both sides of the last relation, we obtain that $c = x_1(0) \prod_{j=3}^K y_j(+\infty) \cdot 0 = 0$, a contradiction.

Case 2: $x_1(0) > x_2(0)$. We first show that $\lim_{t \rightarrow \infty} y_1(t) = 0$. From (53), it follows

$$\begin{aligned} \frac{d}{dt}y_1(t) &< -x_1(0)y_1(t)(q_2 + (1 - q_2)y_1(t))^{K-1} \\ &\leq -x_1(0)y_1(t)q_2^{K-1}. \end{aligned}$$

Now, let $y_1^+(0) = y_1(0)$ and $(d/dt)y_1^+(t) = -x_1(0)y_1^+(t)q_2^{K-1}$. We have that $y_1(t) \leq y_1^+(t)$, for all $t \geq 0$, where $y_1^+(t) = e^{-x_1(0)q_2^{K-1}t}$. From our hypothesis $x_1(0) > x_2(0)$ we have that $q_2 > 0$ and thus $\lim_{t \rightarrow \infty} y_1^+(t) = 0$, which shows that we must have $y_1(+\infty) = 0$. The theorem now follows by Lemma 7 by noting from (54) that $\lim_{t \rightarrow \infty} x_i(t) = 0$ for $i > 1$, and $\lim_{t \rightarrow \infty} x_1(t) = x_1(0)\rho^{K-1}$.

E. Proof of Lemma 3

The construction yielding the result amounts to build the Markov processes from the same driving Poisson processes. More precisely, take two independent Poisson processes on the line, P_1 and P_2 , with intensity λ suitably large (larger than $X_0^2 + Y'^2$ is enough). Each point of P_1 may trigger a type 1-transition, namely one which involves departure of a type 1-individual: if the rate of the transition is r , then a random variable U , uniformly distributed on $[0, 1]$ is generated, and the transition is enabled if $U < r/\lambda$. The trick then is to use the same uniform random variables for the two Markov processes. It is easily seen that type 1-jumps will occur at the same instants in both processes until either Y or Y' hits zero, but Y' should hit zero first if the desired comparison holds. On the other hand, type-2 transitions occur in the original system at rate

$$Y + X_1 - \frac{\sum_{k=2}^K X_k^2}{Y},$$

where $Y = \sum_{j=2}^K X_j$. This rate is no larger than $Y(K - 2)/(K - 1) + X_1$. Thus, when $Y' = Y$, type-2 transitions are

more likely in the auxiliary system.

F. Proof of Lemma 4

The process X' may be constructed on $[0, t]$ as follows. First construct X'_1 . This may be done by associating i.i.d., exponentially distributed lifetimes T_i to each of the X_1 initial individuals, the parameter of the exponential distribution being 1, and by letting

$$X'_1(s) = \sum_{i=1}^{X_1} \mathbf{1}_{T_i > s}.$$

That $X'_1(t)$ has the desired binomial distribution follows directly.

Then, associate i.i.d., exponentially distributed lifetimes with parameter r to each of the initial Y individuals, namely S_i . If we could ignore the term X'_1 in the transition rate for reductions to Y' , we would have a similar representation for $Y'(t)$ as we have for $X'_1(t)$. Let V be exactly what we would have in that case:

$$V = \sum_{i=1}^Y \mathbf{1}_{S_i > t}.$$

Finally, in order to account for the extra death rate X'_1 in Y' , proceed as follows. Conditionally on X'_1 , and independently of the S_i 's, generate a Poisson process P with time varying intensity $X'_1(s)$, whose points are t_1, \dots, t_m . At each time t_i , $i = 1, 2, \dots, m$, one of the individuals constituting Y' is removed. Now, the random variable W accounts for these removals, for which the corresponding lifetime would have exceeded t . Conditionally on X'_1 and the Poisson process P , the distribution of W is specified as:

$$\mathbf{E}[z^W | P, X'_1] = \prod_{t_i \in P} \left(1 - e^{-r(t-t_i)} + e^{-r(t-t_i)} z\right).$$

Indeed, the probability that the lifetime S_i exceeds t is $e^{-r(t-t_i)}$. Take then expectations with respect to P . Using the formula for Laplace transforms of Poisson processes, one gets

$$\mathbf{E}[z^W | X'_1] = \exp\left(- (1-z) \int_0^t e^{-r(t-s)} X'_1(s) ds\right).$$

This is readily identified as the generating function of the sum of N independent random variables with a conditionally Poisson distribution, with random parameters:

$$\Lambda_i = \int_0^{t \wedge T_i} e^{-r(t-s)} ds. \quad (56)$$

Hence, the expectation of W reads:

$$\mathbf{E}(W) = N \mathbf{E}(\Lambda).$$

One readily sees that

$$\begin{aligned} \mathbf{E}(\Lambda_i) &= e^{-t} \int_0^t e^{-rs} ds + \int_0^u e^{-u} du \int_0^u e^{-r(t-s)} ds \\ &= \frac{e^{-rt} - e^{-t}}{1-r}, \end{aligned}$$

which yields the expression for $\mathbf{E}(W)$ in (33).

The large deviations bound (34) is a standard result for binomial random variables; see e.g. Corollary A.1.14, p. 268 in [4] (in fact, it holds with a factor 2 rather than 4 in front of the exponential).

In order to derive a similar result for W , we proceed in two steps. We first obtain a large deviations bound for the sum $\sum_{i=1}^N \Lambda_i$. This is readily done by noting that the r.v. Λ_i lies in the interval $[0, (1 - e^{-rt})/r]$. It is an easy result that any probability distribution restricted to this interval and with given mean is stochastically dominated, in the sense of convex stochastic ordering (see e.g. [3]), by that which puts mass only at the two end points of this interval and has the same mean; see e.g. [3]. Thus, provided the interval length $(1 - e^{-rt})/r$ is uniformly bounded, we may obtain a large deviations result for $\sum_{i=1}^N \Lambda_i$ from that available for binomial distributions. The interval length is indeed uniformly bounded for $K \geq 3$ by $1/r = (K-1)/(K-2) \leq 2$.

The large deviations bound for W then follows by writing:

$$\begin{aligned} \mathbf{P}(|W - \mathbf{E}(W)| \geq \epsilon \mathbf{E}(W)) &\leq \\ &\mathbf{P}\left(\left|\sum_{i=1}^N \Lambda_i - \mathbf{E}(W)\right| \geq \epsilon/2 \mathbf{E}(W)\right) \\ &+ \mathbf{P}(\text{Poi}((1 + \epsilon/2)\mathbf{E}(W)) \geq (1 + \epsilon)\mathbf{E}(W)) \\ &+ \mathbf{P}(\text{Poi}((1 - \epsilon/2)\mathbf{E}(W)) \leq (1 - \epsilon)\mathbf{E}(W)), \end{aligned}$$

where we have used the fact that W is conditionally Poisson, with parameter $\sum_{i=1}^N \Lambda_i$. As Poisson random variables also satisfy a large deviations bound of type (34), with a constant of 2 in front of the exponent, this yields a similar result for W , with a constant of 4 instead.

PLACE
PHOTO
HERE

Laurent Massoulié is a researcher with Thomson Research Paris, France. His research interests include management of overlay networks, peer-to-peer system design, congestion control, and probabilistic modelling and performance analysis of communication systems in general.

PLACE
PHOTO
HERE

Milan Vojnović is a researcher with systems and networks group at Microsoft Research Cambridge, United Kingdom. His research interests are in architecture and performance of networks and computer systems that include wireless & mobile systems, information dissemination systems, information systems, and congestion control.