

Successfully Mapping DASH over a P2P Live Streaming Architecture

Laura Natali and Maria Luisa Merani

Abstract—This work proposes to combine Peer-to-Peer (P2P) and Dynamic Adaptive Streaming over HTTP (DASH), leveraging the scalability and self-organization properties of the former and the capability of the latter to adapt the rate of the delivered video to varying operating conditions. The devised P2P-DASH architecture is a multi-overlay P2P platform, where every peer implements a decentralized rate control strategy that steers its transitions from one overlay to another, so as to achieve a good viewing quality and also preserve the good functioning of the entire system. At peer’s site, the new architecture displays the following salient features: (i) an augmented pool of neighbors, to provide every node with the knowledge of the streaming process within adjacent overlays; (ii) a priority mechanism to speed up the delivery of video chunks to peers that newly enter the system or switch from the current to a new overlay, in order to quickly turn them into active peers; (iii) some alternative rate control algorithms to guarantee, with different degrees, more resources to overlays that deliver higher streaming rates, therefore reducing the misalignment among streaming processes which occur in different overlays, and ultimately improve the viewing quality of the single user. System behavior is extensively investigated through simulation and some clear design guidelines are provided.

keywords: DASH, peer-to-peer, video streaming, switching delay

I. INTRODUCTION

The usage of video as the primary communication form over the Internet is a constantly growing experience for the majority of contemporary societies: the trend manifests itself through several emblematic service offerings, some of which are very recent and pervasive. YouTube and NetFlix, the most popular web platforms for sharing videos and distributing movies and television programs, respectively, are by now the forerunners of the phenomenon, which is unceasingly enriched by new applications and features: Periscope and Meerkat are offering their users a way to directly stream videos from the smartphone and likewise to follow videos being streamed by others; even Facebook is introducing video as profile picture. This steep increase in multimedia content fruition has a profound impact on the amount of traffic that revolves around Internet video servers: Youtube and Netflix consume about 50 percent of North American downstream bandwidth during peak hours and 30 percent in Europe [1] [2], and this poses powerful challenges on the quality of the offered services, also threatened by the varying traffic conditions in the underlying

network infrastructure. In parallel, Content Delivery Networks have boomed, and one of their transversal features is to be heavily based on the combined use of HTTP and TCP. This entails several benefits, most notably, the employment of commodity servers and stateless services.

To cope with increasing video requests from the users, varying network conditions and to preserve HTTP servers via an interoperable standardized approach, Dynamic Adaptive Streaming over HTTP (DASH) has been developed [3], [4]. A DASH server typically encodes video content at multiple bit rates and provides a manifest file to describe the structure of the media segments and the different video versions that it makes available; it is up to the client to decide which version, also termed representation, to stream. The large diffusion of DASH is testified by its adoption from companies that are leaders in video on demand and video streaming and according to the survey in [5], most of the European broadcasters will provide DASH contents in the very next future too. Yet, servers and available bandwidth remain critical elements of such architectures: scalability and network resources emerge as the real challenges, which await for a sound solution. The peer-to-peer (P2P) paradigm, often put aside and discarded with conceit, can answer this quest. This is the premise embraced by the current work, whose proposal rests on a multi-overlay P2P architecture, where there are as many swarms as the number of available DASH representations and every swarm is responsible for distributing one of them; some alternative control strategies rule the transitions of the peers from one overlay to another. The control logic is implemented at the peer’s site and hinges on locally measured parameters, as well as on macroscopic system indicators, to secure the quality that the single peer experiences, as well as the good functioning of the entire platform.

For the proposed architecture, the goal of the current study is multifaceted:

- on the DASH side, given the system is thought for near real-time services, the aim is to understand how and to what extent the flexibility that DASH offers in switching from one representation to another, which can occur with a time granularity as low as the segment duration, plays in favor of a timely migration among overlays;
- on the P2P side, the goal is to put forth and analyze the performance of various system solutions and of two different control algorithms steering the node movements among overlays, so as to ensure that the video is streamed in the smoothest way and that the node experiences no interruptions when moving from one DASH representa-

Manuscript received October, 2015; revised January, 2016; accepted February, 2016. Date of publication YYYY.

L. Natali and M.L. Merani are with the Engineering Department “Enzo Ferrari” of the University of Modena and Reggio Emilia, Italy (e-mail:laura.natali@unimore.it; marialuisa.merani@unimore.it).

tion to a different one. At the same time, the examined solutions strive to have the user view the video at a rate which is as close as possible to the desired one, with a satisfying quality.

We propose a system that exhibits the following salient features: (i) an increased number of neighboring peers from different overlays to boost the exchange of buffer maps, that consent the peer to know the streaming position of its target overlay in advance, as well as which video chunks are currently circulating in adjacent overlays; (ii) the prioritized scheduling of video chunks by parent peers, which serve first the children peers that have just switched to a new overlay; (iii) two different flavors of the rate control algorithm, to either judiciously grant more resources to overlays distributing video at higher streaming rates or to privilege the user's satisfaction.

System behavior is extensively explored, placing an emphasis on the switching delay, i.e., the delay incurred by the peer when transiting from one DASH representation to another, and our major findings are summarized as follows:

- the time granularity of DASH plays in favor of P2P, as it allows to reuse a significant fraction of the streaming buffer content when the peer migrates from one overlay to a new one;
- the buffer reuse mechanism and the increased neighborhood positively combine to reduce the switching delay, to an extent that depends on the misalignment of the streaming processes among overlays;
- the introduction of priority scheduling reduces the switching delay regardless of the alignment among overlays;
- the rate control algorithm can be tuned so as to achieve a better alignment, hence a low switching delay, combined with good operating conditions, at the expense of a suboptimal streaming rate experienced by the peers.

In the remaining part of the paper, Section II offers an overview of the related work; Section III delineates the essential DASH features and it is preparatory to Section IV, that illustrates the proposed architecture resting on DASH and peer-to-peer, its constituent elements and the merit figures examined to assess its performance. Section V discusses the numerical results and Section VI draws the conclusions.

II. RELATED WORK

As regards existing contributions in literature, the bulk of the work on DASH focuses on client-server architectures and in this setting, different algorithms are proposed to dynamically choose the most appropriate DASH representation, in order to cope with bandwidth variations and guarantee the final user the best Quality of Experience (QoE). Several rate adaptation techniques are based on the analysis of bandwidth variations and the streaming buffer level. Referring to bandwidth-based solutions, [6] proposes a strategy that relies upon the TCP download throughput of the client to determine the available bandwidth in the presence of congestion, whereas it constantly probes the network and adapts to its new conditions otherwise; in [7], the authors use feedback control theory to design an adaptive streaming algorithm with the objective to maximize the bit rate in the download and to guarantee a stable buffer

level; [8] suggests a rate adaptation algorithm relying on the smoothed HTTP throughput to determine if the current media bit rate matches the end-to-end network bandwidth. In other cases, the rate adaptation algorithm is based on the status of the playback buffer, as in [9], that strives to guarantee the minimum buffer length. Still other solutions have been recently explored: in [10], fuzzy logic is the core of the rate controller; in contrast to most DASH solutions, [11] implements a server-side open loop rate control (OLAC). Overall, whether the rate control is implemented at client or server side, current conditions are monitored in order to maximize video quality, the focal point being the client perspective: the client is opportunistic in its decision taking process, as the aim is to optimize its viewing experience only. To the authors' knowledge, in the algorithms proposed in literature neither the user is aware of the implications that its local decision has on the entire system, nor it possesses any notion of friendliness towards concurrent connections, this task being left to TCP. Unlike the previous investigations, our proposal strives to guarantee a good performance to the user without loosing the overall system sight. The reason is that in a system built upon the notion of mutual cooperation, as any P2P architecture happens to be, the actions that the single peer performs immediately reflect on all other peers: it is therefore mandatory that, in its decision-making process, the peer takes into account not only its perspective, but also a global prospect. Our former study in [12] already began moving along this path: in a multioverlay scenario, it put forth a rate control algorithm and adopted an Integer Linear Programming approach to provide a benchmark to system performance, also demonstrating that P2P users can experience a very satisfying performance, better than in a conventional P2P setting where DASH is not employed. In [20] we performed an extensive numerical analysis to demonstrate the effectiveness of the envisioned solution, even in the critical occurrence of flash-crowds, finding that they are rapidly detected and the peers steered to the most appropriate overlay.

In the present work, not only we perfection the formerly proposed algorithm, but also conceive innovative structural features to enhance system effectiveness.

There are some explicit contributions in literature that jointly deal with P2P and DASH and that we deem worth being mentioned, namely, [13], [14] and [15]. The work in [13] proposes a standard compliant solution to integrate peer-assisted streaming in conventional DASH client-server systems and it represents a suitable approach in a CDN type of environment. Unlike in the present work, the adaptation logic employed in [13] to modify the requested video rate solely serves the purpose of achieving a reduced system complexity and to guarantee satisfactory savings of server bandwidth. The system heavily relies upon the server presence, and the number of clients considered in the simulation setting is modest, of the order of magnitude of few tens. On the contrary, in our study the server contribution is minimal and we have several thousands of nodes to accommodate: accordingly, the rate control algorithm that we propose has to be carefully crafted. In [14], P2P is integrated in a Personalized Learning System (PLS), the basic idea being to exploit the cooperation among

peers in those settings where a given number of users request the same video content, as it is typically the case for a learning environment such as a campus network. However, content is always made available by the server to a gateway and P2P distribution is restricted to the local site, with a very low number of participating nodes being involved; in our proposal no restriction is posed on the user location and no intermediary devices are foreseen. [15] applies DASH to a P2P architecture and exploits cooperation game theory to rule the node switching process among different representations. The authors present a cost-sharing mechanism that favors the cooperation among peers watching the same representation. Again, the server role is predominant, the size of the peers group is modest and no attention is paid to the complete statistical characterization of system performance.

Overall, as regards the comparison between the current contribution and the correlated works existing in literature, the most notable remarks are that we do not examine a VoD system whose population size is very limited, nor we interpret P2P as a secondary, although beneficial feature of the video distribution architecture. On the contrary, in order to tackle the scalability issue, we believe that a pure P2P solution is the real challenge; our study therefore investigates the feasibility and the achievable performance of a system that delivers a live streaming channel to a large floor of users, with very little contribution from the server. Hence, we explore more sophisticated algorithms and system solutions, assessing system performance from numerous new facets.

III. DASH FUNDAMENTALS

Dynamic Adaptive Streaming over HTTP (DASH) [3] is a standard that serves the purpose of delivering media content through the network via conventional HTTP servers to HTTP clients. As clients may request data from standard web servers with no DASH-specific capabilities, the standard is centered on the data format and on the way to describe – alias present – the media content, as opposed to client-server actions.

In DASH, the media content is composed of a sequence of consecutive temporal periods and each period in turn consists of multimedia components of different types, e.g., the audio component and the video component, that are typically available in different versions: as an example, there might be video versions exhibiting different resolutions and, in the case of pre-recorded material, audio versions in different languages. Within a period, the multimedia content is arranged into adaptation sets, and each set describes a group of interchangeable versions of one or several media content components, that are termed representations. Clients can dynamically switch from one representation to another in response to, e.g., varying network conditions: if the switching takes place properly, it is experienced as seamless by the user.

Within every representation, content is further divided into segments, that contain media data and/or metadata to access, decode and playout the media content; each segment is made accessible providing a URL. In a representation, segments normally exhibit the same duration, which may vary from representation to representation. For streaming services, the

segment duration is rather short, of the order of a few seconds, not to aggravate the end-to-end delay budget.

The Multimedia Presentation Description (MPD), a document in the form of a manifest file, details each period, providing its starting time, its duration and the available adaptation sets; it is the MPD that also provides the URL where segments can be referenced at, together with a segment availability interval, that is, the temporal window in wall-clock time during which the segment can be accessed via the HTTP-URL. The MPD is initially passed to the DASH client, which parses it and next proceeds to issue HTTP get requests to retrieve the segments of the desired representations.

The DASH standard does not indicate which algorithm to employ to rule the switching from representation to representation, that is left entirely open. However, for the dual purpose of enabling synchronization and flawless switching, DASH adopts a common timeline shared among different media components and versions: the segments contain precise timing information about the presentation of media content, that in principle allows clients to switch from representation to representation at an arbitrary point in time; in practice however, this turns out complex for different reasons, most notably because there are coding dependencies within the same representation and because for live, streaming services, segments get available over time, depending on their position in the media presentation timeline. Hence, the most immediate choices suggested by the standard to ease the switching process are: (i) to have each segment begin with a Stream Access Point (SAP), a position in the media stream that enables the playout to start relying only upon the information in the representation data from that position onwards, and (ii) to keep the boundaries of segments aligned across different representations in the same adaptation set. The current work adheres to these suggestions and examines the case of live streaming services provided through DASH to a large population of users.

IV. THE PROPOSAL

A. *Its Rationale*

P2P is the significant feature that rules the distribution of multimedia content within the proposed architecture: users are therefore organized in a virtual network which is built at application layer and are actively involved in the video diffusion process. They relay portions of the content that resides in their local streaming buffer to other users – their peers – through the virtual links of the overlay. We consider the case where a single video channel has to be streamed to a population of users whose size is N and assume that K DASH representations of the video are instantaneously available: unlike VoD client-server architectures, where K takes on relatively high values, in the examined scenario K is intentionally confined to lower values, not to excessively complicate the design of a live system which has to respect tighter latency constraints.

We assume that the bit rate is what distinguishes the different available DASH representations, each being distributed within a separate P2P overlay: the j -th representation displays

a streaming rate r_j and is delivered within the j -th overlay, equivalently termed swarm, $1 \leq j \leq K$; without loss of generality, we set $r_j < r_{j+1}$, $\forall j$, $j = 1, 2, \dots, K - 1$. As Fig.1 indicates, the server is responsible for initially igniting the diffusion process: it delivers each representation to the corresponding overlay as a sequence of video chunks, relatively small fragments of the encoded video. Every peer belongs to one overlay at a time and contributes to sustain the process with its own upload bandwidth, further spreading the video chunks that it receives to other peers that belong to the same overlay and that request them. C distinct peer classes are assumed, each exhibiting different upload and download capacities: hence, peers contribute to the system good functioning differently, depending on the class they belong to.

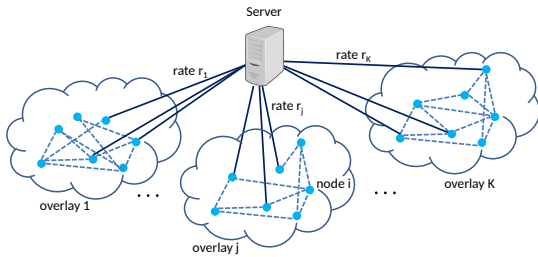


Fig. 1: System architecture

The examined P2P overlays are mesh-based, and within each of them every peer exchanges video chunks with M randomly chosen neighbors, implementing a pull-based protocol. The peer periodically informs its neighbors about the video chunks that it holds in its streaming buffer by forwarding them its buffer maps. Moreover, the peer asks its neighbors for the chunks that are missing within its current request window W . Such window identifies the range where chunks that can be requested by the node fall: its right edge coincides with the highest chunk sequence number that the node is aware of from the inspection of the buffer maps it received, its left edge is X s behind. The request window slides forward whenever the peer learns through the new buffer maps that it receives from its neighbors that a chunk with a higher sequence number is available; when a chunk that is yet to be received falls outside the current request window, it will not be claimed any longer, as outdated. For the sake of clearness, Table I collects the definition of the most relevant adopted terms, as they will frequently appear all throughout the paper.

The DASH representation a peer desires is the representation featuring the rate that the peer aims at viewing the video content at: we observe that it depends on the combined limits set by the user's access network and by the display characteristics of the user's terminal; the user's preferences might also contribute to determine it, as it is the case when different charging rates are applied for distinct video qualities.

A peer requesting the video channel for the first time begins watching representation 1, that features the lowest bit rate, r_1 , hence the lowest quality: this choice allows to start the payout in a reasonably short time, and positively influences the overall

Term	Definition
<i>node</i>	generic entity belonging to a network
<i>peer</i>	entity belonging to a P2P network, where all nodes have the same importance, that is, they are all "peers"
<i>swarm/overlay</i>	group of peers connected through virtual links at application layer
<i>video chunk</i>	small fragment of a video sequence, lasting a few seconds at the most
<i>buffer map</i>	a map representing which video chunks are currently available within the node buffer
<i>representation</i>	specific version of the source multimedia content. This work assumes that each representation is identified by a different bitrate

TABLE I: Glossary

quality that the user experiences [16]. After joining overlay 1, if the peer's desired representation is not the video alternative at rate r_1 , the peer attempts to move upward and possibly succeeds; however, during the streaming the peer might also have to move downward to a lower quality representation, depending on current system conditions, and then it will dynamically attempt to advance again. These actions translate into migrations from one overlay to another: we require that a peer exclusively moves from its current overlay to an adjacent one, i.e., from overlay j to $j + 1$ or to $j - 1$ if $2 \leq j \leq K - 1$, from overlay j to overlay $j + 1$ if $K = 1$, and from K to $K - 1$ if $j = K$; this guarantees the minimal gap between two consecutive representation playouts, as recommended in [17] [18] to confine the amplitude of rate variations and their negative effects on the perceived video quality. Furthermore, we assume that every DASH segment has the same duration and it is made of the same number n of video chunks, so that the chunk duration t_{chunk} is the same in every overlay. This translates into a different amount of media bits being placed in every chunk: indicating by L_j such size for the chunks that are distributed within overlay j , we have

$$L_j = t_{chunk} \cdot r_j. \quad (1)$$

The choice of equal duration segments allows a smooth transition between different representations, as switching involves playing to the end of a segment in a representation and then playing from the beginning of next segment in the new representation. In the realistic circumstance of a time lag between different overlays, the constant segment duration still protects the smoothness of the transitions, as long as the user's buffer can absorb time shifts among different representations.

B. Core Algorithm

In a conventional client-server scenario, DASH provides the user with the functionalities needed to perform adaptive streaming and leaves open the selection of the algorithm that rules the switching among different representations. The client is assigned the task to check its current conditions via the monitoring of various parameters, and to ask the server for a different representation, if needed. As previously observed, proposals in literature rely on the observation of several indicators and put forth different criteria to govern the switching. However, it is exclusively the local perspective that

plays a role in the design of the control algorithm: within the current proposal, although the primary aim of the peer remains to experience a satisfying streaming quality and to efficiently deploy network resources, the user's decisions cannot any longer be taken in isolation, as the peer has to be fair to all other peers, given its viewing quality is strongly influenced by the quality of the streaming process of the entire overlay, that is, it depends on the status of all other peers. It is therefore required that every peer employs two distinct types of status indicators, *local* and *global*, to decide if, when and where to migrate; whereas the former indicators indirectly supply information about the peer's video quality and are locally measured by the peer itself, the latter parameters provide a clue about the current health of the overlays and are periodically handed out by the server to all peers, that rely upon both to enforce the adaptive rate selection algorithm.

Among local indicators, we utilize:

- the Delivery Ratio (DR), defined as the ratio between the number of video chunks that meet the playback deadline over the total number of chunks that a peer should receive, measured with a periodicity of \tilde{t} s. Such ratio indicates the throughput that the peer experiences and indirectly signals the quality of the received video in the very recent past;
- the Request Window State (RWS), defined as the ratio between the number of downloaded video chunks within the current request window W and the size of such window measured in number of video chunks, $0 \leq RWS \leq 1$. This indicator provides an indirect forecast of video quality in the near future, as its value reflects the imminent status of the playout buffer at the peer's site.

Among global indicators, we select those that the monitoring server of a P2P platform typically determines and that can distribute to peers with very little effort, namely:

- the instantaneous resource index [19] of the j -th overlay, $\sigma_j(t)$, $j = 1, 2, \dots, K$, defined as

$$\sigma_j(t) = \frac{C_{Sj} + \sum_{i \in N_j(t)} c_i}{|N_j(t)| \cdot r_j}, \quad (2)$$

where C_{Sj} is the capacity that the server commits to the j -th overlay to distribute the representation with rate r_j , c_i is the upload capacity of the i -th peer belonging to overlay j , $N_j(t)$ is the set of active nodes within such overlay at time t and $|N_j(t)|$ is the set cardinality;

- the instantaneous efficiency of the j -th overlay, defined as

$$E_j(t) = \frac{U_{Sj}(t) + \sum_{i \in N_j(t)} u_i(t)}{|N_j(t)| \cdot r_j} \quad (3)$$

where $U_{Sj}(t)$ is the actual upload rate that the server provides to overlay j at time t and $u_i(t)$ is the actual upload rate at time t of the i -th user within the same overlay.

When the instantaneous resource index $\sigma_j(t)$ takes on a value greater than or equal to 1, in principle the j -th overlay can successfully guarantee the video delivery to all of its members, whereas when its value falls below 1, the overlay operates in a critical regime; hence, σ_j is a high-level indicator of the

overlay health. However, the instantaneous efficiency $E_j(t)$ provides a more accurate picture, as it captures some system behavior that would otherwise go unseen, if σ_j only were examined. To depict a scenario where this happens, it suffices to think of a flash crowd of viewers that abruptly enters the system, wishing to stream the video: newly incoming peers initially act as free riders, as they momentarily have no video chunks to share. As a consequence, the value of the instantaneous efficiency drops, revealing a potentially critical operating condition that the resource index cannot seize.

Let us next consider the generic peer i within overlay j and indicate by $r_d(i)$ the streaming rate of its desired representation; the steps that the algorithm enforced by the peer goes through every Δt seconds are listed below:

- (i) peer i checks its current streaming rate r_j against $r_d(i)$ and if $r_j < r_d(i)$, i.e., if the peer is not satisfied, it first verifies whether it can leave its current overlay or it has to defer its departure. This last circumstance occurs if overlay j is not in good health and node i upload capacity is beneficial to it, that is, if σ_j is lower than 1 and if the peer upload capacity c_i is greater than the streaming rate r_j ; in this case, it is convenient that the peer does not move upward to overlay $j + 1$. If on the contrary nothing prevents the departure, peer i further verifies if its future contribution to the target overlay $j + 1$ will be positive, i.e., if its upload capacity c_i is greater than the streaming rate r_{j+1} . If so, the peer migrates, as it will be beneficial to overlay $j + 1$. If not, the node further checks that overlay $j + 1$ has abundant overall upload capacity and that the video diffusion process within the overlay is taking place in an efficient and noncritical manner, hence the target overlay is able to accommodate a new peer, regardless of it being a relatively "poor" contributor. In detail, the peer moves to overlay $j + 1$ if $\sigma_{j+1} > 1$ and $E_{j+1}(t) > E_{thres}$, where E_{thres} is a properly set threshold.
- (ii) peer i also verifies its local status quality indicators, updating the weighted average of its delivery ratio and request window state in the following manner:

$$DR_i^{(t)} = w_D \cdot DR_i^{(t)} + (1 - w_D) \cdot DR_i^{(t-\Delta t)} \quad (4)$$

and

$$RWS_i^{(t)} = w_W \cdot RWS_i^{(t)} + (1 - w_W) \cdot RWS_i^{(t-\Delta t)}, \quad (5)$$

respectively, where w_D and w_W are the weight coefficients. If $DR_i^{(t)}$ and $RWS_i^{(t)}$ are below their predefined thresholds, DR_{thres} and RWS_{thres} respectively, the viewing quality is not deemed acceptable and the peer scales down to a lower rate representation, hence it moves to overlay $j - 1$.

The pseudo code describing the algorithm locally implemented by the peer to steer its movements among overlays is reported in Alg.1.

Note that in the algorithm the upgrade decision is conservatively influenced by the peer's local metrics – its current rate r_j and its desired streaming rate $r_d(i)$ – and also by the global indicators of the current and destination overlays, whereas the

Algorithm 1 Rate Switching Control Algorithm

```

Node  $i$  in overlay  $j$  every  $\Delta t$  seconds
;verifies its satisfaction
if (  $r_j < r_d(i)$  ) then
;verifies the current overlay status
if (  $(\sigma_j < 1)$  and (  $c_i \geq r_j$  ) ) then
do not migrate to overlay  $j + 1$ ;
else
;verifies the destination overlay status
if ( (  $c_i > r_{j+1}$  ) or (  $\sigma_{j+1} > 1$  and  $E_{j+1} > E_{thres}$  ) )
then
migrate to overlay  $j + 1$ ; exit;
end if
end if
;verifies its viewing quality
if ( (  $DR_i^{(t)} < DR_{thres}$  ) and (  $RWS_i^{(t)} < RWS_{thres}$  ) )
then
migrate to overlay  $j - 1$ ; exit;
else
stay in overlay  $j$ ; exit;
end if

```

downgrade process is exclusively governed by local indicators. The rationale behind this choice is that global indicators point to the overall status of the system and therefore let the peer learn if the migration from its current swarm to the overlay distributing a higher bit rate representation can be safely performed. On the other hand, through local indicators the peer indirectly measures its streaming quality and decides if it is time to scale down to a lower DASH representation, in order to preserve a good viewing experience.

C. Silent Neighborhoods, Buffer Reuse and Priority Queuing

In this subsection we discuss some structural features we introduce in the design of the P2P platform. We begin observing that, in the transition from the current overlay to an adjacent one, the peer loses some time in disconnecting from its old neighbors, building up the new neighborhood, then waiting for buffer maps. Only at this point in time the peer is able to request video chunks. These actions are necessary for the successful insertion of the peer into its target overlay, but inevitably require some time, that might in turn translate into a discontinuous playout: in order to confine such delay, that we term setup time, the proposed system relies upon an additional feature we call *the silent neighborhood*, whose underlying idea is to have the peer maintain contacts with nodes both in its current and also in its adjacent overlays. Namely, we request the peer in overlay j to create and simultaneously maintain a neighborhood in three overlays, the j -th, the $j + 1$ -th and the $j - 1$ -th, and to exchange buffer maps with the nodes belonging to all three neighborhoods: however, as illustrated in Fig.2, only one of such sets of nodes is active at a time, i.e., the peer exchanges both buffer maps and video chunks with the nodes that belong to the “active” neighborhood in the j -th overlay where the peer resides, whereas it only

exchanges buffer maps with its neighboring peers in the two adjacent, “silent” overlays. The construction of the silent neighborhood obeys the same rules as the active one: M peers are randomly chosen within the overlay and whenever a node departs, it is immediately replaced by a new one. Moreover, a reciprocal relationship holds: if node i belongs to node h silent neighborhood, then node h is a silent neighbor of i .

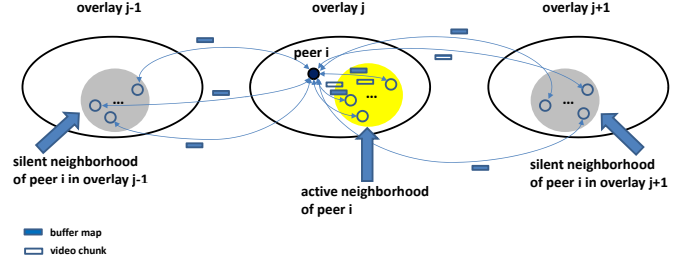


Fig. 2: The silent neighborhood mechanism

There is one first and immediate benefit that the silent neighborhood solution brings in: when the peer migrates to an adjacent overlay, its setup time is zeroed, as the node can immediately ask its silent neighbors to accept it as an active node, assigning some of their actual upload bandwidth to the communication and beginning to deliver content to the peer. Within the process, if the uplink bandwidth of some among the potential new neighbors is exhausted, the peer will look for alternative partners; however, the peer will very unlikely start from scratch in building up its new neighborhood; rather, it will swiftly begin receiving the video chunks currently distributed in the target overlay.

It is indisputable that the silent neighborhood solution implies some signaling overhead, as the generic peer constantly transmits and receives more buffer maps than in the conventional architecture, but the burden is more than acceptable, as it will be quantified in Section V.

As regards the streaming process within the target overlay, note that it can be “early” with respect to the origin overlay of the peer, meaning that the swarm distributes newer video chunks, or it can be “late”, that is, the overlay distributes older video chunks, and the node has to adapt to both conditions: for doing so, it determines the newest advertised video chunk in the target overlay and from there it begins to request chunks backward for the entire width W of its request window; assuming that the head of the newest chunk corresponds to time t_{curr} within the video stream, the requests will fall within the $R = [t_{curr} - W, t_{curr}]$ interval. Both cases are illustrated in Fig.3. In analogy to what happens when the node joins the system for the first time and has to collect a reasonable amount of content before the playout begins, we consider the migration to the new overlay successfully completed when the node has accumulated a given number, say S , of consecutive video chunks within its buffer and we call the delay the node incurs to complete this process the *switching delay*. Luckily, because of the previous assumptions, DASH moves from one representation to another with a time granularity which is equal to the size of a segment, so that at the time the switching

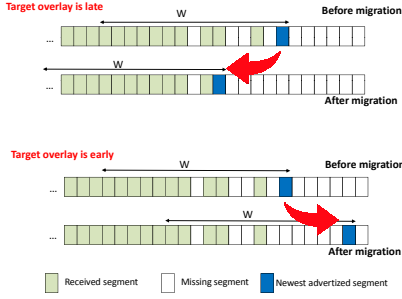


Fig. 3: “Early” and “late” target overlays

takes place the node can profitably inherit all the segments – not necessarily contiguous – within its buffer that fall within R . This buffer content reuse contributes to avoid video stalls and eases the transition to the new video representation. Note that if the buffer reuse mechanism only were introduced, with no silent neighborhoods, the video chunks already present in the peer’s buffer at the time of the transition would not be exploited in the most effective manner: as a matter of fact, to identify which DASH segments can be reused a node needs to know the streaming position of the target overlay; to this end it has to construct the new active neighborhood within the target overlay and then wait for some buffer maps; during this time interval some of the video segments stored in the streaming buffer might become outdated, and if this happens, they have to be discarded to make room for more recent video chunks. On the contrary, when the silent neighborhood mechanism is implemented, the migrating peer is already hooked to the newest advertised chunk in the target overlay and immediately proceeds to check what chunks can be profitably inherited to accumulate S seconds of consecutive video.

We aim at quantifying the benefit that the buffer reuse allowed by DASH guarantees and to have a sound term of comparison, we consider the reference case where the node joining the new overlay is forced to remove from its buffer the video chunks of the representation it was previously receiving and to collect the video chunks of the new representation from scratch. As every DASH segment is composed by n video chunks, at migration time it might happen that the node has not received all chunks of one or more segments: segment fractions are useless and have to be discarded, whereas the completely received segments only are inherited. This statement is made clear by the example portrayed in Fig.4, where the segment with sequence number 135 is missing after the migration process. Yet, non-contiguous segments numbered 133, 134 and 136 can be reused.

A further technique that we propose to ameliorate system performance is the introduction of priority for the delivery of video chunks from the parent node to the requesting peers. This solution implies that, when the peer moves to an adjacent overlay, its requests for video chunks are taken care of by its parent nodes with higher priority than the requests raised by nodes already resident within the swarm, so as to

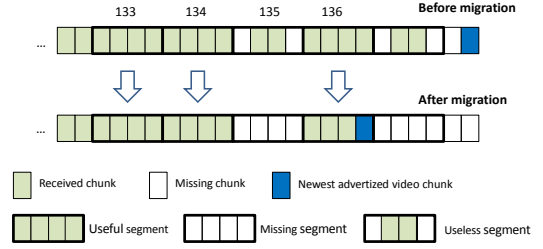


Fig. 4: Example of missing fragments

favor a swift acquisition of video content. This privileged condition lasts for an adequate, although finite, time interval, which is again equal to the time needed to accumulate S seconds of consecutive video. Accordingly, in every overlay chunk requests are categorized as either low or high priority; moreover, a low priority chunk request being received while the parent node is busy is lost, whereas some limited waiting is allowed for high priority requests, that are lost only when the corresponding queue, i.e., buffer, is full. In order not to excessively penalize the video distribution process towards nodes already present within the overlay, the size of such queue is limited, its value depending on the upload capacity of the distributing node and on the size in bytes of the video chunks being spread in the overlay. Indicating by $Q(l, j)$ the queue length of the generic parent node belonging to class l , with upload capacity c_l , within overlay j , and by t_{req} s the periodicity by which the nodes issue new chunk requests, we set

$$Q(l, j) = \frac{t_{req}}{c_l \cdot r_j}, \quad \forall l = 1 \dots C, \quad \forall j = 1 \dots K, \quad (6)$$

so as to guarantee that if a video chunk is queued, it will be sent to the requesting peer at most after t_{req} s. For this feature too, in next Section we will numerically demonstrate what effects it introduces on system behavior.

D. Diverse Requirements for Diverse Peer Classes and Overlays: the Modified Algorithm

As our interest is in devising mechanisms that allow DASH a fluid playout, we finally conceive a change to the core algorithm so as to guarantee an even sleeker overlay transition. The idea is to force the streaming process within the different overlays to be as much as possible aligned: the transition will therefore be facilitated, since a larger portion of the peer buffer content can be advantageously reused. A way to achieve this is to harden the conditions that rule the peer transition to an overlay that distributes video at a higher rate than the peer’s capacity, and to make such conditions stiffer when the rate of the target overlay increases. Accordingly, the original algorithm Alg.1 is modified as detailed in Alg.2, where we introduce a matrix of resource index thresholds whose generic element is $\sigma_{thres}(l_i, j)$, l_i being the class that node i belongs to and j the index of the overlay, obeying the following rules: $\sigma_{thres}(l_i, j) > \sigma_{thres}(l_h, j)$, under the hypothesis that peer i belongs to a class whose upload capacity is lower than the

Algorithm 2 Modified Switching Control Algorithm

```

Node  $i$  in overlay  $j$  every  $\Delta t$  seconds
;verifies its satisfaction
if (  $r_j < r_d(i)$  ) then
;verifies the current overlay status
if ( (  $\sigma_j < \sigma_{thres}(k_i, j)$  ) and (  $c_i \geq r_j$  ) ) then
do not migrate to overlay  $j + 1$ ;
else
;verifies the destination overlay status
if ( (  $c_i > r_{j+1}$  ) or (  $\sigma_{j+1} > \sigma_{thres}(k_i, j + 1)$  ) and
(  $E_{j+1} > E_{thres}$  ) ) then
migrate to overlay  $j + 1$ ; exit;
end if
end if
end if
;verifies its viewing quality
if ( (  $DR_i^{(t)} < DR_{thres}$  ) and (  $RWS_i^{(t)} < RWS_{thres}$  ) )
then
migrate to overlay  $j - 1$ ; exit;
else
stay in overlay  $j$ ; exit;
end if

```

class peer h belongs to, and $\sigma_{thres}(l_i, j + 1) > \sigma_{thres}(l_i, j)$, recalling overlay $j + 1$ distributes video at a higher rate than overlay j . When compared to the original algorithm, the new approach forces weaker peers to reside within the overlays that stream the video at the lower rates, hence in some cases it prevents them to reach the desired swarm. Yet, we will numerically provide evidence that the overall performance they experience is more rewarding, which is the final result every system should strive to achieve.

E. Merit Figures

To assess the performance of the proposed architecture and gain a thorough understanding of the effects that its different features bring in, various merit figures have to be considered. In detail, we have focused upon and evaluated the following indicators:

- the Delivery Ratio, DR, experienced by system users within the different overlays, that indicates how efficiently the streaming process takes place;
- the playback delay, defined as the interval that elapses from the point in time when the video chunk is generated by the streaming server up to the instant when it is received by the viewing user, so as to gain an insight on the time lag that the P2P distribution inevitably introduces, and also to verify to what extent the solution is able to fulfill the users' requests of near-real time services;
- the distribution of the users within the overlays, to understand how closely their original streaming rate requests can be satisfied;
- finally and most importantly, the switching delay incurred by the generic user when it moves from the current overlay to a new one, an action which DASH adoption inherently favors.

V. NUMERICAL RESULTS

A. Simulation Setup

The behavior of the proposed system is investigated with the help of an event-driven simulator, already employed in [12]. The simulator builds a replica of a multi-overlay P2P system, where every swarm distributes one DASH representation. The system is populated by $N = 2000$ active peers, that dynamically enter and leave; nodes join the system within the first 20 s and in this interval their interarrival times are exponentially distributed with an average of 0.1 s; after the first 20 s, the interarrival times are modulated so as to keep nearly constant the number of peers in the system. Session times are exponentially distributed, with an average duration of 1500 s. Every simulation lasts for 3000 s. Nodes belong to $C = 4$ different classes whose upload and download capacity values stem from the current European Internet connection offerings [22] [23] [24]; the percentages of users belonging to the different classes are drawn from the Akamai European average connection speed report [25]; the employed values are reported in Table II. $K = 4$ video representations are available at rates 700, 1500, 2500 and 3500 kbit/s; such values have been chosen having in mind the typical streaming rates of Internet Standard Definition video (SD) and High Definition video (HD). The streaming server allocates to each overlay only a small amount of its upload capacity, equal to four times the rate of the video, i.e., $C_{Sj} = 4 \cdot r_j$, indicating that the focus is on a pure P2P system. Moreover, the size of the current request window that every peer works with is $X = 20$ s, the request interval for video chunks is set to $t_{req} = 0.8$ s, buffer maps are sent out every $t_{BM} = 1$ s, the segment duration is 2 s and every segment is divided into $n = 10$ chunks. Video chunks are then fragmented and carried inside packets whose payload is 1250 bytes, transmitted over the underlying wide area network, where they suffer latencies that are captured via a suitable delay matrix: each pair of nodes in the overlay is randomly mapped to a pair of elements of the latency matrix, with an average end-to-end delay of 79 ms. Further, the delivery ratio is locally computed every $\tilde{t} = 5$ s and the threshold values we employ in the core and the modified algorithm are the following: $E_{thres} = 0.9$, $DR_{thres} = 0.55$, $RW_{thres} = 0.3$; we have picked these values after a careful tuning, where we verified that they guarantee the optimal system performance, that is, the lowest playback delay and the highest delivery ratio [20]. The coefficients for the computation of the weighted average of DR_i and RWS_i are $w_D = \frac{1}{3}$ and $w_W = \frac{2}{3}$, respectively, revealing that, with regard to the local delivery ratio we privilege stability in the estimate in (4), whereas for the request window state in (5) we give priority to what happens at current time. Last, the periodicity that the rate control algorithm employs is $\Delta t = 4$ s.

We investigate system behavior when every peer aims at streaming the video at the best possible representation, i.e., at the highest bit rate which results lower than the peer download capacity: in this circumstance, all peers aim at streaming representation 4, except for class 1 peers, that are confined to representation 2. Within this scenario it is impossible to fulfill

	Class 1	Class 2	Class 3	Class 4
Upload capacity (kbit/s)	704	1024	1500	10000
Download capacity (kbit/s)	2048	8192	10000	50000
% of peers	20	21	42	17

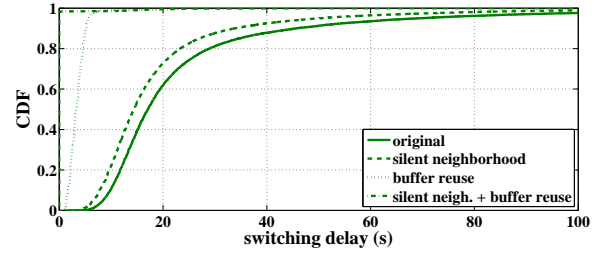
TABLE II: User capacity profiles and percentages

all users' requests: if peers were uncritically placed within the overlays distributing the desired representations (2 and 4), the aggregate upload capacity of each swarm would not be enough to satisfyingly deliver all peers the video, the resource index values being markedly lower than 1, namely, $\sigma_2 = 0.48$ and $\sigma_4 = 0.91$.

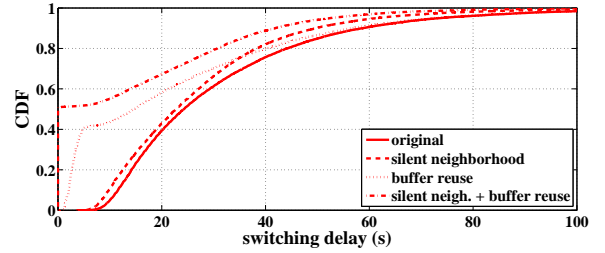
To estimate the switching delay introduced in Subsection IV-C, we set $S = 8$ s, that we deem a reasonable duration (similar values are encountered in literature when estimating the playout delay).

B. Silent Neighborhood and Buffer Reuse

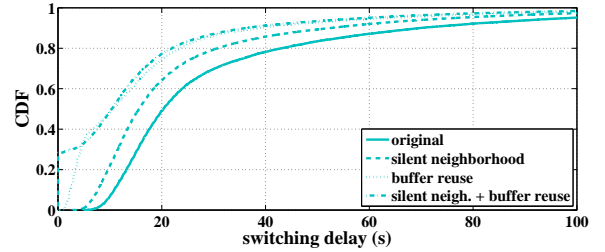
Figs.5(a)-(c) show the cumulative distribution function (CDF) of the switching delay for the system that leverages either the silent neighborhood or the buffer reuse mechanism, or both, and for a reference system, that we call the original system, that does not employ them. The figures reveal that such features statistically improve the switching delay, although in a different manner; moreover, these actions are significant for all overlays, except for overlay 1, where the majority of inputs are caused by the arrival of new peers in the system, which explains why the corresponding results are omitted. When the system employs the silent neighborhood mechanism only, the improvement consists in reducing to zero the setup time: in the original system this time is modeled as a random variable, uniformly distributed between 500 and 4000 ms. Hence the gain is moderate, and the horizontal lag between the corresponding curves (solid vs dashed lines) is of the order of magnitude of its average, equal to 2250 ms. We also observe that the combined use of the silent neighborhood mechanism and the buffer reuse (dash-dotted lines) leads to the best results: the former action allows the nodes to immediately know which complete segments can be reused, on the basis of the streaming position within the migration overlay; indeed, every silent neighbor advertises what video chunks it owns, so that a migrating node already knows the newest video chunk distributed within the target overlay and immediately proceeds to request the missing video chunks. As an example, Fig.5(a) referring to overlay 2 indicates that a node in the original system (solid line) experiences a switching delay equal to or greater than 25 s with probability 0.8, while in the system that introduces the buffer reuse only (dotted line) the delay decreases to 4.5 s and it is 0 s for the system that jointly combines the two actions (dash-dotted line). The exclusive introduction of the buffer content reuse (dotted lines) allows to reduce the switching delay in a different way depending on the considered overlay, namely, the advantage decreases when the distributed video streaming rate increases, mainly due to the mismatch between the streaming point of adjacent overlays. The fact that overlay 3 is a transit overlay, where there are not



(a) overlay 2



(b) overlay 3



(c) overlay 4

Fig. 5: Switching delay CDF

stable nodes, contributes to misaligning the streaming process between overlay 2 and overlay 3, as well as between overlay 3 and overlay 4. Notably, this is observed in Fig.5(c), that refers to overlay 4: here, to know in advance how many and which DASH segments the node can reuse leads to a minimal benefit, since there is a high probability that the number of inheritable segments is very low.

The silent neighborhood mechanism inevitably introduces some overhead, as a generic node has to construct and manage three neighborhoods in place of one, and several control packets are necessary to connect, disconnect, replace the silent neighbors, as well as to exchange buffer maps. So, in addition to quantifying the significant benefits this mechanism introduces when employed in conjunction with the buffer reuse, it is also mandatory to assess the price to be paid. To this aim, we define the control overhead as the ratio between the control bytes and the overall bytes sent by a node calculated with a time periodicity of $\Delta = 5$ s. Figs.6(a)-(d) plot the overhead averaged over all nodes in each overlay as a function of the simulation time. In the original system (solid lines) the control overhead decreases from overlay 1 to overlay 4: this had to be expected, as the size of the control packets is the same in all overlays, whereas the size of the video chunks increases in the overlays that distribute the video at higher streaming rates, as

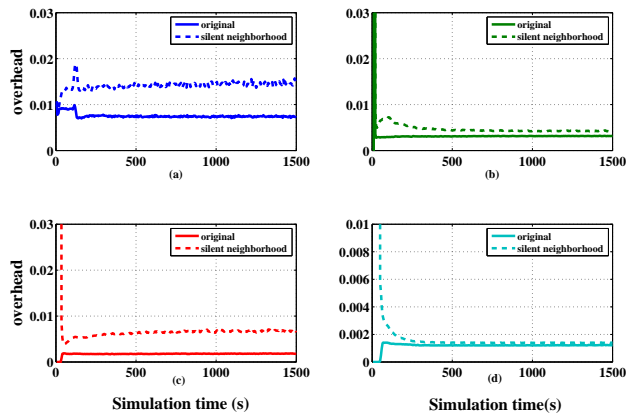


Fig. 6: Control overhead with and without silent neighborhoods: (a) overlay 1; (b) overlay 2; (c) overlay 3; (d) overlay 4

indicated in (1). When the silent neighborhood is introduced, the overhead increases, in particular for those overlays where the joining events take place more frequently, i.e., for overlay 1, as all nodes joining the system enter from here, and for overlay 3, an overlay where a node transits toward overlay 4, but where it often comes back if its viewing quality is not satisfying. Even though the values of overhead double and triple for overlay 1 and for overlay 3, in the worst cases they are confined to about 1 – 2%, so we can safely affirm that the percent increase of control packets due to the silent neighborhood management is abundantly counterbalanced by the switching delay reduction.

C. Priority Queuing

This subsection reports the results obtained when the scheduling policy adopted by the parent nodes to provide video content to the requesting peers is modified, in order to assign higher priority of retrieving video chunks to those users that have just entered the overlay and need to fill up their streaming buffer, to either start or continue the playout at a different rate. As regards the $Q(l, j)$ values defined in (6) that have been employed, some numerical examples are reported next: for a class 4 node with upload capacity $c_4 = 10000$ kbit/s belonging to overlay 1 that distributes video at $r_1 = 700$ kbit/s, the queue length is $Q(4, 1) = 57$, whereas if a node of the same class belongs to overlay 4 that distributes video at $r_4 = 3500$ kbit/s, its queue length is $Q(4, 4) = 11$.

Figs.7 (a) through (c) show the CDF of the switching delay experienced by the peers in overlays 2 through 4; as before, the viewgraph referring to overlay 1 is omitted. The figures allow to compare the original system (solid line) to the architecture that introduces the priority policy (dashed line), to the system where the silent neighborhood and the buffer reuse techniques are introduced (dash-dotted line), and finally to the one that jointly combines the priority mechanism to the silent neighborhood management and buffer reuse. Comparing the original system to the one that introduces the priority scheduling policy, from Fig.7(a) referring to overlay

2 we observe that a generic node within the latter system accumulates 8 s of video in 12.3 s or fewer seconds with probability 0.8, against the 26.6 s required by a peer in the original system. Similar gains are observable for overlay 3 and 4, as Figs.7(b)-(c) show. When examining the system that combines the priority scheduling policy with the silent neighborhood management and the buffer reuse, we conclude that here too a statistical improvement of the switching delay is achieved, as we note comparing the dotted and the solid lines in Fig7(a)-(c). The delay reduction is different for the three overlays: it is huge in overlay 2, where with probability 0.9 the switching delay becomes 0 when the buffer reuse strategy is introduced, and still very appreciable for overlays 3 and 4. However, the combined adoption of priority and neighborhood plus the buffer reuse mechanism has a contrasting effect:

- for low values of switching delay, the curves that refer to the system with no priority (dash-dotted lines) are better;
- for high values of switching delay, the curves that refer to the system with priority scheduling (dotted lines) are better.

With reference to the first remark, we anticipate that the priority scheduling introduction to some degree deteriorates the streaming process in overlay 3, which results more misaligned with respect to the adjacent overlays: numerical evidence of this statement will be provided by next figures. Any misalignment implies a lower efficiency of the silent neighborhood and buffer reuse techniques, since there are not so many useful DASH segments to inherit; for this reason we assist to a decrease in the probability of a generic peer to experiment no switching delay at all and to an initial worsening of the CDF curves. This effect is predominant for low values of the switching delay, that is, for the nodes that migrate with an almost full buffer and take a little time to fill it. However, we can also state that the adoption of the priority scheduling accelerates the collection of the video chunks needed to fill up the buffer, so the slope of the curves of the system with priority (dotted lines) is steeper than that of the system with the silent neighborhood and the buffer reuse (dash-dotted lines). This effect is more evident for high values of the switching delay, that is, for the nodes that join the new overlay with an almost empty buffer and that cannot benefit of the buffer reuse effect. Moreover, these figures indicate that the system with priority guarantees lower values of the maximum switching delay.

We next complete the investigation on the effects of priority scheduling on system performance and determine the average delivery ratio and the CDF of the playback delay. Fig.8 reports the average delivery ratio of each overlay for the system that introduces the priority scheduling and compare it to the original architecture. The values of delivery ratio have been calculated as averages over the total number of peers and also over the last 1000 s of simulation, when the system exhibits a stable behavior, for 10 different simulation runs. We notice that only the delivery ratio of overlay 3 decreases, whereas the remaining values in the two systems are comparable. This occurs because overlay 3 is a poorly populated overlay, subject to frequent join events and departures, as no peer actually wants to reside in it: the few “older” nodes are busy serving

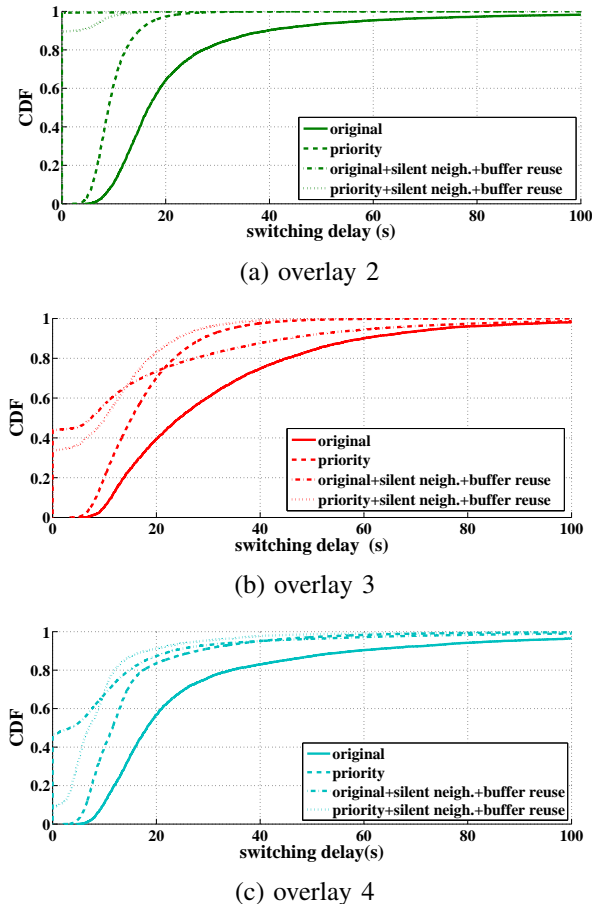


Fig. 7: Switching delay CDF: (a) overlay 2; (b) overlay 3; (c) overlay 4

the numerous newcomers with higher priority, at the expense of the video diffusion process towards its old peers.

Fig.9 reports the CDF of the playback delay experienced by a generic node in the different overlays, and as anticipated before, it reveals that the spreading process in the “odd” overlay moderately worsens, due to the introduction of the priority scheduling policy, and that the same holds for overlay 1, the system entry point. Overall, it is noteworthy to remark that the priority policy applied in isolation represents an effective tool to reduce the switching delay; when combined with other features, it nevertheless guarantees to confine such delay to modest values with high probability.

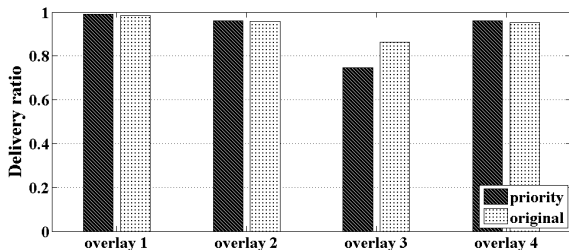


Fig. 8: Average delivery ratio for the original system and for the system with priority scheduling

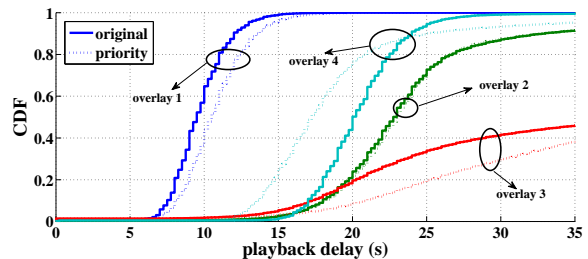


Fig. 9: Playback delay CDF for the original system and for the system with priority scheduling

D. Modified Control Algorithm

As an alternative to the previous system with all its possible variations, we next investigate the behavior of the modified algorithm introduced in subsection IV-D. To this aim we employ the matrix given below; as required for the algorithm to be effective, the $\sigma_{thres}(l, j)$ values increase or do not vary moving from left to right and from top to bottom.

$$\sigma_{thres}(l, j) = \begin{bmatrix} 1.0 & 1.2 & 1.6 & 1.8 \\ 1.0 & 1.1 & 1.4 & 1.6 \\ 1.0 & 1.0 & 1.2 & 1.4 \\ 1.0 & 1.0 & 1.0 & 1.0 \end{bmatrix}$$

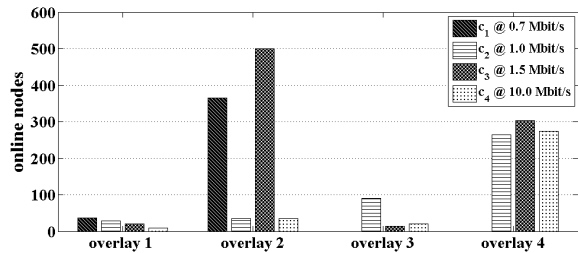
The adoption of the new algorithm has an immediate impact on the average resource index of each overlay, as it is shown in Table III, as well as on the distribution of nodes among the various overlays, as reported in Figs.10(a)-(b). Table III highlights that, as desired, more resources are redirected towards overlays 3 and 4, those that distribute the video at higher streaming rates, at the expense of overlay 1.

	overlay 1	overlay 2	overlay 3	overlay 4
original system	2.5	1.0	1.0	1.1
σ_{thres} system	1.4	1.0	1.4	1.5

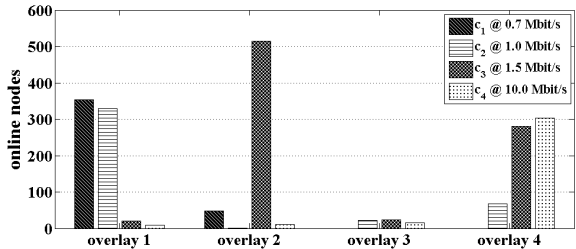
TABLE III: Average values of the resource index

Figs.10(a)-(b), that report the node distribution in the original system and in the system with differentiated σ_{thres} thresholds, reaffirm that the modified algorithm tends to place the weaker nodes within the overlays that distribute the video at a streaming rate lower than or equal to their upload capacity. Here too, the reported values have been obtained as averages over the last 1000 s of the simulation time and over 10 runs. The introduction of tighter thresholds limits the possibility of weaker nodes to reach the overlays that distribute the higher streaming rates, even if they desire it; at the same time it guarantees a better content diffusion process, i.e., higher values of delivery ratio and lower playback delays, as the following table and figures will show.

Figs.11(a)-(b) allow to compare the CDF of the playback delay of the two systems: we note that the system that employs differentiated σ_{thres} thresholds is able to guarantee closer values of playback delay among overlays, which equivalently means more aligned streaming processes, exactly as desired. Moreover, all the curves improve (i.e., move to left) for the system with differentiated σ_{thres} thresholds, except for overlay



(a) Original system



(b) System with differentiated resource index thresholds

Fig. 10: Per class node distribution within each overlay

1, where a reduction of resources occurs, as indicated by Table III. To complete the picture, Fig.12 reports the Complementary Cumulative Distribution Function (CCDF) of the delivery ratio in the two compared systems, highlighting that the platform implementing the modified switching algorithm consistently achieves high values of delivery ratio with higher probability.

distance	original system	σ_{thres} system
0	60.4	35.0
1	8.1	20.7
2	28.6	26.4
3	2.9	17.9

TABLE IV: Percentage of nodes that stream a representation n levels away from the desired

Taking a different perspective, for both systems Table IV reports the percentage of nodes as a function of the distance between the overlay where the node actually resides and the overlay streaming the representation the node would like to view: $d = 0$ means that the peer is streaming the desired representation, whereas $d = 1, 2, 3$ indicates the peer is 1, 2, 3 overlay(s) away from its target, respectively. The modified system is able to place fewer peers than the original platform within the desired overlay; yet, Figs.11 and 12 indirectly suggest that the viewing quality the nodes experience is more remarkable, and this is the key advantage. In other words, it is better to stream at a lower bit rate with good quality, without noticing significant switching delays in the transition from one representation to another and experiencing good operating conditions, than to stream high rate videos in scarce conditions, as also suggested in [26].

To complete the investigation, Figs.13(a)-(c) report the switching delay CDF and indicate that we obtain a considerable improvement if we differentiate the σ_{thres} thresholds on the basis of the class the node belongs to and of the destination

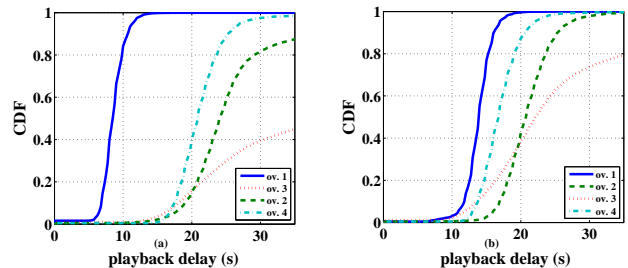


Fig. 11: Playback delay CDF: (a) original system; (b) system with differentiated resource index thresholds

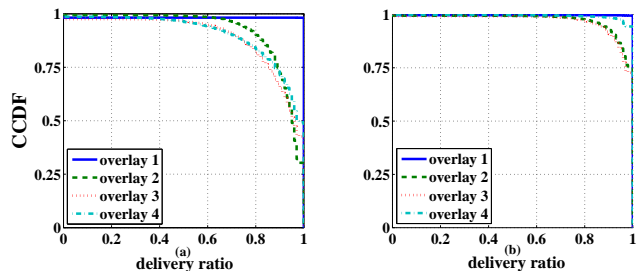


Fig. 12: Delivery ratio CCDF: (a) original system; (b) system with differentiated resource index thresholds

overlay (dashed lines opposed to solid lines). As a remarkable example, we observe that in overlay 4 for the probability value 0.8 the switching delay reduces from 36.7 s in the original system (solid lines) to 16.2 for the σ_{thres} system (dashed lines); in overlay 3 it reduces from 45.5 s to 20.3 and for overlay 2 the gain is approximately 10 s; only for overlay 1 the CDF of the switching delay slightly deteriorates, as resources are less abundant than in the original system.

With the help of the same figures, it is interesting to investigate how the priority scheduling policy and the silent neighborhood management with buffer reuse influence the system that employs differentiated σ_{thres} . Fig.13 reveals that the buffer reuse action (dotted lines) achieves a significant improvement in terms of switching delay, thanks to the greater alignment among overlays that the modified algorithm introduces. We report some numerical examples to emphasize the good results: with probability 0.8 a node within overlay 4 gathers 8 s of video chunks in 3.2 s, far less than the required 16.2 s if the buffer reuse is not active; and for overlays 2 and 3 with probability 0.8 a node experiences no delay at all. If in addition to the buffer reuse mechanism we also activate the newest priority scheduling policy (dash-dotted lines), the improvement is negligible (in Fig.13(c) it is just noticeable for high CDF values). This is due to the fact that if the video diffusion processes of the overlays are adequately aligned, the buffer reuse action allows to inherit a very high percentage of previously gathered chunks, reducing to a minimum the number of new chunks to collect, so that the priority scheduling policy has a very low impact.

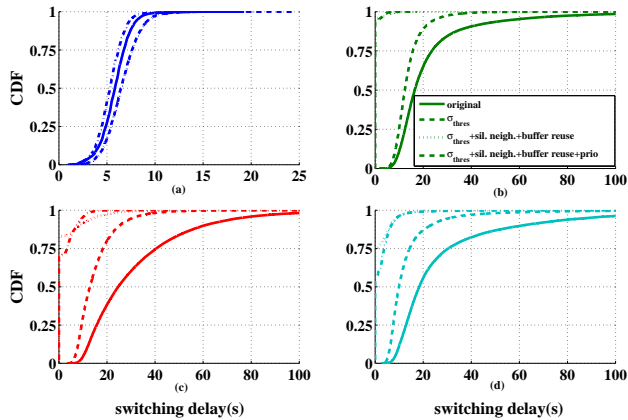


Fig. 13: Comparison between the switching delay CDFs: (a) overlay 1; (b) overlay 2; (c) overlay 3; (d) overlay 4

E. A Different Scenario

At last, we test the performance of the P2P-DASH platform equipped with all the proposed features, that is, augmented neighborhood, priority scheduling, buffer reuse and the modified control strategy, for a different profile of user capacities and requests. We consider higher upload capacities with respect to the previous scenario, as shown in Table V. Moreover, according to the bitrates employed by Netflix [21] to distribute SD, HD and full HD videos, we set the rates of the available representations at 1000, 2500, 4500 and 5800 kbit/s.

		Class 1	Class 2	Class 3	Class 4
Upload	capacity	1024	2048	3072	10000
Download	capacity	3072	8192	10000	50000
	% of peers	20	21	42	17

TABLE V: New user capacity profiles and percentage

Even in this new setting, the simulation results show that the P2P-DASH architecture grants very high values of delivery ratio and confined playback delays: Table VI indicates that all overlays exhibit average delivery ratios above 0.96; also the average playback delays are comparable to those previously computed.

	Average Delivery Ratio	Average Playback Delay (s)
Overlay 1	0.997	13.5
Overlay 2	0.995	21.9
Overlay 3	0.960	29.2
Overlay 4	0.989	26.7

TABLE VI: Average delivery ratio and playback delay values

Moreover, Fig.14 demonstrates that the switching delay is successfully confined too, confirming the validity of our proposal when diverse operating conditions are examined.

VI. CONCLUDING REMARKS

This work has put forth a P2P-DASH live streaming architecture that delivers different representations of the same video

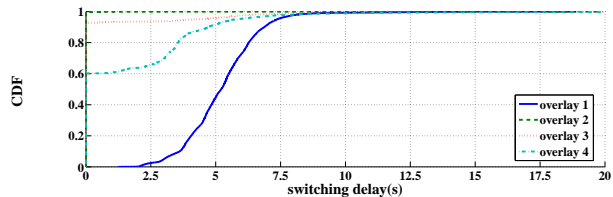


Fig. 14: Switching delay CDF

content, each being streamed by a single swarm of a multi-overlay platform. The study has been mainly centered on the reduction of the switching delay, a critical metric that markedly affects how smooth the streaming process is perceived by the user. The proposed system exhibits some peculiar features that aim at confining such delay and simultaneously guarantee a good functioning of the entire platform: there is a silent extension to the peer neighborhood; DASH segments that are completely received and stored within the peer streaming buffer at transition time are reused; the scheduling of video chunks assigns higher priority to requests raised by peers that newly join the system or migrate to a different overlay; both a core algorithm and a modified alternative are proposed, where the latter more conservatively allocates weaker peers to overlays that stream video at lower rates, therefore reducing their misalignment and improving the experienced performance.

System behavior is extensively investigated through simulation and the major findings are as follows: the benefits attained by the introduction of the silent neighborhood mechanism and the buffer content reuse are the more impressive, the more the streaming processes in adjacent overlays are aligned; priority scheduling is the only approach able to reduce the switching delay regardless of such misalignment; the modified thresholds introduced by the alternative control algorithm guarantee better alignment and a very good viewing quality. We demonstrated that when such rate control algorithm is employed in conjunction with the silent neighborhood and the buffer reuse mechanisms, an excellent performance is achieved, i.e., high values of delivery ratio, low values of playback and switching delays: however, a trade-off is inevitable, as the users that exhibit the lowest upload capacity stream the video - with good quality - at a reduced rate.

REFERENCES

- [1] "Global Internet Phenomena, Latin America and North America," Sandvine Inc., Canada, May 2015. [Online]. Available: <https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/global-internet-phenomena-report-latin-america-and-north-america.pdf>
- [2] "Global Internet Phenomena, Asia-Pacific and Europe," Sandvine Inc., Canada, May 2015. [Online]. Available: <https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/global-internet-phenomena-report-apac-and-europe.pdf>
- [3] *Information Technology - Dynamic Adaptive streaming over HTTP(DASH) - Part 1: Media presentation description and segment formats*, ISO/IEC Std. 23 009-1, 2014.
- [4] T. Stockhammer, "Dynamic adaptive streaming over HTTP - standards and design principles," in *Proc. ACM MMSys*, Feb. 2011, pp. 133-144.
- [5] "Survey of European Broadcasters on MPEG-DASH," Survey, DASH Industry Forum (DASH-IF), May 2013. [Online]. Available: <http://dashif.org/wp-content/uploads/2015/04/Survey-of-the-European-Broadcasters-on-MPEG-DASH-Whitepaper-V2.1.pdf>

- [6] Z. Li, X. Zhu, J. Gahm, R. Pan, H. Hu, A. Begen, and D. Oran, "Probe and adapt: Rate adaptation for HTTP video streaming at scale," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 4, pp. 719–733, Apr. 2014.
- [7] L. De Cicco, V. Caldaralo, V. Palmisano, and S. Mascolo, "ELASTIC: a client-side controller for dynamic adaptive streaming over HTTP (DASH)," in *Proc. IEEE 20th Int. Work. PV*, Dec 2013, pp. 1–8.
- [8] C. Liu, I. Bouazizi, and M. Gabbouj, "Rate adaptation for adaptive http streaming," in *Proc. ACM MMSys*, Feb. 2011, pp. 169–174.
- [9] R. K. P. Mok, X. Luo, E. W. W. Chan, and R. K. C. Chang, "QDASH: a QoE-aware DASH system," in *Proc. ACM MMSys*, Feb. 2012, pp. 11–22.
- [10] D. Vergados, A. Michalas, A. Sgora, and D. Vergados, "A fuzzy controller for rate adaptation in MPEG-DASH clients," in *Proc. IEEE 25th Int. Symp. PIMRC*, Sep. 2014, pp. 2008–2012.
- [11] Y. Shuai, G. Petrovic, and T. Herfet, "OLAC: an open-loop controller for low-latency adaptive video streaming," in *Proc. IEEE ICC*, Jun. 2015, pp. 6874–6879.
- [12] L. Natali and M. L. Merani, "A novel rate control strategy for adaptive video streaming in P2P overlays," in *Proc. IEEE Globecom*, Dec. 2015, (to appear).
- [13] S. Lederer, C. Muller, and C. Timmerer, "Towards peer-assisted dynamic adaptive streaming over HTTP," in *Proc. IEEE 19th Int. Work. PV*, May 2012, pp. 161–166.
- [14] L. Rovcanin and G.-M. Muntean, "A DASH-based performance-oriented adaptive video distribution solution," in *Proc. IEEE Int. Symp. BMSB*, Jun. 2013, pp. 1–7.
- [15] G. Tian, Y. Xu, Y. Liu, and K. Ross, "Mechanism design for dynamic P2P streaming," in *Proc. 13th IEEE P2P*, Sep. 2013, pp. 1–10.
- [16] A. Zambelli, "IIS smooth streaming technical overview," Microsoft Corporation, Tech. Rep., 2009.
- [17] L. Yitong, S. Yun, M. Yinian, L. Jing, L. Qi, and Y. Dacheng, "A study on quality of experience for adaptive streaming service," in *Proc. IEEE ICC*, Jun. 2013, pp. 682–686.
- [18] M. Zink, O. Künzel, J. Schmitt, and R. Steinmetz, "Subjective impression of variations in layer encoded videos," in *Proc. 11th IEEE-ACM IWQoS'03*, Jun. 2003, pp. 137–154.
- [19] D. Wu, C. Liang, Y. Liu, and K. Ross, "View-upload decoupling: A re-design of multi-channel P2P video systems," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 2726–2730.
- [20] L. Natali and M. L. Merani, "Adaptive Streaming in P2P Live Video Systems: A Distributed Rate Control Approach," arXiv preprint arXiv: 1510.04472.
- [21] Netflix. [Online]. Available: <http://www.netflix.com>
- [22] Swisscom. Swisscom internet offers. [Online]. Available: <http://www.swisscom.com>
- [23] Orange. Orange offres internet. [Online]. Available: <http://www.orange.fr>
- [24] T.Net. T.Net connettività e voce. [Online]. Available: <http://www.tnet.it>
- [25] Akamai, "State of the internet q2-2014," Tech. Rep. [Online]. Available: <http://www.akamai.com/stateoftheinternet/>
- [26] S. Medjiah, T. Ahmed, R. Boutaba, "Avoiding Quality Bottlenecks in P2P Adaptive Streaming," in *IEEE J. Sel. Areas Commun.*, vol. 32, no. 4, pp. 734–745, Apr. 2014.