

Stochastic Analysis of Self-Sustainability in Peer-Assisted VoDSystems

Original

Stochastic Analysis of Self-Sustainability in Peer-Assisted VoDSystems / Ciullo, Delia; Martina, Valentina; M., Garetto; Leonardi, Emilio; G., Torrisi. - ELETTRONICO. - (2012), pp. 1539-1547. (Intervento presentato al convegno IEEE INFOCOM 2012 tenutosi a Orlando nel April 2012) [10.1109/INFOCOM.2012.6195521].

Availability:

This version is available at: 11583/2499043 since:

Publisher:

IEEE / Institute of Electrical and Electronics Engineers Incorporated:445 Hoes Lane:Piscataway, NJ 08854:

Published

DOI:10.1109/INFOCOM.2012.6195521

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Stochastic Analysis of Self-Sustainability in Peer-Assisted VoD Systems

Delia Ciullo^{*}, Valentina Martina^{*}, Michele Garetto[†], Emilio Leonardi^{*}, Giovanni Luca Torrisi[‡]

^{*} Dipartimento di Elettronica, Politecnico di Torino, Torino, Italy

[†] Dipartimento di Informatica, Università di Torino, Torino, Italy

[‡] Istituto per le Applicazioni del Calcolo, CNR, Roma, Italy

Abstract—We consider a peer-assisted Video-on-demand system, in which video distribution is supported both by peers caching the whole video and by peers concurrently downloading it. We propose a stochastic fluid framework that allows to characterize the additional bandwidth requested from the servers to satisfy all users watching a given video. We obtain analytical upper bounds to the server bandwidth needed in the case in which users download the video content sequentially. We also present a methodology to obtain exact solutions for special cases of peer upload bandwidth distribution. Our bounds permit to tightly characterize the performance of peer-assisted VoD systems as the number of users increases, for both sequential and non-sequential delivery schemes. In particular, we rigorously prove that the simple sequential scheme is asymptotically optimal both in the bandwidth surplus and in the bandwidth deficit mode, and that peer-assisted systems become totally self-sustaining in the surplus mode as the number of users grows large.

I. INTRODUCTION

Video traffic over the Internet is growing at an impressive rate, being responsible nowadays for the majority of data volume flowing through networks. According to Cisco forecasts [1], the combination of all forms of video (live streaming, video-on-demand and P2P file sharing) will exceed 90% of all global consumer Internet traffic by 2015, posing a tremendous challenge to both content providers and network operators.

Thanks to the proliferation of proxy servers and the adoption of efficient distributed caching solutions, traditional (client-server) Content Delivery Networks (CDNs) play a fundamental role in the distribution of bandwidth-hungry contents such as video, since they permit to “move” contents close to the users, thus reducing the impact on the transport network backbone, and improving the user-perceived performance (*e.g.*, reducing the latency). Increasing traffic volumes, however, require large investments to continuously upgrade the CDNs infrastructure. Observe, indeed, that the aggregate resources required at the servers (bandwidth/storage/processing), and the corresponding costs incurred by content providers, scale linearly with the user demand and the data volume.

For this reason, the P2P approach, originally adopted by file-sharing applications, has become very attractive also for video distribution, as a viable solution to offload the servers and dramatically reduce the bandwidth costs of video publishers [2], [3]. In peer-assisted solutions, there are still servers permanently storing the contents and supervising the video distribution, but peers contribute their upload bandwidth to the redistribution of downloaded pieces to other peers, making the additional bandwidth requested from the servers potentially very small and independent of the numbers of watching users.

In peer-assisted Video-on-Demand (VoD) systems, users browse a catalog of available videos and asynchronously

issue requests to watch a given content, which are ideally immediately satisfied by the system, with the optional support for VCR actions such as pause and jump forward/backward. It is important to distinguish VoD systems from live streaming, in which users join the distribution of a given TV channel at random points in time, but peers connected to the same channel watch the content almost synchronously.

Peer-assisted applications for VoD and live streaming are typically inspired by Bit-Torrent: the content is divided into chunks which are disseminated among peers (or retrieved from servers) in a fully distributed (swarming) fashion based on the exchange of bitmaps. However, differently from traditional file-sharing, chunk and peer selection strategies for peer-assisted video distribution must account for the fact that users watch while downloading. In particular, to avoid service interruptions/degradations: i) a minimum average download rate equal to the video playback rate must be guaranteed; ii) an “almost in order” delivery of chunks is required.

Several peer-assisted systems have already been deployed, attracting millions of users, such as PPLive, GridCast, PP-Stream, TVU, SopCast [4]. However, despite the wide popularity gained by existing applications, several fundamental questions remain unanswered about the design of video streaming systems and the potential benefits of the peer-assisted approach.

In our work we consider peer-assisted VoD systems, focusing our analysis on a given video whose distribution is supported both by concurrently downloading users and by peers caching a copy of it. Our main contribution is a stochastic fluid framework that allows to derive general upper bounds to the bandwidth requested from the servers in the case of sequential chunk delivery, thus permitting to analytically bound the performance achievable by any chunk distribution scheme. In particular, our bounds allow to tightly characterize the behavior of peer-assisted VoD systems as the number of users increases, proving in a rigorous way the self-sustainability (*i.e.*, the desirable operating point at which no additional server bandwidth is needed) of systems operating in the bandwidth surplus mode.

We also show that the simple sequential delivery scheme is asymptotically optimal both in the surplus and in the deficit mode, under weak assumptions on the user behavior. For the sequential scheme, we also propose a methodology to derive exact estimate of the bandwidth requested from the servers under the assumption that the peer upload bandwidth is exponentially (or more generally phase-type) distributed.

We emphasize that our analysis is orthogonal to two ongoing streams of theoretical research: the one targeting opti-

mal replication strategies and push/pull schemes for content (re)placement [5], [6], [7], and the one dealing with practical issues related to chunk scheduling [8] and peer selection [9]. This will become clear later on in the paper.

The paper is organized as follows. In Sect. II we describe the system model; in Sect. III we first introduce the stochastic fluid framework used by our theoretical analysis; then we obtain general bounds in Sect. III-B and exact solutions for particular cases in Sect. III-C; in Sect. IV we give an overview of the related literature. We conclude the paper in Sect. V.

II. MODEL

A. System assumptions

We model a peer-assisted VoD system where users¹ run applications that allow them to browse an online catalog of videos. When a user selects a video, we assume that the request is immediately satisfied and the selected video can be watched uninterruptedly till the end, *i.e.*, the system is able to steadily provide to the user a data flow greater than or equal to the video playback rate. Users contribute their upload bandwidth to the video distribution, thus they can retrieve part of the video (or even the entire video) from other peers, saving servers resources.

We focus on a given video of size L bytes. We assume that the video is downloaded by each user at constant rate d bytes/s, equal to the playback rate. Let $\tau_d = L/d$ be the time needed to download the whole video. In general, the download rate of a peer could be adapted over time, and even depend on some peer's characteristics (such as its upload bandwidth). By assuming a constant download rate d at each user we greatly simplify the analysis, while obtaining a conservative prediction with respect to the case in which the download rate can be adapted over time maintaining an average value equal to d .

The amount of upload bandwidth with which peers contribute to the redistribution of the video that they are downloading, instead, may or may not be under the control of the system. In our analysis, we assume that the upload bandwidth available at a peer is a random variable with a given distribution. This way, we encompass both the realistic case of users with heterogeneous Internet connections (*i.e.*, ADSL, fiber, LAN) and cross-traffic fluctuations, and the case in which the peer upload bandwidth allocated to the given video is tuned by the system (such as in universal streaming architectures). More specifically, the amount of upload bandwidth with which users contribute at a given time to the redistribution of the considered video is modeled as a random variable U with cumulative distribution function $F_U(w)$ and mean \bar{U} . The random variables denoting the instantaneous upload bandwidths of the users are assumed to be i.i.d. (independent and identically distributed).

B. Peers dynamics

We assume that the arrival process of requests for the considered video follows a Poisson process of intensity λ . Assuming that at a given time the arrival process is Poisson is reasonable, since users behave independently of each other, and their requests are immediately satisfied. Although in this paper we consider a constant request rate, our analysis can

be easily extended to the case of an inhomogeneous Poisson process with time-varying intensity $\lambda(t)$. The impact of non-stationary traffic conditions on the system performance is studied in our companion paper [10].

As soon as users issue their request to watch the considered video, they start downloading it and assisting other peers. We define as *activity period* the duration of the interval during which a peer contributes its upload bandwidth to the distribution of the requested video. Activity periods of the users are highly heterogeneous, as observed in several measurement studies [3]: some users stop watching the video after a very short time since the beginning, because they realize they are no longer interested in it; most users who decide to watch the video shut down the computer/Internet-TV towards the end of it; some of them keep the application running for prolonged time after the end of the video; those running set-top-boxes can be considered to be always active and serving other peers (while the set-top-box keeps a copy of the considered video, contributing to its distribution). We adopt a fairly general model of peer churn (or set-top-box behavior), assuming that the activity period of each user is described by an arbitrary random variable T with finite mean \bar{T} and cumulative distribution function $F_T(x)$. The activity periods of the users are assumed to be i.i.d.

It follows from our assumptions that the number of active users N at a certain time instant is distributed as the number of customers in an M/G/ ∞ queue, hence it follows a Poisson distribution with mean $\bar{N} = \lambda\bar{T}$. In our analysis we need to distinguish two classes of active users: those who are still *downloading* the video, and those who have completed the download (referred to as *seeds* in the following). The number of *downloading* peers at a given time instant, N_d , follows a Poisson distribution of mean $\bar{N}_d = \lambda\bar{T}_d$, where $\bar{T}_d = \int_0^{\tau_d} (1 - F_T(x)) dx$ is the average time spent downloading the video. Observe that in general \bar{T}_d can be shorter than τ_d due to peer churn (premature abandons).

Then, standard properties of Poisson processes allow to say that the number of *seeds* at a given time, N_{seed} , follows a Poisson distribution of mean $\bar{N}_{seed} = \bar{N} - \bar{N}_d = \lambda(\bar{T} - \bar{T}_d)$.

We define as system load the quantity

$$\gamma \triangleq \frac{\bar{N}_d d}{\bar{U} \cdot \bar{N}} \quad (1)$$

which is the ratio between the average aggregate rate requested by downloading peers, and the average aggregate upload rate provided by all active users. Borrowing the terminology adopted in previous work [2], [11] we say² that the system operates in *deficit* mode if $\gamma > 1$, and in *surplus* mode if $\gamma < 1$.

C. Performance metrics

A fundamental goal of a VoD system is to minimize the bandwidth requested from the servers. To save server bandwidth, the system tries to exploit as much as possible the upload capacity of the peers, under the strict constraints of video distribution (*i.e.*, delay, minimum rate). Let S be the random variable denoting the additional bandwidth that the servers must supply at a given time to satisfy all active downloads of the considered video. Let $F_S(w)$ be the cumulative distribution of S . At last, we denote by \bar{S} the mean of S . Since

¹In this paper we use the terms peer and user interchangeably.

²In this paper we do not consider the special case $\gamma = 1$.

TABLE I
NOTATION

Symbol	Definition
d	user download rate
\bar{U}	average user upload bandwidth
λ	arrival rate of new requests for the video
\bar{T}_d	average time spent downloading the video (s)
\bar{T}	average user activity period (s)
\bar{N}	average number of users
\bar{N}_d	average number of <i>downloading</i> users
\bar{N}_{seed}	average number of <i>seeds</i>
\bar{S}_d	average bandwidth requested by downloading users
\bar{S}_{seed}	average bandwidth offered by seeds
\bar{S}	average bandwidth requested from the servers
γ	system load

in practice there are multiple videos to be served concurrently by the system, statistical multiplexing arguments suggest that a good design goal is to minimize the mean value \bar{S} of the servers bandwidth required by a single video. Therefore, this will be the main metric that we will look at in our performance analysis.

We define a VoD system to be *self-sustaining* if $\lim_{\lambda \rightarrow \infty} \bar{S} = 0$. Notice that servers must necessarily provide some bandwidth when a new video is inserted in the system. However, once the video has been replicated at peers, it is possible that no additional server bandwidth is needed from the servers, since the swarm can sustain itself. In our analysis we neglect transient effect, and look at the stationary system behavior. Table I summarizes the notation of our model.

III. ANALYSIS

A. Preliminaries and universal lower bound

Our goal is to characterize the random variable S denoting the bandwidth requested from the servers at a given time. We denote by S_d the aggregate bandwidth requested by the downloading users, and by $S_{seed} = \sum_{i=1}^{N_{seed}} U_i$ the aggregate upload bandwidth offered by the seeds. The bandwidth requested from the servers is then given by the difference between the bandwidth requested by the downloading users and the bandwidth offered by seeds, provided that such difference is positive:

$$S \triangleq \max\{0, S_d - S_{seed}\}. \quad (2)$$

We observe that S_{seed} is simple to characterize, since N_{seed} is a random variable with a Poisson distribution. Then S_{seed} is a compound Poisson random variable whose moment generating function is $\mathbb{E}[e^{tS_{seed}}] = e^{\bar{N}_{seed}(\phi_U(t)-1)}$, where $\phi_U(t) \triangleq \mathbb{E}[e^{tU}]$ is the moment generating function of the peer upload bandwidth U , which is known. In particular, the average of S_{seed} is $\bar{S}_{seed} = \bar{N}_{seed}\bar{U}$. Notice also that S_{seed} is independent of S_d .

The difficulty then lies in characterizing the random variable S_d . To analyze S_d , we first condition it on the number of downloading users, defining

$$S_d(k) \triangleq (S_d | N_d = k) \quad (3)$$

Focusing our attention to $S_d(k)$, we observe that an easy lower bound to this quantity can be obtained assuming that the upload bandwidth of each downloading user can be always fully utilized by the system, irrespective of the arrival time of

the user into the system. We obtain

$$S_d(k) \geq \max\{d, k d - \sum_{i=1}^k U_i\} \geq k d - \sum_{i=1}^k U_i$$

and thus $\mathbb{E}[S_d(k)] \geq k(d - \bar{U})$. Deconditioning with respect to k we obtain $\mathbb{E}[S_d] \geq \bar{N}_d(d - \bar{U})$.

At last, we observe that by construction $\bar{S} = \mathbb{E}[\max\{0, S_d - S_{seed}\}] \geq \max\{0, \mathbb{E}[S_d - S_{seed}]\} \geq \max\{0, \bar{N}_d(d - \bar{U}) - \bar{N}_{seed}\bar{U}\} = \max\{0, d\bar{N}_d - \bar{U}\bar{N}\} \quad (4)$

which provides a universal lower bound to \bar{S} for any chunk distribution scheme. The above lower bound is trivially zero for $\gamma < 1$, whereas it is equal to $d\bar{N}_d - \bar{U}\bar{N}$ for $\gamma \geq 1$.

B. Upper bounds

An upper bound to the bandwidth requested from the servers can be obtained assuming that all peers must download the video chunks sequentially. We observe that many implemented applications inspired by BitTorrent allow also non-sequential chunk dissemination in a swarm-like fashion, although this is typically only enabled within a limited portion of the video to meet the hard delay constraints of individual chunks. Actually, an almost in-order download is the only choice when the download rate d is close to the video playback rate (and the start-up delay is small).

Besides being analytically tractable (as we will see), the sequential download is also simple to implement in a peer-assisted VoD system, as it does not require the complex chunk/peer selection mechanisms which are necessary in BitTorrent-like swarms. In any case, the main point is that the server bandwidth required under sequential download is an upper bound to the bandwidth required by a more general (non-sequential) download scheme.

Below we show how to obtain analytical upper bounds to \bar{S} , the average of S , in the case of sequential download, obtaining upper bounds valid for any distribution scheme.

We start looking at quantity $S_d(k)$ defined in (3). We observe that, if all peers download the video sequentially at common rate d , a peer can only redistribute video pieces to peers arrived later on in time.

When there is only one downloading user, we trivially have $S_d(1) = d$. If there are two downloaders, the first arrived makes its entire upload bandwidth available to the second, and we have

$$S_d(2) = d + \max\{0, d - U_1\} = \max\{d, 2d - U_1\}$$

where d represents the external bandwidth necessary to sustain the download of the first arrived peer and $\max\{0, d - U_1\}$ represents the bandwidth needed to sustain the download of the second arrived peer.

When there are three downloaders, the last arrived can exploit the upload bandwidth of the second plus the residual upload bandwidth of the first, *i.e.*, a total upload rate of $U_2 + \max\{U_1 - d, 0\}$. Summing up the download rates needed by the three peers, we obtain

$$\begin{aligned} S_d(3) &= d + \max\{0, d - U_1\} + \max\{0, d - U_2 - \max\{U_1 - d, 0\}\} \\ &= d + \max\{0, d - U_1\} + \max\{0, \min\{2d - U_1 - U_2, d - U_2\}\} \\ &= \max\{S_d(2), 3d - U_1 - U_2\} \end{aligned} \quad (5)$$

The last equation descends from the fact that if $d - U_2 < 2d - U_1 - U_2$ then $d - U_1 > 0$ and thus $\max\{0, d - U_1\} + \max\{0, \min\{2d - U_1 - U_2, d - U_2\}\} = \max\{0, d - U_1, 2d - U_1 - U_2\}$.

In general the k -th downloader (assuming downloaders to be numbered in order of arrival) can receive the content from every other downloader preceding it. However, if preceding peers are not able to fully support the download of the k -th downloader in addition to previous downloads, (i.e., if $\sum_{i=1}^{k-1} U_i - kd < 0$) the missing bandwidth must be provided either by servers or by seeds. We obtain the following recursive equation for $S_d(k)$:

$$S_d(k) = \begin{cases} d & k = 1 \\ \max\{S_d(k-1), kd - \sum_{j=1}^{k-1} U_j\} & k > 1 \end{cases} \quad (6)$$

If we iterate back up to $S_d(1)$ we can obtain an explicit expression for $S_d(k)$ in terms of the upload bandwidths of peers U_i , for $i < k$, and of the download rate d , as:

$$\begin{aligned} S_d(k) &= d + \max\left\{0, d - U_1, 2d - (U_1 + U_2), \right. \\ &\quad \left. 3d - (U_1 + U_2 + U_3), \dots, (k-1)d - \sum_{i=1}^{k-1} U_i\right\} \\ &= d + \max_{1 \leq j \leq k-1} \left\{0, \sum_{i=1}^j (d - U_i)\right\} \end{aligned} \quad (7)$$

We emphasize that (7) has already been obtained in [2], [11], however in previous work authors have resorted to Monte-Carlo approaches to evaluate it. To the best of our knowledge, we are the first to provide an analytical characterization of the solution to (7), and of the resulting server bandwidth S defined in (2), considering also the impact of the seeds.

To proceed, we define the auxiliary variable $Z_d(k)$, that will come in handy in the following:

$$Z_d(k) \triangleq \max_{1 \leq j \leq k} \left\{ \sum_{i=1}^j (d - U_i) \right\} \quad (8)$$

where $Z_d(k) = 0$ if $k = 0$. Then $S_d(k)$ can be expressed in terms of $Z_d(k-1)$ according to

$$S_d(k) = d + \max\{0, Z_d(k-1)\}. \quad (9)$$

Now, $Z_d(k)$ can be regarded as the maximum value (up to time k) reached by a unidimensional random walk with increments $X_i = d - U_i$. Thus we can exploit the existing literature on random walks, and especially their application to risk theory, to characterize the distribution of $Z_d(k)$.

For our purposes, we need the following classic result, known as the Lundberg's inequality (see for example [12]).

Lemma 1: (Lundberg inequality) Consider a sequence of i.i.d. variables $(X_i)_{i \geq 1}$, satisfying the following three properties: i) $\mathbb{E}[X_1] < 0$; ii) $\mathbb{P}(X_1 > 0) > 0$; iii) $\mathbb{E}[e^{tX_1}]$ is finite in a neighborhood of the origin. Define the r.v. $Q(k) \triangleq \sum_{i=1}^k X_i$, $k \geq 1$, $Q(0) \triangleq 0$. Then, denoted θ^* the strictly positive solution of $\mathbb{E}[e^{\theta^* X_1}] = 1$, which exists unique under i), ii), and iii), we have, for all $n \geq 1$:

$$\mathbb{P}\left(\max_{1 \leq k \leq n} Q(k) > w\right) \leq e^{-\theta^* w}, \quad \forall w \geq 0. \quad (10)$$

Remark: condition iii) requires X_1 to be light-tailed (i.e., to have a tail that decays at least exponentially fast).

For completeness, in our technical report [13] we report a proof of Lemma 1 based on a Martingale approach.

Lundberg inequality can be generalized and adapted to our context, to obtain an upper bound to $\mathbb{P}(S_d > w)$:

Theorem 1: Assume the following properties hold for U : i) $\bar{U} > 0$, ii) $\mathbb{E}[e^{tU}]$ is finite in a neighborhood of the origin, iii) $F_U(w) > 0$ for every $w > 0$. For $\epsilon \in [(\bar{U} - d)^+, \bar{U}]$, define $A := d - \bar{U} + \epsilon$ (note that $\max\{0, d - \bar{U}\} \leq A < d$). Let θ^* be the unique strictly positive solution of the equation $\mathbb{E}[e^{\theta(d-U-A)}] = 1$. For any $w \geq 0$, it holds

$$\mathbb{P}(S_d > w) \leq \begin{cases} \min\{C_2 e^{-\theta^*(w-d)}, C_3\} & w \geq d \\ C_1 & 0 \leq w < d, \end{cases}$$

where $C_1 \triangleq 1 - e^{-\bar{N}_d}$, $C_2 \triangleq e^{-\theta^* A} e^{-\bar{N}_d} (e^{\bar{N}_d e^{\theta^* A}} - \bar{N}_d e^{\theta^* A} - 1)$ and $C_3 \triangleq 1 - e^{-\bar{N}_d} - \bar{N}_d e^{-\bar{N}_d}$.

A detailed proof of Theorem 1 is reported in Appendix A. Observe that, when $d < \bar{U}$, we can obtain an upper bound on $\mathbb{P}(S_d > w)$ applying the Lundberg inequality to $\mathbb{P}(S_d(k) > w)$ for any k . Instead, when $d > \bar{U}$, since $\mathbb{E}[d - U_i] > 0$, we cannot apply Lundberg inequality directly to $\mathbb{P}(S_d(k) > w)$. Therefore we need to define an auxiliary sequence of random variables, tightly related to $S_d(k)$, on which Lundberg bound can be applied. Then we can derive a bound on $\mathbb{P}(S_d(k) > w)$. The approach of the auxiliary sequence of variables is generalized also to the case $d < \bar{U}$, to obtain a possibly tighter upper bound.

Exploiting the result in Theorem 1, we derive an upper bound to the average bandwidth \bar{S}_d requested by the downloading peers:

Corollary 1: The average bandwidth requested by downloading peers satisfies:

$$\bar{S}_d \leq \begin{cases} C_1 d + C_3(w^* - d) + C_3/\theta^* & \text{if } C_2 > C_3 \\ C_1 d + C_2/\theta^* & \text{o.w.} \end{cases} \quad (11)$$

where $w^* \triangleq \frac{1}{\theta^*} \log\left(\frac{C_2}{C_3}\right) + d$.

The proof of Corollary 1 can be found in our technical report [13].

Remark: note that (1) and (11) hold under an arbitrary choice of $\epsilon \in [(\bar{U} - d)^+, \bar{U}]$. The tightest bound is obtained by minimizing the expressions (1) and (11) with respect to ϵ .

From Corollary 1 it immediately follows that the average bandwidth requested by downloaders is finite even when $\lambda \rightarrow \infty$ (i.e., $\bar{N}_d \rightarrow \infty$), provided that $\bar{U} > d$. Indeed, by selecting $\epsilon = \bar{U} - d$ (and thus $A = 0$), we obtain $\bar{S}_d \leq d + \frac{1}{\theta^*}$, being θ^* the unique positive solution to $\mathbb{E}[e^{\theta(d-U)}] = 1$.

By taking into account also the impact of the seeds, we obtain an upper bound to the average bandwidth requested from the servers, according to (2):

Theorem 2: The following bound holds for the average bandwidth requested from the server:

$$\begin{aligned} \bar{S} \leq \min \left\{ C_1 F_{S_{\text{seed}}}(d) d + C_2 \frac{e^{\theta^* d}}{\theta^*} \mathbb{E}[e^{-\theta^* S_{\text{seed}}}], \right. \\ \left. C_1 F_{S_{\text{seed}}}(d) d + C_3 \left(\frac{2}{\theta^*} + w^* - \bar{S}_{\text{seed}} \right) \right. \\ \left. + \mathbb{E}[e^{\theta^* S_{\text{seed}}}] e^{-\theta^* w^*} / \theta^* \right\} \end{aligned} \quad (12)$$

where $w^* = \frac{1}{\theta^*} \log\left(\frac{C_2}{C_3}\right) + d$.

The proof of Theorem 2 is reported in our technical report [13]. We state now our main result that completely characterizes the system behavior as the number of users increases.

Theorem 3: Assume U not constant. Then, as $\lambda \rightarrow \infty$, the following asymptotic regimes hold for any chunk distribution scheme: For $\gamma < 1$ and, additionally $\bar{T} > \bar{T}_d$, the average bandwidth requested from the servers tends to zero, *i.e.*, $\lim_{\lambda \rightarrow \infty} \bar{S} = 0$. For $\gamma > 1$, the average bandwidth requested to the servers grows linearly with the number of users. In particular, $\lim_{\lambda \rightarrow \infty} \frac{\bar{S}}{(\bar{N}_d d - \bar{U} \bar{N})} = 1$.

The proof of Theorem 3 can be found in Appendix B. Notice that for $\gamma > 1$ the upper bound becomes asymptotically tight to the lower bound (4).

Theorem 3 suggests that, for very popular contents (large number of users concurrently watching the same video), a peer-assisted video distribution is *self-sustaining*, provided that the system is in surplus mode (*i.e.*, $\gamma < 1$) and users stay in the system, on average, for a time larger (by an arbitrarily small constant) than the time needed to download the whole video ($\bar{T} > \bar{T}_d$). This holds for any chunk distribution scheme, including the simple sequential scheme.

In the deficit mode, *i.e.*, $\gamma > 1$, the system is obviously not *self-sustaining*, since an additional bandwidth at least equal to the bandwidth deficit $(\bar{N}_d d - \bar{U} \bar{N})$ must be provided by servers. However, in this case there is (asymptotically) no gain in adopting a non-sequential chunk delivery scheme with respect to the simple sequential download. Indeed, as the number of users grows large, the system in which users download the content sequentially performs as well as an ideal (unfeasible for VoD applications) system in which the content can be downloaded in arbitrary order.

C. Exact solutions

In this section we derive a methodology to obtain an exact solution of (7) when the bandwidth distribution is exponentially or phase-type distributed.

The first step consists in deriving an integral equation satisfied by the cumulative distribution function of the quantity $Z_d(k)$ defined in (8). Observe that $Z_d(k)$ can be written as:

$$Z_d(k) = \max \left\{ d - U_1, \max_{2 \leq j \leq k} \left\{ \sum_{i=2}^j (d - U_i) \right\} + d - U_1 \right\}.$$

Now since the U_i are i.i.d., we can permute the indices i of $(U_i)_{i \geq 1}$, obtaining a new random variable $\hat{Z}_d(k)$ defined as

$$\hat{Z}_d(k) \triangleq \max \left\{ d - U_k, \max_{1 \leq j \leq k-1} \left\{ \sum_{i=2}^j (d - U_i) \right\} + d - U_k \right\}$$

which has the same distribution of $Z_d(k)$. Note that $\hat{Z}_d(k)$ can be written as:

$$\hat{Z}_d(k) = \max \left\{ d - U_k, \hat{Z}_d(k-1) + d - U_k \right\}. \quad (13)$$

Denoting by $F_Z(w | k)$ the cumulative distribution function of $Z_d(k)$ (and thus of $\hat{Z}_d(k)$), we have:

$$\begin{aligned} F_Z(w | k) &= \mathbb{P}(\hat{Z}_d(k) \leq w) = \\ &= \mathbb{P} \left(\max \left\{ d - U_k, \hat{Z}_d(k-1) + d - U_k \right\} \leq w \right) \end{aligned} \quad (14)$$

We now condition on the value assumed by $X_k \triangleq d - U_k$:

$$\begin{aligned} F_Z(w | k) &= \\ &= \int_{-\infty}^{\infty} \mathbb{P} \left(\max \left\{ X_k, \hat{Z}_d(k-1) + X_k \right\} \leq w | X_k = \alpha \right) dF_{X_k}(\alpha) \\ &= \int_{-\infty}^w \mathbb{P} \left(\hat{Z}_d(k-1) + \alpha \leq w \right) dF_{X_k}(\alpha) \\ &= \int_{-\infty}^w F_Z(w - \alpha | k-1) dF_{X_k}(\alpha). \end{aligned} \quad (15)$$

Observing that by construction $F_Z(\alpha | 1) = F_{X_k}(\alpha)$, from (15) we get:

$$F_Z(w | k) = \int_{-\infty}^w F_Z(w - \alpha | k-1) dF_Z(\alpha | 1) \quad (16)$$

Explicit solutions of the above functional equation can be given when the peer upload bandwidth is phase-type distributed. In particular, in the case of peer upload bandwidth exponentially distributed, we obtain (see Appendix C):

$$\begin{aligned} F_Z(w | k) &= F_Z(d | k) e^{\frac{w-d}{\bar{U}}} \mathbb{I}_{w < 0} + \mathbb{I}_{w > kd} + \\ &+ e^{\frac{w-d}{\bar{U}}} \sum_{i=0}^k (-1)^i F_Z(d | k-i) \frac{(w-id)^i}{\bar{U}^i i!} e^{-id/\bar{U}} \mathbb{I}_{id \leq w \leq kd}, \end{aligned} \quad (17)$$

where \mathbb{I} is the indicator function. In (17) the constants $F_Z(d | k)$ can be obtained imposing the condition $F_Z(kd | k) = 1$ for all k , as shown in Appendix C.

From $F_Z(d | k)$ we immediately obtain $\mathbb{P}(S_d(k) > w)$:

$$\mathbb{P}(S_d(k) > w) = \begin{cases} 1 - F_Z(w - d | k-1) & \text{if } w \geq d \\ 1 & \text{if } w < d \end{cases} \quad (18)$$

Finally, we can derive the average server bandwidth \bar{S}_d requested by the downloading peers, and the average server bandwidth \bar{S} requested from the servers.

In the case $\bar{U} > d$, since the sequence of increasing random variables $Z_d(k)$ converges w.p.1 to a finite random variable $Z_d(\infty)$ (as direct consequence of Lemma 1), we can find the distribution of $Z_d(\infty)$ from the stationary version of (16). We state this result in the following theorem.

Theorem 4: Under the condition $\bar{U} > d$, the cumulative function of $Z_d(\infty)$ satisfies the stationary version of (16), *i.e.*,

$$F_Z(w | \infty) = \int_{-\infty}^w F_Z(w - \alpha | \infty) dF_X(\alpha) \quad (19)$$

When the U_k are exponentially distributed the solution of (19) can be obtained following the approach described in Appendix C, obtaining:

$$\begin{aligned} F_Z(w | \infty) &= F_Z(d | \infty) e^{\frac{w-d}{\bar{U}}} \mathbb{I}_{w < 0} + \\ &+ F_Z(d | \infty) e^{\frac{w-d}{\bar{U}}} \sum_{i=0}^{\infty} (-1)^i \frac{(w-id)^i}{\bar{U}^i i!} e^{-id/\bar{U}} \mathbb{I}_{w > id} \end{aligned}$$

where

$$F_Z(d | \infty) = \lim_{k \rightarrow \infty} \frac{1}{\sum_{i=0}^k (-1)^i \frac{[(k-i)d]^i}{i!} e^{(k-i)d/\bar{U}}}$$

Notice that $Z_d(\infty)$ provides a tight bound to the distribution of the server bandwidth requested by a large number of downloading users, when the system operates in the surplus mode.

We emphasize that the approach described in Appendix C can be generalized in a rather straightforward way to obtain the exact solution of (16) and (19) under any phase-type distribution of peers upload bandwidth. On this regard, recall that any distribution whose moment generating function is finite in a neighborhood of the origin (*i.e.*, it is light-tailed) can be approximated by a phase-type distribution with an arbitrary degree of accuracy (see [12, Ch.4]). Thus the methodology presented in this section can be applied to derive analytical approximations of the bandwidth requested from the servers in the case in which the peer upload bandwidth is arbitrarily distributed.

As a concluding remark, we wish to emphasize that upper bounds obtained in Section III-B and exact solutions presented in this section are complementary tools for the analysis of peer-assisted VoD systems properties. Indeed, the upper bounds presented in Section III-B provide very general and easy-to-handle expressions from which we can derive qualitative/asymptotic properties of the system. The methodology described in this section provides more accurate estimates (which are exact for phase-type distributions) of the bandwidth requested from the servers, but are computationally more expensive, especially for large numbers of users.

D. Numerical Illustration

We provide a graphical illustration of our results considering a scenario in which the average activity period of the users is twice the time spent downloading the movie, representing users who tend to keep their application/devices active after watching the movie. We normalize the parameters $d = 1$ and $\bar{T}_d = 1$, and thus we set $\bar{T} = 2$.




Fig. 1. Average server bandwidth \bar{S} versus the average number of users \bar{N} , for different values of the system load γ , in the case $d = 1$, $\bar{T}_d = 1$, $\bar{T} = 2$, and exponentially distributed upload bandwidth.

Figure 1 reports on a log-log scale the average server bandwidth \bar{S} as function of the average number of users \bar{N} , for different values of the system load γ . The peer upload bandwidth is exponentially distributed with mean $\bar{U} = 1/(2\gamma)$. We compare the upper bound (12) (labeled UB) with the exact solution presented in Section III-C. We also report for $\gamma > 1$ the lower bound (4).

In compliance to Theorem 3, we observe that, as the number of users grows large, the average server bandwidth decreases to zero for $\gamma < 1$ (self-sustainability), whereas it tends to increase linearly with \bar{N} for $\gamma > 1$, approaching asymptotically the lower bound.

In the surplus mode ($\gamma < 1$), the average server bandwidth reaches a maximum value for a certain number of users, which increases as γ increases. Comparing the upper bound with the exact solution, we observe that, although the bound can be pessimistic up to a factor about 4, the bound captures well the qualitative behavior of the exact curve.




Fig. 2. Average server bandwidth \bar{S} versus the average number of users \bar{N} , for different values of the coefficient of variation (CV) of user upload bandwidth, in the case $d = 1$, $\bar{T}_d = 1$, $\bar{T} = 2$, $\gamma = 0.9$.

To show the impact of peer bandwidth heterogeneity, we consider the same scenario as before, keeping the load fixed to $\gamma = 0.9$ (surplus mode), and varying the coefficient of variation (CV) of the upload bandwidth distribution of the users. In particular, we assume that the upload bandwidth is distributed according to a second-order hyper-exponential distribution with balanced means, which could well describe the situation in which we have many peers with low upload bandwidth (*e.g.*, behind ADSL lines) and few peers with large upload bandwidth (*e.g.*, connected with fiber or LAN). We observe the strong impact of the CV on the resulting server bandwidth. Since $\gamma = 0.9 < 1$, Theorem 3 guarantees that \bar{S} goes to zero as $\bar{N} \rightarrow \infty$, however the maximum value of \bar{S} is achieved for quite large number of users (in the order of thousands) for large values of the CV. Again, the analytical upper bound follows well the qualitative behavior of the system, in all considered cases.

IV. RELATED WORK

We restrict ourselves to mentioning theoretical performance studies of peer-assisted content distribution systems, which are closely related to our work.

Stochastic fluid models for BitTorrent-like file-sharing system, accounting for the dynamics of downloaders and seeds, have been proposed for both transient and steady-state regimes [14], [15], but they are not directly applicable to streaming

systems. In [16], authors adapt the fluid model in [15] to VoD systems, investigating the impact of different piece selection policies (rarest-first and in-order) on download latency and startup delay, in the case of homogeneous peers. In contrast to [16], we focus on the self-sustainability of VoD systems with strict service guarantees and heterogeneous user upload bandwidths.

A stochastic fluid approach to analyze peer-assisted video distribution has been proposed in [17] in the context of live streaming, in which (heterogeneous) peers download and playback content synchronously. Here we apply the stochastic fluid approach to VoD systems, whose dynamics are quite different from live streaming, since users can watch the video asynchronously.

The mathematical formulation (7) for the server bandwidth needed by a VoD system based on sequential delivery, appeared in [3], in which authors resort to a Monte Carlo approach to get basic insights into the system behavior (like surplus and deficit modes).

The same formulation (7) has been considered in [11], where authors explore by simulation the effectiveness of different replication strategies to minimize the server load in the slightly surplus mode, as well as distributed replacement algorithms to achieve it. To the best of our knowledge, we are the first to analytically study the stochastic process (7), establishing its connection with random walks and risk theory.

In [18], a per-chunk capacity model is developed to show the tradeoff that exists between system throughput, sequentiality of downloaded data and robustness to heterogeneous network conditions. Optimal content placement strategies to maximize the upload capacity of (homogeneous) set-top-boxes (and thus minimize the servers workload) in VoD systems have been recently investigated in [7] under many-user asymptotic.

V. CONCLUSION

We have developed a stochastic fluid methodology that allows to derive analytical upper bounds to the bandwidth requested from the servers in peer-assisted VoD systems, studying the performance achieved by the simple sequential video distribution scheme. Our bounds hold under the only assumption that the upload bandwidth distribution of peers is light-tailed. We have also proposed an analytical methodology to exactly estimate the bandwidth requested from servers when peer upload bandwidth is phase-type distributed. Besides being analytically tractable, the simple sequential delivery scheme is also an attractive solution in real systems, for two main reasons: i) it allows users to immediately start watching the requested movie; ii) it is simple to manage and control. Moreover, we have proved that the sequential delivery scheme leads to an asymptotically optimal exploitation of the peers' upload bandwidths as the number of users grows large. Indeed, our bounds tightly characterize the asymptotic performance of large-scale peer-assisted VoD systems employing both sequential and non-sequential delivery schemes.

REFERENCES

- [1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010–2015," White paper published on Cisco web site, 2011.
- [2] C. Huang, J. Li, and K. W. Ross, "Can Internet Video-on-Demand Be Profitable?" in *ACM SIGCOMM*, 2007.

- [3] Y. Huang, T. Z. J. Fu, D. ming Chiu, J. C. S. Lui, and C. Huang, "Challenges, Design and Analysis of a Large-scale P2P VoD System," in *ACM SIGCOMM*, 2008.
- [4] "PPLive, <http://www.gridcast.cn/>. GridCast, <http://www.gridcast.cn/>. PP-Stream, <http://www.ppstream.com/>. TVU, <http://www.tvunetworks.com/>. SopCast, <http://www.sopcast.com/>."
- [5] S. Borst, V. Gupta, and A. Walid, "Distributed Caching Algorithms for Content Distribution Networks," in *IEEE INFOCOM*, 2010.
- [6] K. Suh, C. Diot, J. Kurose, L. Massoulié, C. Neumann, D. Towsley, and M. Varvello, "Push-to-Peer Video-on-Demand system: design and evaluation," *JSAC*, vol. 25, no. 9, pp. 1706–1716, 2007.
- [7] B. Tan and L. Massoulié, "Optimal Content Placement for Peer-to-Peer Video-on-Demand Systems," in *IEEE INFOCOM*, 2011.
- [8] T. Bonald, L. Massoulié, F. Mathieu, D. Perino, and A. Twigg, "Epidemic Live Streaming: Optimal Performance Trade-Offs," in *ACM SIGMETRICS*, 2008.
- [9] S. Liu, R. Zhang-Shen, W. Jiang, J. Rexford, and M. Chiang, "Performance Bounds for Peer-Assisted Live Streaming," in *ACM SIGMETRICS*, 2008.
- [10] D. Ciullo, V. Martina, M. Garetto, E. Leonardi, and G. L. Torrisi, "Performance Analysis of Non-stationary Peer-assisted VoD Systems," in *IEEE INFOCOM Mini-Conference*, 2012.
- [11] W. Wu, J. Lui, "Exploring the Optimal Replication Strategy in P2P-VoD Systems: Characterization and Evaluation," in *IEEE INFOCOM*, 2011.
- [12] S. Asmussen, *Applied Probability and Queues*. 2nd ed. Springer-Verlag, New York, 2003.
- [13] D. Ciullo, V. Martina, M. Garetto, E. Leonardi, and G. L. Torrisi, "Stochastic Analysis of Self-Sustainability in Peer-Assisted VoD Systems," technical report, available at <http://www.telematica.polito.it/~emilio/papers/Infocom12-1.pdf>.
- [14] X. Yang and G. de Veciana, "Service Capacity of Peer to Peer Networks," in *IEEE INFOCOM*, 2004.
- [15] D. Qiu and R. Srikant, "Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks," in *ACM SIGCOMM*, 2004.
- [16] N. Parvez, C. Williamson, A. Mahanti, and N. Carlsson, "Analysis of BitTorrent-like Protocols for On-Demand Stored Media Streaming," in *ACM SIGMETRICS*, 2008.
- [17] R. Kumar, Y. Liu, , and K. Ross, "Stochastic Fluid Theory for P2P Streaming Systems," in *IEEE INFOCOM*, 2007.
- [18] B. Fan, D. Andersen, M. Kaminsky, and K. Papagiannaki, "Balancing Throughput, Robustness, and In-Order Delivery in P2P VoD," in *ACM CoNEXT*, 2010.

APPENDIX A PROOF OF THEOREM 1

Define $X_i \triangleq d - U_i - A$, for all $i \geq 1$ and $Q(k) \triangleq \sum_{i=1}^k X_i$. Since $\max\{0, d - \bar{U}\} \leq A < d$ we have $\mathbb{E}[X_i] < 0$ and $\mathbb{P}(X_i > 0) > 0$. Note that by (8), $Z_d(0) = 0$, while, for $k > 0$, it holds:

$$\begin{aligned} Z_d(k) &= \max_{1 \leq j \leq k} \sum_{i=1}^j (d - U_i) \\ &\leq \max_{1 \leq j \leq k} \left\{ \sum_{i=1}^j (d - U_i) + (k - j)A \right\} = \left(\max_{1 \leq j \leq k} Q(j) \right) + kA. \end{aligned}$$

For $k > 0$

$$\mathbb{P}(Z_d(k) > w + kA) \leq \mathbb{P}(\max_{1 \leq j \leq k} Q(j) + kA > w + kA) \leq e^{-\theta^* w},$$

and then $\mathbb{P}(Z_d(k) > w) \leq e^{-\theta^* w} e^{\theta^* kA}$.

By (9), $S_d(k) = d + \max\{0, Z_d(k - 1)\}$. It is easy to prove that the event $\{\max\{0, Z_d(k - 1)\} > w'\}$ is equal to the event $\{Z_d(k - 1) > w'\}$, for all $w' \geq 0$. Therefore, $\mathbb{P}(S_d(k) > w) = \mathbb{P}(Z_d(k - 1) > w - d)$ for all $w \geq d$ and $\mathbb{P}(S_d(k) > w) = 1$ for $w < d$. Thus, for $w < d$ it holds

$$\begin{aligned} \mathbb{P}(S_d > w) &= \sum_{k=1}^{\infty} \mathbb{P}(S_d(k) > w) \mathbb{P}(N_d = k) \\ &= \sum_{k=1}^{\infty} \mathbb{P}(N_d = k) = 1 - \mathbb{P}(N_d = 0) = 1 - e^{-\bar{N}_d} = C_1. \end{aligned}$$

On the other hand, for $w \geq d$, we have

$$\begin{aligned} \mathbb{P}(S_d > w) &= \sum_{k=1}^{\infty} \mathbb{P}(S_d(k) > w) \mathbb{P}(N_d = k) \\ &= \sum_{k=1}^{\infty} \mathbb{P}(Z_d(k-1) > w-d) \frac{\bar{N}_d^k e^{-\bar{N}_d}}{k!} \\ &\leq \mathbb{P}(Z_d(0) > w-d) \bar{N}_d e^{-\bar{N}_d} + \\ &\quad e^{-\theta^*(w-d)} e^{-\bar{N}_d} e^{-\theta^* A} \sum_{k=2}^{\infty} \frac{(e^{\theta^* A} \bar{N}_d)^k}{k!} \\ &= e^{-\theta^*(w-d)} e^{-\bar{N}_d} e^{-\theta^* A} \sum_{k=2}^{\infty} \frac{(e^{\theta^* A} \bar{N}_d)^k}{k!} = C_2 e^{-\theta^*(w-d)}. \end{aligned}$$

Moreover, if $w \geq d$, we can get another simple bound as follows:

$$\begin{aligned} \mathbb{P}(S_d > w) &= \sum_{k=1}^{\infty} \mathbb{P}(S_d(k) > w) \mathbb{P}(N_d = k) \\ &= \sum_{k=1}^{\infty} \mathbb{P}(Z_d(k-1) > w-d) \frac{\bar{N}_d^k e^{-\bar{N}_d}}{k!} \\ &= \sum_{k=2}^{\infty} \mathbb{P}(Z_d(k-1) > w-d) \frac{\bar{N}_d^k e^{-\bar{N}_d}}{k!} \\ &\leq \sum_{k=2}^{\infty} \frac{\bar{N}_d^k e^{-\bar{N}_d}}{k!} = 1 - e^{-\bar{N}_d} - \bar{N}_d e^{-\bar{N}_d} = C_3. \end{aligned}$$

Therefore

$$\mathbb{P}(S_d > w) \leq \begin{cases} \min\{C_2 e^{-\theta^*(w-d)}, C_3\} & w \geq d \\ C_1 & 0 \leq w < d \end{cases}$$

APPENDIX B PROOF OF THEOREM 3

We preliminarily recall that by Theorem 2, we have

$$\bar{S} \leq C_1 F_{S_{\text{seed}}}(d) d + C_2 e^{\theta^* d} \mathbb{E}[e^{-\theta^* S_{\text{seed}}}] / \theta^* \triangleq \bar{S}_{\text{up},1}. \quad (20)$$

Moreover, the first term in the sum above goes to zero, as $\lambda \rightarrow \infty$. Indeed, since $\lim_{\lambda \rightarrow \infty} \bar{N}_d = \infty$, we have $C_1 \rightarrow 1$ as $\lambda \rightarrow \infty$, and the claim follows noticing that $F_{S_{\text{seed}}}(d)$ tends to zero, as $\lambda \rightarrow \infty$ (i.e., the mean number of seeds \bar{N}_{seed} and their offered bandwidth tend to infinity).

We first consider the case $\gamma < 1$. If $d \leq \bar{U}$, by Theorem 1, ϵ may be freely chosen in the interval $[\bar{U} - d, \bar{U})$. We set $\epsilon = \bar{U} - d$, obtaining $A = 0$. For the second term in the sum (20), note that due to $A = 0$ we have $C_2 = e^{-\bar{N}_d} (e^{\bar{N}_d} - \bar{N}_d - 1) = 1 - e^{-\bar{N}_d} \bar{N}_d - e^{-\bar{N}_d}$. So, using again $\lim_{\lambda \rightarrow \infty} \bar{N}_d = \infty$, we deduce that $C_2 \rightarrow 1$ as $\lambda \rightarrow \infty$. Combining this with the relations:

$$\mathbb{E}[e^{-\theta^* S_{\text{seed}}}] = e^{\lambda \bar{T}_{\text{seed}} (\phi_U(-\theta^*) - 1)}$$

and $\phi_U(-\theta^*) - 1 < 0$ (this latter inequality holds since $\phi_U(-\theta^*) = \mathbb{E}[e^{-\theta^* U}] < 1$), we easily have that even the second term in the sum (20) tends to zero as $\lambda \rightarrow \infty$. Consequently, for $\gamma < 1$ and $d \leq \bar{U}$, we get $\lim_{\lambda \rightarrow \infty} \bar{S}_{\text{up},1} = \lim_{\lambda \rightarrow \infty} \bar{S} = 0$. Now, suppose $d > \bar{U}$. Since U is not constant the equation in t : $e^{-t\epsilon} \mathbb{E}[e^{t(\bar{U}-U_1)}] = 1$ has a unique solution, say $\theta^*(\epsilon)$. The properties of the function $\epsilon \mapsto \theta^*(\epsilon)$ are given in Proposition 1 below. For λ large, consider the sequence

$$\{\epsilon_\lambda\} \subset (0, \bar{U} - \text{ess inf } U_1) \text{ defined by} \\ \epsilon_\lambda := (\theta^*)^{-1}(\lambda^{-1/2}).$$

Note that by Proposition 1 $\epsilon_\lambda \rightarrow 0$ and $\theta(\epsilon_\lambda) \rightarrow 0$, as $\lambda \rightarrow \infty$. Furthermore $\lambda \theta^*(\epsilon_\lambda) \rightarrow \infty$ as $\lambda \rightarrow \infty$. We neglect again the first term in (20), and we obtain³ $\bar{S}_{\text{up},1} \sim C_2 e^{\theta^* d} \mathbb{E}[e^{-\theta^* S_{\text{seed}}}] / \theta^* = C_2 e^{\theta^* d} e^{\lambda \bar{T}_{\text{seed}} (\phi_U(-\theta^*) - 1)} / \theta^*$. We can say that $\bar{S}_{\text{up},1} \rightarrow 0$ if and only if $\log \bar{S}_{\text{up},1} \rightarrow -\infty$. Thus, we consider

$$\log \bar{S}_{\text{up},1} \sim \log C_2 - \log \theta^* + \theta^* d + \lambda \bar{T}_{\text{seed}} (\phi_U(-\theta^*) - 1).$$

Using the Taylor expansion (and neglecting the term $e^{-\theta^* A}$ that tends to 1 as $\lambda \rightarrow \infty$), we obtain that

$$\log C_2 \leq -\theta^* A + \bar{N}_d (e^{\theta^* A} - 1)$$

as $\lambda \rightarrow \infty$. Using the Taylor expansion as $\lambda \rightarrow \infty$, and the choice of ϵ_λ above (thus $\theta^* A \rightarrow 0$ and $\theta^* \geq \lambda^{-1/2}$), we have

$$\begin{aligned} \log \bar{S}_{\text{up},1} &\leq -\theta^* A + \bar{N}_d (e^{\theta^* A} - 1) - \log \theta^* + \theta^* d + \\ &\quad \bar{N}_{\text{seed}} (\phi_U(-\theta^*) - 1) \\ &\leq \lambda \bar{T}_d (\theta^* A + o(\theta^* A)) - \log(\theta^*) \\ &\quad + \lambda \bar{T}_{\text{seed}} (-\bar{U} \theta^* + o(\theta^*)) \\ &\sim \lambda \bar{T}_d \theta^* A - \log(\theta^*) - \lambda \bar{T}_{\text{seed}} \bar{U} \theta^* \\ &\leq \lambda \theta^* (\bar{T}_d A - \bar{T}_{\text{seed}} \bar{U}) - 1/2 \log \lambda \\ &\quad \rightarrow -\infty \end{aligned}$$

since in the regime $\gamma = \frac{\bar{T}_d d}{\bar{U} \bar{T}} < 1$, the quantity $\bar{T}_d A - \bar{T}_{\text{seed}} \bar{U}$ is negative. Therefore, the theorem follows.

We consider the case $\gamma > 1$. By Theorem 2 we have that

$$\begin{aligned} \bar{S} &\leq C_1 F_{S_{\text{seed}}}(d) d + C_3 \left(\frac{2}{\theta^*} + w^* - \bar{S}_{\text{seed}} \right. \\ &\quad \left. + \mathbb{E}[e^{\theta^* S_{\text{seed}}}] e^{-\theta^* w^*} / \theta^* \right) \triangleq \bar{S}_{\text{up},2}. \end{aligned} \quad (21)$$

In (21) we can neglect the term $C_1 F_{S_{\text{seed}}}(d) d \sim 1$. Note that as $\lambda \rightarrow \infty$, $C_3 \sim 1$. Thus, we obtain:

$$\bar{S}_{\text{up},2} = \frac{2}{\theta^*} + w^* - \bar{S}_{\text{seed}} + \mathbb{E}[e^{\theta^* S_{\text{seed}}}] e^{-\theta^* w^*} / \theta^*$$

The quantity w^* , as $\lambda \rightarrow \infty$ becomes $w^* = (1/\theta^*) \log(C_2/C_3) + d \sim \bar{N}_d A \sim \lambda \bar{T}_d (d - \bar{U})$. Therefore, as $\lambda \rightarrow \infty$, it is easy to prove that $\log(\mathbb{E}[e^{\theta^* S_{\text{seed}}}] e^{-\theta^* w^*} / \theta^*) \rightarrow -\infty$; as before, we can conclude that $\mathbb{E}[e^{\theta^* S_{\text{seed}}}] e^{-\theta^* w^*} / \theta^* \rightarrow 0$.

Finally, we obtain, for $\lambda \rightarrow \infty$, $\bar{S}_{\text{up},2} \sim -\bar{S}_{\text{seed}} + w^* + 1/\theta^* \sim \lambda(-\bar{T}_{\text{seed}} \bar{U} + \bar{T}_d (d - \bar{U})) + \lambda^{1/2} \sim \lambda(\bar{T}_d (d - \bar{U}) - (\bar{T} - \bar{T}_d) \bar{U}) = \lambda(\bar{T}_d d - \bar{U} \bar{T})$,

or, equivalently:

$$\lim_{\lambda \rightarrow \infty} \frac{\bar{S}_{\text{up},2}}{(\bar{N}_d d - \bar{U} \bar{N})} = 1.$$

Note that the quantity $\bar{N}_d d - \bar{U} \bar{N}$ is a lower bound for \bar{S} , as described in (4). Therefore, necessarily

$$\liminf_{\lambda \rightarrow \infty} \frac{\bar{S}}{(\bar{N}_d d - \bar{U} \bar{N})} \geq 1. \text{ Recalling that } \bar{S} \leq \bar{S}_{\text{up},2},$$

we obtain,

$$1 \leq \lim_{\lambda \rightarrow \infty} \frac{\bar{S}}{(\bar{N}_d d - \bar{U} \bar{N})} \leq \lim_{\lambda \rightarrow \infty} \frac{\bar{S}_{\text{up},2}}{(\bar{N}_d d - \bar{U} \bar{N})} = 1$$

³With abuse of notation we will use the expression $f \sim g$ to indicate that $f \in \Theta(g)$, i.e., f is bounded both above and below by g asymptotically

and the theorem follows.

Proposition 1: If $d > \bar{U}$, then the equation in t $\mathbb{E}[e^{t(\bar{U}-U_1-\epsilon)}] = 1$ admits a unique solution for $\epsilon \in (0, \bar{U} - \text{ess inf } U_1)$. Furthermore, $\theta^*(\epsilon) = \arg_{t>0}(e^{-t\epsilon}\mathbb{E}[e^{t(\bar{U}-U_1)}] = 1)$ is strictly increasing and C^1 on the interval $(0, \bar{U} - \text{ess inf } U_1)$. Finally, it holds $\lim_{\epsilon \rightarrow 0} \theta^*(\epsilon) = 0$.

Proof: We define the function $f(t, \epsilon) = \mathbb{E}[e^{t(\bar{U}-U_1-\epsilon)}]$. Observe that $f(t, \epsilon)$ is analytic in the domain $t \geq 0$ and $\epsilon \geq 0$, as immediate consequence of the fact that $\bar{U} - U_1 \leq \bar{U} < \infty$.

Observe also that i) $f(0, \epsilon) = 1$ and $f'(0, \epsilon) = -\epsilon < 0$ for any $\epsilon > 0$; ii) $f(t, \epsilon)$ is convex in t , since $\frac{\partial^2 f(t, \epsilon)}{\partial t^2} = \mathbb{E}[(\bar{U} - U_1 - \epsilon)^2 e^{t(\bar{U}-U_1-\epsilon)}] > 0$; iii) $\lim_{t \rightarrow \infty} f(t, \epsilon) = \infty$ for any $\epsilon < \bar{U} - \text{ess inf } U_1$. This because, for all $c \in \mathbb{R}$,

$$\begin{aligned} f(t, \epsilon) &= \int_{-\infty}^{\infty} e^{t(\bar{U}-w-\epsilon)} dF_U(w) \geq \\ &\geq \int_{-\infty}^c e^{t(\bar{U}-w-\epsilon)} dF_U(w) \geq e^{t(\bar{U}-c-\epsilon)} \Pr(U_1 < c). \end{aligned}$$

Since $\epsilon < \bar{U} - \text{ess inf } U_1$, there exists $a > 1$ such that $(\bar{U} - \epsilon)/a > \text{ess inf } U_1$. Defining $c = (\bar{U} - \epsilon)/a$, we have $e^{t(\bar{U}-c-\epsilon)} \rightarrow \infty$ while $\Pr(U_1 < c) > 0$.

As a consequence of i) , ii) and iii), recalling that $f(t, \epsilon)$ is continuous w.r.t. t for any $\epsilon \geq 0$ and $t \geq 0$ there is a unique solution $\theta^*(\epsilon) = \arg_{t>0}(e^{-t\epsilon}\mathbb{E}[e^{t(\bar{U}-U_1)}] = 1)$.

The regularity of $\theta^*(\epsilon)$ with respect to ϵ immediately follows by the implicit function theorem. At last the monotonicity of $\theta^*(\epsilon)$ can be derived again from the implicit function theorem, according to which $\frac{d\theta^*(\epsilon)}{d\epsilon} = -\frac{\frac{\partial f(\theta^*(\epsilon), \epsilon)}{\partial \epsilon}}{\frac{\partial f(\theta^*(\epsilon), \epsilon)}{\partial t}}|_{t=\theta^*(\epsilon)}$. Note indeed that $\frac{\partial f(\theta^*(\epsilon), \epsilon)}{\partial \epsilon} < -\theta^*(\epsilon)f(\theta^*(\epsilon), \epsilon) < 0 \forall \epsilon > 0$, while $\frac{\partial f(t, \epsilon)}{\partial t}|_{t=\theta^*(\epsilon)} > 0$ by construction, since $f(t, \epsilon)$ is convex w.r.t. t and $f(t, \epsilon) < 1$ for $0 < t < \theta^*(\epsilon)$ and $f(t, \epsilon) > 1$ for $t > \theta^*(\epsilon)$.

At last, it is immediate to see that also for $\epsilon \rightarrow 0$ $\theta^*(\epsilon) \rightarrow 0$, in light of the fact that $f(0, 0) = 1$, and $f(t, 0) > 1$ for $t > 0$. \blacksquare

As immediate consequence of the fact that $\theta^*(\epsilon)$ is strictly increasing (and thus invertible) and continuous over the domain $(0, \bar{U} - \text{ess inf } U_1)$ with $\lim_{\epsilon \rightarrow 0} \theta^*(\epsilon) = 0$, we have that the following proposition holds.

Proposition 2: Provided that $d > \bar{U}$, and U not constant, the image of $\theta^*(\epsilon)$ for $0 < \epsilon < \bar{U} - \text{ess inf } U_1$ is the open interval $(0, \delta)$, with $\delta = \lim_{\epsilon \rightarrow \bar{U} - \text{ess inf } U_1} \theta^*(\epsilon)$.

APPENDIX C

DERIVATION OF THE EXACT SOLUTION IN (17)

In this section we derive the solution of (16) for the case in which the bandwidths U are exponentially distributed. We point out that the proof of Theorem 4 can be obtained leveraging exactly the same arguments. Furthermore we wish to emphasize that the same approach can be extended to the case in which the bandwidth U has a general phase-type distribution.

When the bandwidth is exponentially distributed we have: $dF_Z(\alpha | 1) = \frac{1}{\bar{U}} e^{-\frac{d-\alpha}{\bar{U}}} \mathbb{I}_{\alpha \leq d} d\alpha$. Thus, from (15), we have

$$\begin{aligned} F_Z(w | k) &= \int_{-\infty}^w F_Z(w - \alpha | k - 1) \frac{1}{\bar{U}} e^{-\frac{d-\alpha}{\bar{U}}} \mathbb{I}_{\alpha \leq d} d\alpha \\ &= \int_{-\infty}^{\min\{w, d\}} F_Z(w - \alpha | k - 1) \frac{1}{\bar{U}} e^{-\frac{d-\alpha}{\bar{U}}} d\alpha \end{aligned} \quad (22)$$

If $w < d$, making the substitution $y = w - \alpha$ in the integral, we obtain:

$$F_Z(w | k) = \frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_0^\infty F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy.$$

Note that the integrand function does not depends on w , thus the whole integral can be regarded as a constant. Moreover, notice that

$$\frac{1}{\bar{U}} \int_0^\infty F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy = F_Z(d | k). \quad (23)$$

Thus, it holds:

$$F_Z(w | k) = e^{-\frac{w-d}{\bar{U}}} F_Z(d | k) \quad \forall w \leq d \quad (24)$$

Now consider equation (22) when $d \leq w \leq 2d$:

$$\begin{aligned} F_Z(w | k) &= \frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_{w-d}^\infty F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy \\ &= \frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_0^\infty F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy + \\ &\quad - \frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_0^{w-d} F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy. \end{aligned} \quad (25)$$

Note that the first term in the sum is equal to $e^{-\frac{w-d}{\bar{U}}} F_Z(d | k)$. For the second integral, since $d \leq w \leq 2d$, variable y is such that $0 < y < w - d < d$. Thus, using (23) we can write:

$$\begin{aligned} &\frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_0^{w-d} F_Z(y | k - 1) e^{-\frac{y}{\bar{U}}} dy \\ &= \frac{e^{-\frac{w-d}{\bar{U}}}}{\bar{U}} \int_0^{w-d} e^{-\frac{y-d}{\bar{U}}} F_Z(d | k - 1) e^{-\frac{y}{\bar{U}}} dy \\ &= \frac{(w-d)}{\bar{U}} e^{-\frac{w-2d}{\bar{U}}} F_Z(d | k - 1). \end{aligned} \quad (26)$$

Thus, for $d \leq w \leq 2d$, we obtain

$$F_Z(w | k) = e^{-\frac{w-d}{\bar{U}}} F_Z(d | k) - \frac{(w-d)}{\bar{U}} e^{-\frac{w-2d}{\bar{U}}} F_Z(d | k - 1).$$

Now considering $2d \leq w \leq 3d$, we can still use (25) to express $F_Z(w | k)$ in terms of $F_Z(y | k - 1)$ over a domain in which $y \leq w - d \leq 2d$. Again we know explicitly the expression of $F_Z(w | k)$ over the considered domain in terms of the two constants $F_Z(d | k - 1)$ and $F_Z(d | k - 2)$. It turns out:

$$\begin{aligned} F_Z(w | k) &= e^{-\frac{w-d}{\bar{U}}} F_Z(d | k) - \frac{(w-d)}{\bar{U}} e^{-\frac{w-2d}{\bar{U}}} F_Z(d | k - 1) + \\ &\quad + \frac{(w-2d)^2}{2\bar{U}^2} e^{-\frac{w-3d}{\bar{U}}} F_Z(d | k - 2) \quad 2d \leq w \leq 3d \end{aligned} \quad (27)$$

Proceeding in a similar way we can express $F_Z(w | k)$ for any $w \leq kd$ in terms of the constants $F_Z(d | 1) \dots F_Z(d | k)$, while for $w > kd$ we have trivially $F_Z(w | k) = 1$.

The constants $F_Z(d | k)$ can be obtaining forcing $F_Z(kd | k) = 1$. Indeed, by imposing $F_Z(w | 1)|_{w=d} = 1$ we immediately obtain $F_Z(d | 1) = 1$. Imposing $F_Z(w | 2)|_{w=2d} = 1$ we obtain an algebraic linear equation between $F_Z(d | 2)$ and $F_Z(d | 1)$, from which we can derive $F_Z(d | 2)$. In general imposing $F_Z(w | k)|_{w=kd} = 1$ we obtain a linear algebraic equation containing all constant $F_Z(d | i)$ with $i \leq k$. This equation can be exploited to derive $F_Z(d | k)$ as function of $F_Z(d | i)$ with $i < k$.