

On the Impact of Playout Scheduling on the Performance of Peer-to-Peer Live Streaming

Constantinos Vassilakis^a, Nikolaos Laoutaris^b, and Ioannis Stavrakakis^a

^aDept of Informatics and Telecommunications, University of Athens, Greece;

^bTelefonica Research, Barcelona, Spain

ABSTRACT

In this paper we examine the impact of the adopted playout policy on the performance of P2P live streaming systems. We argue and demonstrate experimentally that (popular) playout policies which permit the divergence of the playout points of different nodes can deteriorate drastically the performance of P2P live streaming. Consequently, we argue in favor of keeping different playout points “near-in-time”, even if this requires sacrificing (dropping) some late frames that could otherwise be rendered (assuming no strict bidirectional interactivity requirements are in place). Such nearly synchronized playout policies create “positive correlation” with respect to the available frames at different playout buffers. Therefore, they increase the number of upstream relay nodes from which a node can pull frames and thus boost the playout quality of both single-parent (tree) and multiple-parent (mesh) systems. On the contrary, diverging playout points reduce the number of upstream parents that can offer a gapless relay of the stream. This is clearly undesirable and should be avoided as it contradicts the fundamental philosophy of P2P systems which is to supplement an original service point with as many additional ones presented by the very own users of the service.

1. INTRODUCTION

P2P streaming using tree or mesh overlays: Distributing a live video stream using a P2P streaming system has the advantage over a point-to-point client/server system of offering more resources to clients by effectively turning each one of them into a secondary server that assists in the distribution of the stream. These additional resources can yield improved scalability and/or resilience, depending on the design of the system. Organizing the nodes into a tree shoots for scalability by requiring only n overlay links, where n denotes the number of receivers. Having a minimal number of overlay links reduces the *stress* of the underlying physical links, i.e., the number of times that the same information can flow over the same physical link (this can happen as multiple overlay links may go through the same physical link). It also minimizes the amount of overlay link monitoring overhead for detecting congestion/churn and triggering handoffs in-time to avoid disruption of playout (assuming that the same amount of monitoring expenditure is paid at each link). These observations hold true for both single-tree/no-coding and multiple-tree/multiple-description-coding architectures.

Meshes on the other hand shoot for resilience to congestion/churn by providing each node with multiple parents from which it can receive the stream in parallel (using single-, multiple-description-, or network-coding). Since a mesh uses more than n overlay links, it can increase the stress of the underlying physical links and the monitoring overhead (although the latter is not necessarily true as a node can take advantage of the redundancy offered by having multiple parents and performing a more lazy monitoring of each incoming links). The discussion of which combination of topology and encoding

Email addresses: cvassilakis@noc.uoa.gr (C. Vassilakis), nikos@tid.es (N. Laoutaris), ioannis@di.uoa.gr (I. Stavrakakis).

is the “right-one” has been going-on for some time, and although it seems that recent mesh-based systems using coding have several advantages,¹ it all depends in the end on the assumed operating environment and the desired cost/complexity of building and maintaining the system: well-behaving environments (e.g., dedicated cable networks) can benefit from the simplicity/economy offered by tree-based distribution; uncontrolled/variable environments (e.g., under-provisioned parts of the current Internet) can benefit from the redundancy offered by mesh topologies and coding.

The (new) role of playout scheduling in P2P: The playout scheduler is the component of a video receiver that handles the buffering and rendering of received frames. Designing appropriate playout schedulers for video streaming application was one of the central research topics of the multimedia transmission community up to the emergence of P2P streaming systems, at which point the focus shifted onto overlay construction and coding issues. Although fairly well understood in the context of point-to-point video streaming,² playout scheduling has received a rather limited attention in the context of P2P video streaming.³ The new setting, however, perplexes playout scheduling beyond our previous understanding. In addition to achieving the desired tradeoff between interactivity and stream continuity, the playout scheduler must now jointly factor-in that different playout processes become coupled in the context of P2P: buffering, rendering, or dropping a frame affects not only the local process but also downstream ones that might connect and request frames from the local node. In this article we argue that although seemingly subtle compared to topology construction and coding, playout scheduling still deserves some attention as a bad choice with respect to it can impact quite negatively the performance of P2P streaming systems, despite the existence of the other two powerful enablers.

Our contribution: Locating upstream parents and content delivery path maintenance are recognized as two major challenges in the area of P2P streaming.⁴ In this work we argue and demonstrate that the adoption of a playout policy has an impact on the performance since it directly affects the availability of upstream parents.

We consider a *delay preserving* playout policy called *Sync* and a *data preserving* one called *Async* in the context of a P2P streaming system. These two policies lay at the extremes of the spectrum of studied playout policies.² Under *Sync*, the playout scheduler enforces a fixed predefined time offset between the time that a frame is presented at a receiver and the time it was captured at the encoder. To do so, it has to drop “late” frames that arrive after their scheduled playout time, even if they are eventually received correctly and in their entirety. The data preserving *Async* policy on the other hand, imposes an initial buffering delay and then presents frames by draining the buffer at a constant rate. In the event of a buffer underflow, the playout freezes and resumes again upon the reception of the next frame. Not dropping late frames makes the offset between encoding and decoding times variable. In fact, in the absence of losses in the network, the offset increases with each underflow by an amount equal to the duration of the underflow.

We operate each one of these playout policies in a P2P streaming system with the following characteristics: (1) hierarchical structure, (2) threshold-based handoffs (change of upstream parent) based on partial or full information on the remaining network, (3) single-description coding. Such a setting resembles initial P2P streaming systems as the one presented in⁵ and was chosen due to the popularity of such systems, their simplicity, and most importantly, in order to protect our evaluation of playout from issues that are orthogonal to it. In a sense, our chosen setting is the most fragile one as it includes a minimum amount of redundancy. Certainly one can design an over-provisioned system based on a dense overlay graph with multiple reception points and elaborate coding, but this would obscure the effects of playout policy which is what we want to isolate in this work.

We develop a simulation environment for the above policies and setting and use it to compare them across different levels of network load and heterogeneity with respect to link capacities. Our evaluation is based on “direct” metrics such

as *Discontinuity*, which captures the percentage of time a user spends viewing some frozen frame, and *Loss*, which captures the percentage of content never presented to a user (both metrics are defined precisely later). To explain the observed results on these metrics we introduce a new “indirect” one – called *Availability* – which roughly amounts to the number of available upstream parents to which a node can perform a smooth handoff at a time of poor reception quality from its current parent. Based on several simulation scenarios for our control variables (load, heterogeneity, information on remote nodes, number of past frames kept) we arrive at the following main observations and conclusions:

- Sync performs consistently better than Async with respect to both *Discontinuity* and *Loss* under a wide spectrum of load and heterogeneity. The improved performance can be explained by the fact that Sync maintains higher *Availability* and thus is able to perform smooth handoffs at times of poor reception. Under Async, the underflows contribute to the time divergence of playout points and the de-correlation of buffer contents. Thus when a node seeks a handoff it becomes difficult to find a parent with the missing frames for a smooth transition.
- Sync is effective even under limited knowledge of remote nodes (used for performing handoffs). Having the playout nodes nearly synchronized means that any one of them can offer more or less the missing frames, so we don’t need to have a global view of buffer contents – tracking a small set of alternative parents suffices for handoff operations. Async on the other hand needs to know the buffer contents of remote nodes so as to identify the one (if any) whose playout point is at the right distance for a gapless handoff.
- Similarly, Sync is relatively immune to constraints on the number of downstream nodes that a parent can support. Having the playout points of different nodes near in time creates a natural load-balancing with respect to the handoffs because all nodes hold approximately the same frames and are equally good from the standpoint of a node seeking for a new parent. Contrary to this, in Async there are many cases where few nodes exist that are at the “correct” time distance from many other nodes, but cannot accommodate all of them due to these constraints and thus the seeking nodes are forced to perform handoffs that induce gaps in playout.
- Although rather counter intuitive, Async’s performance is favored by randomness in parent selection (imposed by restrictions such as the ones described above) since the latter eventually assists in keeping playout points “near-in-time”; peers are forced not to diverge a lot by performing handoffs that induce loss and thus restore, up to a point, their offset.
- Unlike Sync, Async can benefit from keeping frames in the buffer even after they have been displayed locally. This, however, leads to several known complications (how much of it is needed to smooth out the disruption without making the offset exceedingly large) as well as some new ones (copyright restrictions permit the nodes of P2P streaming systems to buffer only a limited time window of copyright protected material.⁶

All the above indicate that Sync is a better option for the considered P2P systems. At the core of its advantage is that it is conforming to the P2P character of the application. Async on the other hand, by virtue of the divergence that it fosters, goes against the P2P paradigm by effectively reducing the number of secondary service points that are available to a node.

The remainder of the article is structured as follows. In Sect. 2 we review related work from the literature. In Sect. 3 we present the details of the considered playout policies and P2P setting. Several simulation scenarios on multiple metrics and control parameters are presented in Sect. 4. Finally, Sect. 5 summarizes our findings, discusses the applicability of the results to multiple-parent systems and concludes the article.

2. RELATED WORK

Several application-layer multicast systems have been proposed for addressing the low deployment of network-layer multicast. Initial application-layer multicast systems mimicked network-layer multicast and thus adopted a single tree topol-

ogy⁷⁻⁹ aiming at providing a similar performance to it with respect to stress and stretch (the stress metric is defined per-link and counts the number of identical packets sent by a protocol over each underlying link in the network, the stretch metric is defined per pair of nodes and captures the ratio of path-length over the overlay to the corresponding path-length over unicast IP).

The first wave of improvements to these systems aimed at addressing the unreliability of end-nodes and at decreasing the control overhead. The Bullet system¹⁰ proposed splitting a stream into multiple blocks and delivering disjoint subsets of these blocks to different nodes over an overlay tree. It then let the nodes search for “missing blocks” and download them from other nodes using additional mesh links. The Zigzag system¹¹ aimed at reducing the control overhead by first clustering peers and then building a multicast tree on top the formed clusters.

Additional improvements aimed at balancing the forwarding load and leveraging bandwidth heterogeneity. SplitStream¹² splits a stream into multiple stripes and sends each one over a different multicast tree. Load balancing in terms of forwarding is achieved by making a node internal into one tree and leaf in all the others. CoopNet¹³ proposed using multiple-description-coding^{*} and transmitting each description layer over a different tree so as to allow heterogeneous peers to join trees according to their bandwidth capacity. ChunkySpread¹⁴ adopted ideas from SplitStream but used an unstructured approach for building the trees whereas SplitStream is based on the Pastry DHT.¹⁵

Proposals like CoolStreaming¹⁶ and PRIME¹⁷ have abandoned trees in favor of BitTorrent-like¹⁸ transmission based on swarming on top of a mesh topology. In these systems the stream is broken into different blocks and nodes obtain such blocks from multiple senders. “Buffer maps” are used for advertising the availability of blocks at each node.

All these works focus on overlay construction and coding but do not look at the details of playout scheduling. Tree-based schemes perform blind forwarding (“push”) of all received content and thus also send out frames which are not needed by playout processes in which these frames have expired. Mesh-based systems generally “pull” frames from the senders, but most of them assume that receivers buffer all the incoming frames even after they have been presented (which is not generally possible as discussed in the introduction). PRIME seems to be the first BitTorrent-like mesh system that integrates into its swarming distribution the real time constraints of live video streams. Contrary to our work, however, PRIME assumes a deterministic delivery of frames resulting from (1) the assumption that there is no jitter at the backbone and (2) manipulating in- and out-degrees of nodes so as to avoid congestion due to access-link bandwidth constraints. In our work we examine playout scheduling in jitter-prone environments and “fragile” (single) tree topologies. In doing so we want to substantiate the maximum damage from using the wrong playout policy in such a setting.

We are aware of only two works directly related to ours. In¹⁹ the authors state that for gapless playout, peer selection should not only be done based on network quality criteria, but also on the buffer status of the candidate parent peer, effectively recognizing the phenomenon of “negative correlation”. However they do not associate this phenomenon with the level of synchronization resulting from different playout schedulers, which is the main contribution of our work. In²⁰ different receivers achieve different synchronization levels with the source as a result of the initial prefetching mechanism. In order to improve a peer’s “liveness”, the playout rate is slightly altered while a parent and a client peer may switch roles if the selected parent is behind in playback to facilitate catch up of the late peer. The connection between the synchronization of different receivers and their ability to cooperate by serving missing frames (what we call “availability” here) is not explored.

^{*}A layered coding scheme in which each layer/description/substream is independently decodable and full stream quality amounts to obtaining all the layers.

3. SYSTEM DESCRIPTION

In this section we present the details of the various components of our evaluation. It must be obvious by now that our purpose is not to introduce a novel delivery system but to adopt a system setting that resembles initial P2P streaming systems which will allow us to isolate and explain the effect of playout scheduling on the performance without any hidden interferences from the other system components. Following, we start with the two playout policies and move on to the details of the single tree hierarchical P2P streaming system which we consider.

3.1 Playout Policies

Let V be the set of all peers in the system. Let $e(k)$ denote the encoding time for the k th frame and $p_i(k)$ be its scheduled playout time at node v_i . We define the following playout schemes:

Sync(D_i): Frames that become available at peer v_i before their scheduled playout time are displayed at their exact playout time p_i . Frames that miss their playout time are skipped. This amounts to synchronous playout between the source and node v_i where by *synchronous* we indicate a *fixed offset* D_i between encoding and playout times. That is:

$$p_i(k) = e(k) + D_i$$

When $D_i = D, \forall v_i \in V$, all nodes display the same frame at the exact same time and a global synchronization is achieved.

Async(D_i): After an initial buffering delay D_i for the first frame, subsequent frames get displayed at the earliest possible time following their previously displayed one. Assuming that no frames are lost in the network, we can define Async recursively as follows:

$$\begin{aligned} p_i(k) &= p_i(k-1) + T + U_i(k-1) \\ p_i(1) &= e(1) + D_i \end{aligned}$$

where T is the nominal duration of a frame and $U_i(k-1)$ is the duration of a possible underflow that follows the presentation of frame $k-1$ at peer v_i . If frame k is readily available after the presentation of frame $k-1$ then $U_i(k-1) = 0$.

Sync and Async stand at the two extremes of the spectrum of playout policies from delay to data preserving. Of course there exist intermediate policies in this spectrum, e.g., those that apply modified playout rates depending on the current buffer occupancy,²¹ but these are rather elaborate and fall outside the scope of the current article.

3.2 Initial Tree Build-Up

We assume that nodes form a single hierarchy rooted at the video source which transmits a single-description stream. A new peer v_i selects randomly a parent peer v_j already in the system and connects to it (we discuss reconnecting peers and handoffs later). Node v_i selects a specific frame from v_j 's playout buffer and starts prefetching it and all subsequent ones for a time interval F_{ij} which leads to the desired offset D_i between the playout of this first received frame at v_i and its encoding time at the source. We describe how to select this frame and the prefetching period F_{ij} in Appendix A. At the end of the prefetching period, a number of \hat{k}_i frames have been received and stored into peer's v_i playout buffer and the playout process starts. It is at this point that the differences between Sync and Async start to materialize – the first one maintains this initial offset whereas the second one lets it increase by accepting and presenting late frames.

3.3 Performing Handoffs

We allow a node v_i to be in either of the following two modes:

Stable mode: A node is stably connected to its parent as long as its current buffer occupancy b_i is above a threshold value B_h and its parent hasn't left the distribution tree.

Handoff mode: A node enters a handoff mode as soon as its buffer occupancy falls beneath B_h or it is abandoned by its parent. The handoff amounts to selecting a new parent and connecting to it for a grace period T_g before returning to a stable mode. The grace period allows for buffer build up thus avoiding cascading handoffs.

Pre-active handoffs are employed to increase the chance for gap-free transitions when the connection to the current parent peer is not good enough, or when the latter has left the system. To perform such handoffs node v_i is supplied with a random subset $V_i \subseteq V : |V_i| = m \leq n$ (n is the number of all peers in the system) of potential parents which it keeps monitoring while in stable mode. Monitoring amounts to exchanging periodic signalling messages with each node $v \in V_i$ containing the identity of the frame currently on display at v , as well as the newest frame in v 's playout buffer. To perform the handoff, v_i partitions V_i into three disjoint subsets V_i^A , V_i^B , and V_i^C , and connects to a random node starting from V_i^A , continuing with V_i^B if $V_i^A = \emptyset$, and with V_i^C if both V_i^A and V_i^B are empty. The three sets are defined as follows:

- V_i^A includes the known peers that have in their buffers the next missing frame for v_i , i.e., the one whose id is higher by 1 from the id of the frame that is on the top of the playout buffer of v_i (position b_i).
- V_i^B includes the known peers that don't have the next missing frame for v_i , but will have it in the future as their playout point hasn't exceeded it yet.
- V_i^C includes peers in $V_i \setminus (V_i^A \cup V_i^B)$ that hold any frames that v_i doesn't have.

Peers are selected randomly within the subsets for load balancing. The handoff is gapless only when it is done towards a node of V_i^A . In Appendix B we present the details of partitioning the nodes of V_i into the three subsets based on the received signalling messages from them.

4. PERFORMANCE EVALUATION

4.1 Metrics

We compare Sync(D) and Async(D) based on the following performance metrics:

Discontinuity: A discontinuity occurs when due to the unavailability of the next fresh frame(s), the last in-time rendered frame remains on display longer than its nominal time. Under Sync, the discontinuity increases by T with each frame that misses its scheduled playout time. Under Async, the discontinuity increases with each underflow, by an amount that equals the duration of the underflow. Let d denote the average discontinuity ratio expressed as the average among all peers of the total time that a peer spends viewing frozen frames to the total playback time.

Let \tilde{T}_i be the total viewing time or the effective duration of the stream at peer v_i . Let \tilde{K}_i be the total number of frames received from v_i during \tilde{T}_i and peer v_0 be the source of the transmission. Recall that \hat{k}_i is the number of frames that have been received and stored into peer's v_i playout buffer during the prefetching period. Then d is expressed as follows:

$$d = \frac{\sum_{i=1}^{n-1} \frac{\sum_{k=\hat{k}_i}^{\tilde{K}_i-1} U_i(k)}{\tilde{T}_i}}{n-1}$$

Loss: Under both Sync and Async, each frame that is not displayed increases the loss by T . Under Sync, discontinuity and loss coincide. Under Async, though, the two are different because a delayed frame causes discontinuity (underflow) but not loss because it is displayed when it eventually arrives. During handoffs, the loss increases when the next missing frames do not exist on the playout buffer of the parent, in which case the scheduler starts pulling and presenting whichever frame is closer (in the future) to the next missing one. We let l denote the average loss ratio expressed as the average among all peers of the total lost playback time experienced by a peer to the total playback time of all frames that should be presented to the user.

Let \dot{K}_i be the difference between the id numbers (sequence numbers) of the last frame that should be displayed at peer v_i at the end of playback and the first frame displayed at v_i when playback started plus 1. Then l is expressed as follows:

$$l = \frac{\sum_{i=1}^{n-1} \frac{\dot{K}_i \cdot T - \tilde{K}_i \cdot T}{\dot{K}_i \cdot T}}{n-1}$$

Thus for the Sync policy $l = d$ since $\tilde{T}_i = \dot{K}_i \cdot T$ and $\dot{K}_i \cdot T - \tilde{K}_i \cdot T = \sum_{k=\hat{k}_i}^{\tilde{K}_i-1} U_i(k)$.

Availability: We define the *peer availability* A to be the average (over an examination period) ratio of compatible pairs over all possible pairs. A pair of peers (v_i, v_j) is deemed compatible when each of them may serve as a parent peer for the other without the client peer experiencing any discontinuity or loss (see Fig.5) given the following assumptions: (1) connection to a parent peer is instant, (2) the parent peer is supplied with new frames at least at the nominal rate, and (3) the available bandwidth between client and parent peers is at least equal to the nominal video rate. It is clear that due to these assumptions A is an upper bound of the real availability that can exist in practice.

Let S be the number of times we observed the system or the number of samples taken during the examination period, t_s be the time the sample with sequence number s is taken, $C_i(t_s)$, $D_i(t_s)$ are a peer's v_i credit (as defined in Appendix A) and offset, respectively, at time t_s and $\mathcal{U}()$ is the unit-step function. Then A is expressed as follows:

$$A = \frac{\sum_{s=1}^S \sum_{i=1}^{n-1} \sum_{\substack{j=0 \\ j \neq i}}^{n-1} \mathcal{U}(C_i(t_s) - D_i(t_s) + D_j(t_s))}{n \cdot (n-1) S}$$

4.2 Description of the Simulation Model

Initial tree formation: We assume discrete time with slot duration set equal to the frame period T . $n = 100$ peers enter the network according to a Poisson arrival process of rate 1 arrival/slot and remain in it for the entire duration of the simulation. We start collecting statistics after all peers have commenced playback.

Video source: The normal playback rate is set to 30 frames/sec, i.e., the video source at the root of the delivery tree makes available a new frame every $T = \frac{1}{30}$ seconds. Frame sizes are extracted from an educational video encoded in MPEG4 format at constant bit rate $R_N = 256$ Kbps (LectureHQ-Reisslein trace file available at ²²).

Available bandwidth of overlay links: We use a simple two-parameter model for obtaining the available bandwidth of each established overlay link: parameter W captures the *average load* (or *congestion level*) in the network, and parameter α captures the *heterogeneity* in terms of the available bandwidth of individual overlay links. We consider that at each time slot of T seconds a directed overlay link L_{ji} from a peer v_j to a peer v_i is “down” (rate 0) with a probability $P_{ji}(\alpha, W)$, which we call the “overlay link drop probability” for a given heterogeneity value α and congestion level value W . At each time slot when an overlay link is “up”, the value of the transmission rate is drawn uniformly at random from $[R_L, R_H]$ Kbps. We present the details of the employed model in Appendix C. This very simple model suffices for an initial qualitative performance comparison between Sync and Async. The overall system being quite complex, it is not clear what more realistic typical workloads look like, so in the end it will take a real prototype to validate our conclusions.

4.3 Experiments

In the following we describe our simulation experiments and present our results.

The number of peers is $n = 100$, the nominal rate of the stream is $R_N = 256$ Kbps, the frame period is $T = \frac{1}{30}$ seconds, when a link is “up” its rate is drawn uniformly at random from the range $[0, 1024]$. All nodes use $D_i = D = 150 \cdot T$ seconds, i.e., initially they prefetch up to 150 frames, which is also their buffer capacity $B_c = 150$. The buffer space is used to store undisplayed (future) frames, thus, unless stated, already displayed frames are discarded since they are no longer useful to a peer. The buffer threshold for triggering a handoff is $B_h = 10$ frames, the grace period is $T_g = 4 \cdot B_h \cdot T$ and the time between a disconnection and reconnection to a new parent peer is $T_h = 5 \cdot T$.

Our control parameters are the heterogeneity α , taking values in $[0.8, 2]$, and the average load W taking values in $[5000, 7000]$. These ranges make a medium to high percentage of overlay links have enough rate to support the nominal stream rate (see Fig. 6(c), 6(d)). Each simulation point in our results is the average of the outcomes of 100 independent runs each of which simulated 50000 time slots of system operation. In each graph the 95th-percentile confidence interval is drawn.

Experiment 1 (Unlimited upload capacity, Global information): In this experiment each peer can serve an unlimited number of downstream peers and has full knowledge of the buffer contents of all other peers i.e., $m = 100$.

In Fig. 1(a) we plot the loss ratio l against the congestion level W , for various heterogeneity values α under Sync(D) and Async(D) payout. As may be seen, Async exhibits a much higher loss throughout the depicted α, W space. For example, for $\alpha = 0.8$ and $W = 7000$ Async has loss ratio 24% whereas Sync has only 1.2%. Increasing congestion hurts both policies as fewer links have enough effective rate (factoring in the up/down transitions) to support the stream; the damage, however, is always worse for Async. Also, for a given congestion level, the loss decreases with increasing heterogeneity, for both policies. This is expected because under high heterogeneity, very few connections have a small available rate, while the majority of them have medium to high rate thus, all together, more overlay links can support the nominal rate under high heterogeneity. In Fig. 1(b) we plot the discontinuity ratio d and the loss ratio l under Async only. Both ratios seem to be pretty close (see discussion later). In Fig. 1(c) we plot the peer availability A , which we use to interpret the previous results on loss and discontinuity. Sync exhibits high availability, close to 1, which decreases slowly as the link rates fall (i.e., with higher W and lower α). Under Async, deteriorating rates de-synchronize the different payout

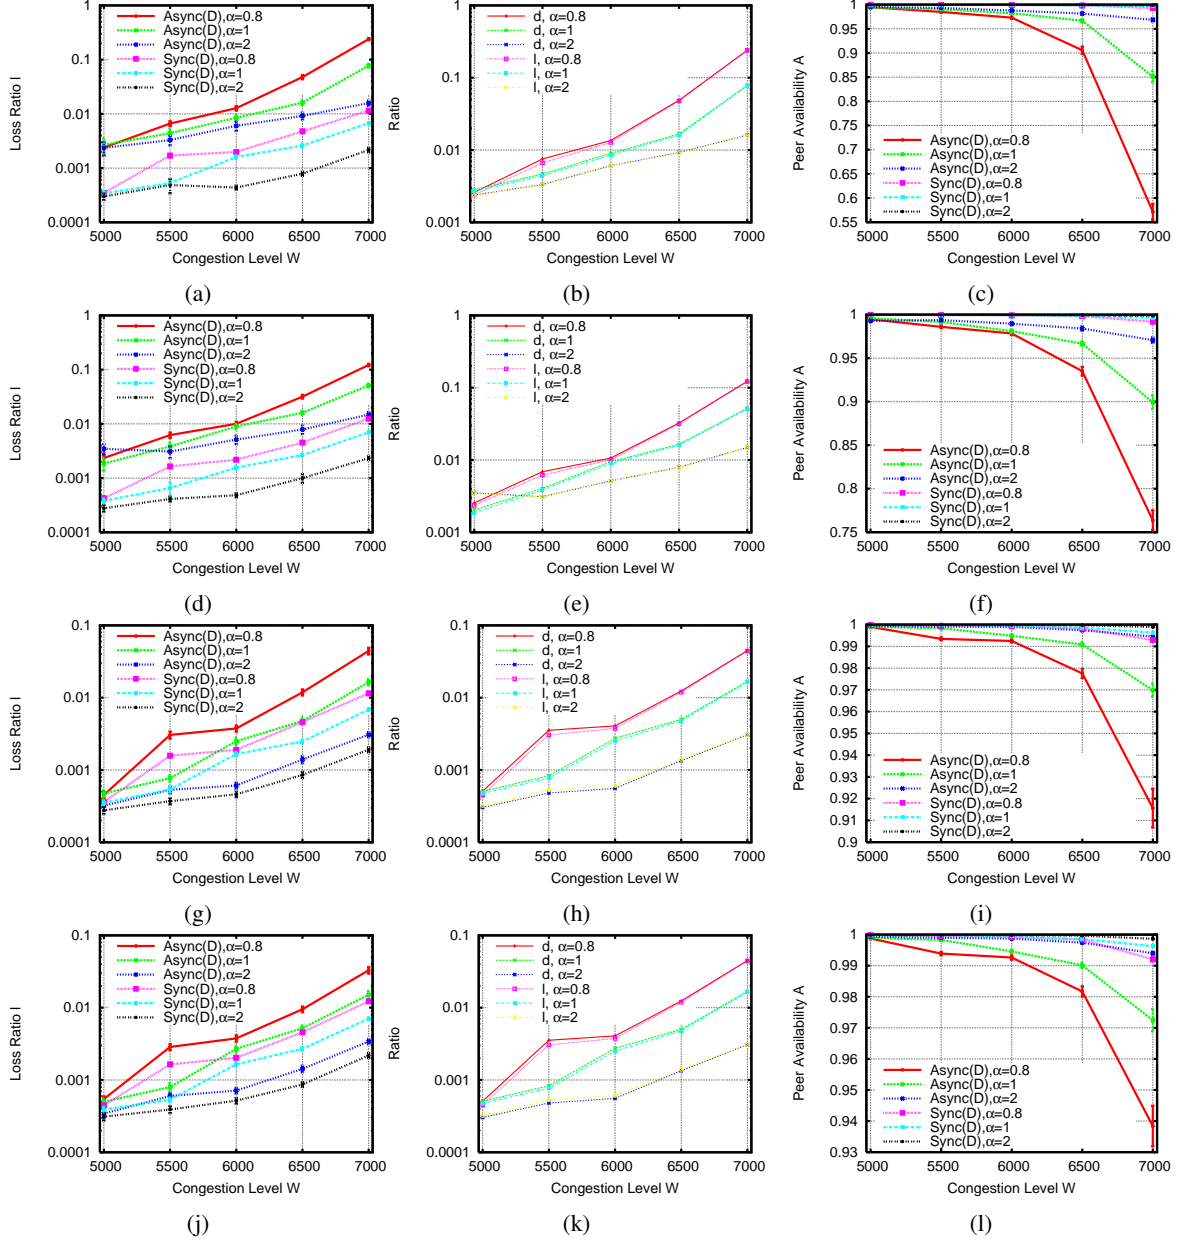


Figure 1. Results from (a)(b)(c) experiment 1,(d)(e)(f) experiment 2,(g)(h)(i) experiment 3 (j)(k)(l) experiment 4.

points, as different nodes face different underflow periods. Consequently, the buffer contents become de-correlated which implies smaller availability of alternative parents which can provide for gapless handoffs. Such, gap-inducing handoffs make Async perform worse with respect to both discontinuity and loss.

Losses occur in Sync whenever a frame is not available at its scheduled playout. Frame lateness can be caused from the artificial jitter we introduce with the aforementioned on/off model for overlay links. For example when the established link to the parent goes off and the buffer level falls beneath the handoff threshold, a node can become underflowed if the new parent cannot feed it immediately with missing frames during the grace period or if the buffer drains during the handoff.

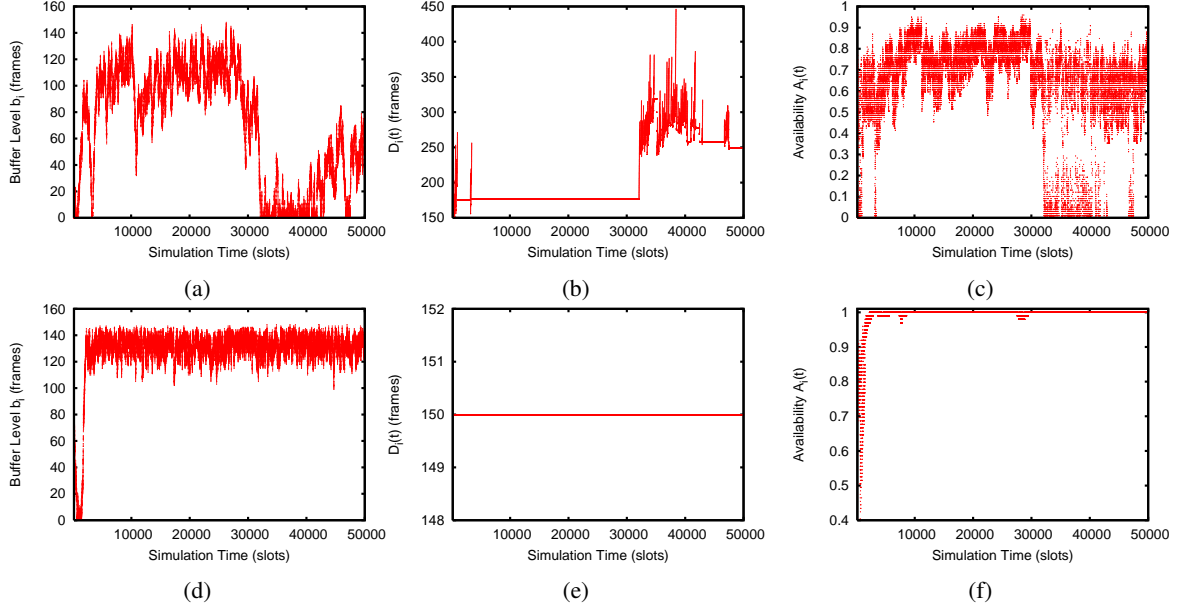


Figure 2. A single peer’s buffer level, offset and availability over time during a single simulation run of experiment 1 for $a = 0.8$, $W = 7000$ under the (a)(b)(c) Async policy ,(d)(e)(f) Sync policy.

In Async such phenomena cause only discontinuity, whereas losses occur only when connection is made to a peer with later than the required frames. Also, recall that for the Sync case a freeze in playback implies also loss of content; thus, a period of discontinuity is also a period of frame loss while frame loss can not exist without discontinuity. The loss ratio in this case represents also the average percentage of playback time that a user experiences an annoying disruption. In the Async case a discontinuity period is not necessarily followed by loss and a loss does not necessarily happen after a discontinuity period. In this case loss ratio represents only the average percentage of content never presented to a peer while the experienced by a user disruption in playback is much higher.

To make the above more clear, we conduct a single simulation run for $\alpha = 0.8$, $W = 7000$ and log at every slot a single peer’s v_i buffer level b_i , current offset from the source $D_i(t)$ and availability $A_i(t)$, for both policies. A peer’s v_i availability $A_i(t)$ at time t is defined to be the ratio of peers which are compatible to v_i to the total number of peers in the system. This is an upper bound to the percentage of peers that may serve as parents for v_i at time t facilitating a disruption-less handoff.

In the Async case a buffer underflow (see Fig.2(a), time period 35000-40000) increases the peer’s offset (see Fig.2(b)) while severely decreasing availability (see Fig.2(c)) since as its playout point diverges from the playout points of other peers; gradually fewer peers are able to serve the peer with the required frames. This leads sooner or later to connection to a parent peer with loss of content; the peer has to consume later frames, which restores its offset at a lower value and availability at a higher value. In this way in the Async case, offset and availability fluctuate through time.

In the Sync case a buffer underflow (see Fig.2(d), time period 0-2000) does not affect the peer’s offset (see Fig.2(e)), as expected, while it only temporarily lowers availability (see Fig.2(f)) which is kept at high levels during all times thereafter due to the fact that since playout points are the same among peers, buffer contents are highly correlated.

In our system setup in order to minimize the probability of experiencing a buffer underflow, a peer seeks another parent when its buffer level falls to the buffer threshold. However, after reconnection to another parent, a buffer underflow may

not be avoided during the grace period even if after the end of that period when buffer level is checked again against the buffer threshold it is found to be above the threshold and thus it is decided that the connection is good enough. Thus, a low connection rate during that period may lead to a buffer underflow and then a rate improvement to enable filling the buffer with missed frames waiting in the transmission queue. This fact explains how collectively buffer contents become decorrelated in the Async case.

Experiment 2 (Limited upload capacity, Global information): In this experiment peers have full knowledge of the buffer contents of all other peers i.e., $m = 100$, but can support only a limited number of downstream peers: the source can support up to 10 peers; all other peers can support up to a number taken uniformly at random from the range $[1,10]$ when joining the streaming hierarchy.

In Fig. 1(d) we plot the loss ratio l under Sync(D) and Async(D) playout. Sync exhibits almost the same loss ratio l as in experiment 1 (see Fig. 1(a)) and seems unaffected by the node outdegree assumed in this experiment. Surprisingly Async exhibits lower loss ratio than before in experiment 1 but still remarkably higher than Sync. The same observation holds for the discontinuity ratio d exhibited by Async which takes similar values with the loss ratio l (see Fig. 1(e)). The behavior changes with α , W as described in experiment 1. In Fig. 1(f) we plot the peer availability A which is close to 1 for the Sync and around 0.75 for Async (i.e., Async is benefited compared to experiment 1 where it got 0.55).

The observed improvement in the performance of Async is attributed to the fact that in the current scenario, because of the limited upload capacity of peers, peers that fall behind due to underflows, do not easily find peers to which they can perform a gapless handoff. Therefore, most of the time they have to induce a gap and connect to a peer that is further ahead in time. The gap hurts momentarily, but in the long run is helpful as it leads implicitly to less divergence of playout points. On the other hand Sync exhibits almost the same performance as in experiment 1 since it manages to always retain a very high availability, thus it is less likely for a peer not to find an appropriate peer to cooperate effectively.

Experiment 3 (Unlimited upload capacity, Partial information): In this experiment each peer can serve an unlimited number of downstream peers and has knowledge of the buffer contents of only a random subset of 10 peers out of all the peers i.e., $m = 10$.

In Fig. 1(g) we plot the loss ratio l under Sync(D) and Async(D) playout. Sync again exhibits almost the same loss ratio l as in experiments 1,2. Async exhibits even lower loss ratio l than in experiment 2 but still higher than Sync. The same observation holds for the discontinuity ratio d (see Fig. 1(h)).

In Fig. 1(i) we plot the peer availability A which is again close to 1 for the Sync case while for the Async takes even higher values (above 0.9).

The performance of Sync and Async are more close in this experiment because peers following the Async policy are more likely to be forced to perform loss inducing handoffs due to the fact that a peer during the peer selection process knows only a random subset of all the peers in the system. This way peers' playout points become almost synchronized and buffer contents become almost correlated, resulting to high availability. Sync as explained above also retains a very high availability thus performance is unaffected.

Experiment 4 (Limited upload capacity, Partial information): In this experiment peers have knowledge of the buffer contents of only a random subset of 10 peers out of all the peers i.e., $m = 10$, while they can support only a limited number of downstream peers: the source can support up to 10 peers; all other peers can support up to a number taken uniformly at random from the range $[1,10]$ when joining the streaming hierarchy.




Figure 3. Loss burst (mean and standard deviation) versus W for $\alpha = 0.8$ in experiment 4.

While Sync performance seems unaffected by restrictions induced in this experiment, Async exhibits a higher performance more close to that of Sync’s (see Fig. 1(j), Fig. 1(k)).

In Fig. 1(l) we plot the peer availability A which is close to 1 for the Sync case while for the Async takes even higher values (above 0.93). The reason for this performance improvement is due to the reasons already mentioned for experiments 2 and 3.

Experiments 2,3 and 4 point to that Async’s performance is favored by randomness in selection since the latter assists in keeping playout points “near-in-time”.

Study of the playback pattern: We define mean *loss burst* L to be the average among all peers of the average number of continuous frames lost during a single loss event over the effective duration of the stream. In Fig.3 we plot loss burst L and its standard deviation for $\alpha = 0.8$ and various W values for experiment 4. We can observe that although for $\alpha = 0.8$ in experiment 4 Sync and Async exhibit pretty close loss ratios, in Async the loss burst is remarkably higher while its standard deviation is close to that of Sync’s. For $W = 7000$ which corresponds to a high congestion level, Sync exhibits a loss burst equal to 11 frames while Async a loss burst equal to 31 frames. This is interpreted that playback under Sync exhibits more skips on the ongoing activity of few sequential frames while under Async less skips of many sequential frames. Also recall that in Async a discontinuity period and a loss period do not necessarily coincide, thus playback is disrupted more often either due to a frame freeze or a frame skip. Although an accurate answer on which playback pattern is more desirable by a viewer would require a subjective quality assessment like ²³ it seems that the adoption of the Async policy results in a more disturbing playback experience even in the cases where both policies exhibit pretty close performance.

The effect of increasing the buffer size: As explained earlier, under the Async policy a rate decrease below the nominal rate during the grace period may lead to a buffer underflow and a freeze in playback. A following drastic rate increase before the end of the grace period enables the peer to fill up the buffer space with the required frames and may lead to a buffer overflow. Thus as the buffer size increases the probability of a buffer overflow decreases. This leads the Async policy to perform slightly better as buffer increases under our rate model while Sync is unaffected by a buffer size increase since its buffer requirements are always constant; in Sync a missed frame is a lost frame, thus the maximum buffer utilization is D/T . It is expected that under a more bursty connection rate model, Async will be significantly favored by a buffer size increase since buffer overflow events are expected to be more often. Avoiding a buffer overflow into a peer not only helps the peer to assure more time ahead of continuous playback but also to serve as a better resource of frames for other peers. However, extensive buffering increases bandwidth usage significantly and this maybe a serious drawback to the deployment of a large scale P2P streaming service. Recent popular P2P streaming systems have embedded this consideration into their design.²⁴ We omit the results of the experiments we conducted varying the buffer




Figure 4. The effect of the number of past frames kept at each peer observed for $\alpha = 0.8$, $W = 7000$, $B_c = 150$ in (a)(b) experiment 1, (c)(d) experiment 3.

size since, although better, Async’s performance is close to the one already presented; it will take a more bursty rate model to substantiate a significantly higher performance.

The effect of keeping past frames: Up to now, peers were assumed to be keeping only a sliding window of future frames. We now let each peer keep a number of past frames i.e., frames that have already been rendered at the local playout scheduler. We retain the terminology used till now and call “buffer” the storage space allocated for undisplayed (future) frames. We use an additional “past frames’ buffer” to keep past frames. We repeat experiment 1 for $\alpha = 0.8$, $W = 7000$, $B_c = 150$ and various sizes of the past frames’ buffer. In Fig.4(a) we plot loss against the number of past frames kept while in Fig.4(b) we plot the availability. One may verify that Sync is oblivious to the number of past frames kept whereas Async is favored from the existence of past frames at each peer. Loss decreases until the number of past frames kept reaches the value 300 at which point it remains flat (this amount of buffer space suffices to mask the worst kind of jitter that can appear under our on/off link model).

Under a similar setup we run experiment 3. The behavior is similar with the only difference that loss ratio is lower than before for past frames kept less than 300 for the reason that divergence of playout points is not favored due to limitations in peer selection (see Fig.4(c),(d)).

Finally, selectively, we conducted simulation runs for peer populations larger than 100 peers. Observations and conclusions presented in this section prove to be also valid when the number of peers in the system scales up.

5. CONCLUDING REMARKS

In this work we examined playout scheduling in jitter-prone environments and “fragile” (single) tree topologies in order to substantiate the maximum damage from using the wrong playout policy in such a setting.

Comparison: We summarize our results and make a side by side comparison between Sync and Async:

- Sync performs consistently better than Async in terms of *discontinuity* and *loss* while both policies exhibit higher performance as heterogeneity increases and congestion level decreases.
- Sync’s performance is unaffected by restrictions in peer selection since it maintains high availability needing more or less only to achieve a good connection to some peer. Async’s performance increases with randomness in peer selection since peers are forced not to diverge a lot by performing handoffs that induce loss.
- Sync policy requires a constant buffer size equal to $\frac{D}{T}$. Async is favored by a buffer size increase. Though, extensive buffering increases bandwidth usage which may be undesirable.
- In Sync a discontinuity period and a loss period coincide. This is not the case in Async where a discontinuity period maybe followed or not by loss while at the same time a loss not necessarily follows a discontinuity period. Thus, playback under Async is disrupted more often either due to a frame freeze or a frame skip.
- In cases where Sync and Async exhibit close performance, playback under Sync exhibits more skips on the ongoing activity of few sequential frames while under Async less skips of many sequential frames.
- Sync is able to satisfy a strict interactivity requirement since it maintains for each viewer a constant offset between encoding and playback times. On the contrary, in Async the offset fluctuates through time in an unpredictable manner.
- While Sync is oblivious to the existence of past frames at each peer, Async is highly favored. However it is not possible to know a priori how many past frames should be kept in order for Async to exhibit a high performance (close to Sync’s performance) since this would require knowledge of the network conditions peers will face.

Applicability of the results to multiple-parent systems: Although in this work we limited our considerations to single distribution trees, we believe that our observations will also carry through to mesh-like topologies.¹ Measurement studies of popular P2P streaming systems supporting distribution over a mesh overlay topology like PPLIVE^{25,26} justify our intuition. PPLIVE employs an Async-like playout policy, in which peers maintain large buffers to store recently presented chunks and chunks scheduled to be played in the future while constantly exchanging “buffer maps”. This measurement study reports that some peers watch frames in a channel minutes behind others while at the same time large average freezing periods close to 1 min are observed. It seems that the adoption of the Sync policy in such a system would improve its performance while it would eliminate the control overhead since availability would be high; thus even a random selection of a list of parents would be sufficient.

Main conclusion: The adoption of a synchronized playout policy in a P2P live streaming system results in “positive correlation” of buffer contents among peers increasing this way the number of upstream relay nodes from which a node can pull frames and thus boost the playout quality of P2P streaming systems. The inherent advantage offered by the P2P architecture is preserved and can be exploited effectively on the benefit of playout quality.

6. ACKNOWLEDGEMENTS

This work has been supported in part by the: i) IST CASCADAS program under contract FP6-027807; ii) IRAKLITOS research fellowship program, co-financed within Op. Education by the ESF (European Social Fund) and National Resources; and iii) NoE CONTENT (IST-384239).

APPENDIX A. SELECTING THE FIRST FRAME AND THE INITIAL BUFFERING DELAY

A peer's v_i playout buffer can be thought to be draining from the bottom (position 1_i holding the currently displayed frame) and filling from the top (the most recently received frame being at position b_i , which is also the buffer occupancy). Let $id\{x_i\}$ denote the id of the frame at position x_i of peer's v_i buffer. Let B_i be the set of the ids of all frames in peer's v_i buffer i.e., $B_i = \{id(1_i), \dots, id(b_i)\}$. Let $C_i(t)$ denote the *credit* of peer v_i at time t , defined to be equal to its remaining discontinuity-free playout time if its input (feeding) rate falls to zero at time t . We can define the credit as follows: $C_i(t) = p_i(id\{b_i\}) + T - t$.

Consider that at time t the new node is v_i and its parent is v_j^\dagger . There are three distinct cases:

- $v_j = v_0$, i.e. v_j is the source of the transmission. Then v_i gets the most recently encoded frame and those that will follow and starts displaying them $F_{ij} = D_i$ time units after the connection time t .
- $v_j \neq v_0$ and assume that v_j is not in its initial prefetching period. v_i starts receiving from v_j the frame at buffer position x_j , $1_j \leq x_j \leq b_j$ and all subsequent ones and starts displaying them $F_{ij} = p_j(id\{x_j\}) - t + D_i - D_j(t)$ time units after the connection time t . $D_j(t)$ denotes the offset between encoding and playout times at time t to node v_j . For the Sync case $D_j(t)$ is constant $\forall t^\ddagger$. The initial targeted offset D_i can only be achieved only if $p_j(id\{x_j\}) - t > D_i - D_j(t)$. Given that this condition is satisfied $p_i(id\{1_i\}) = p_j(id\{x_j\}) + D_i - D_j(t)$. We set $x_j = b_j$ which amounts to retrieving from v_j its newest frame and all subsequent ones. This way the period for buffer buildup subject to the targeted D_i is maximized, giving the chance to prefetch the largest number of frames into the buffer while achieving D_i . This would happen under the assumption that the parent peer will be regularly supplied during that period by its ancestor peer. Thus if $x_j = b_j$, F_{ij} can be also expressed as $F_{ij} = C_j(t) - T + D_i - D_j(t)$.
- $v_j \neq v_0$ and assume that v_j is in its initial prefetching period. Let $F_j(t)$ be the remaining time until node v_j starts playback. Since node v_j will start playback with an initial offset of D_j after $F_j(t)$ then if v_i starts receiving from v_j the frame at buffer position x_j , $F_{ij} = F_j(t) + p_j(id\{x_j\}) - t + D_i - D_j$ given that $F_j(t) + p_j(id\{x_j\}) - t > D_i - D_j$. If $x_j = b_j$ then $F_{ij} = F_j(t) + C_j(t) - T + D_i - D_j$.

APPENDIX B. PARTITIONING V_I

Figure 5 illustrates all the cases of possible playout buffer contents' correlations between a peer v_i performing handoff and a candidate parent peer v_j . The distinct cases become clear in the following where the sets $V_i^A, V_i^B, V_i^C \subseteq V_i$ are determined given the offsets $D_i(t), D_j(t)$ and the credit values $C_i(t), C_j(t)$ at time t of peers v_i and v_j respectively.

- If $D_i(t) = D_j(t)$, peers v_i and v_j have the same offsets between playback and encoding times, thus at time t they both view the same frame. If v_j has more buffered frames for future playback than v_i , $C_i(t) < C_j(t)$ then $v_j \in V_i^A$ since v_i will be able to find the next frame to the one on top of its buffer to peer's v_j buffer (Fig.5,1.a). If $C_i(t) \geq C_j(t)$ then $v_j \in V_i^B$ (Fig.5,1.b,1.c).
- If $D_i(t) > D_j(t)$, peer v_j precedes in playback peer v_i displaying at time t a later frame than the one displayed at v_i . If $C_i(t) \geq D_i(t) - D_j(t) \wedge C_i(t) \leq C_j(t) + D_i(t) - D_j(t)$ then $v_j \in V_i^A$ (Fig.5,2.b, 2.c). If $C_i(t) > D_i(t) - D_j(t) \wedge C_i(t) > C_j(t) + D_i(t) - D_j(t)$ then $v_j \in V_i^B$ (Fig.5,2.d). If $C_i(t) < D_i(t) - D_j(t)$ then $v_j \in V_i^C$ (Fig.5,2.a).

[†] v_i and v_j are considered to have synchronized clocks.

[‡]For $D_i < D_j(t)$ i.e. v_i is less interactive than its parent the received frame will be presented at v_i , $D_j(t) - D_i$ earlier than at v_j while for $D_i > D_j(t)$ the received frame will be presented at v_i , $D_i - D_j(t)$ later than at v_j . For $D_i = D_j(t)$ it will be presented to both the same time.

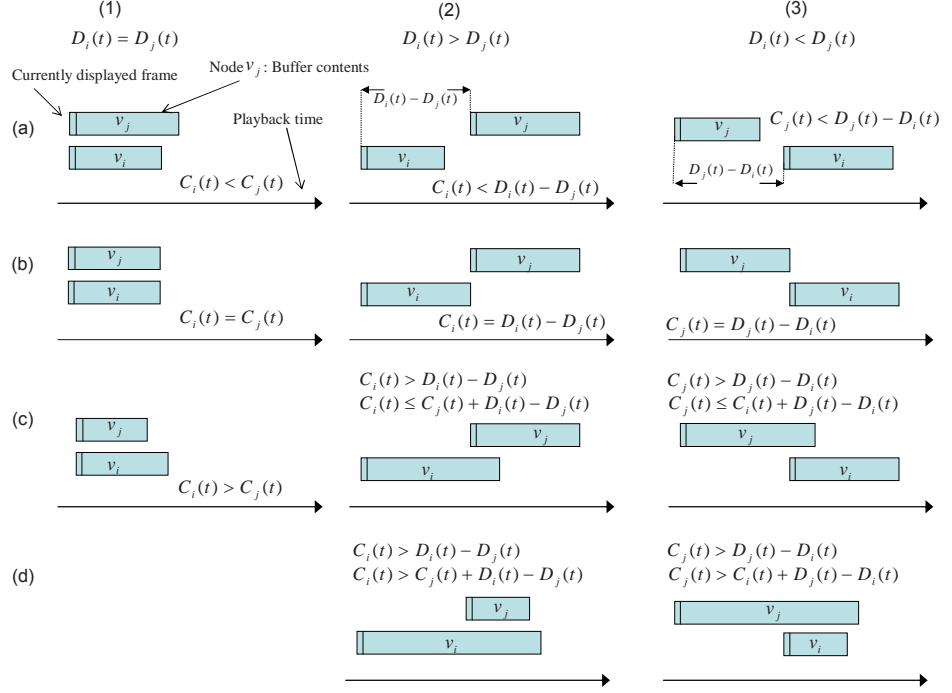


Figure 5. An illustration of all the possible cases of playout buffer contents' correlations between a peer v_i performing handoff and a candidate parent peer v_j .

- If $D_i(t) < D_j(t)$, peer v_i precedes in playback peer v_j . If $C_j(t) \geq D_j(t) - D_i(t) \wedge C_j(t) \leq C_i(t) + D_j(t) - D_i(t)$ then $v_j \in V_i^B$ (Fig.5,2.b, 2.c). If $C_j(t) > D_j(t) - D_i(t) \wedge C_j(t) > C_i(t) + D_j(t) - D_i(t)$ then $v_j \in V_i^A$ (Fig.5,2.d). If $C_j(t) < D_j(t) - D_i(t)$ then $v_j \in V_i^B$ (Fig.5,2.a).

APPENDIX C. DETAILED MODEL FOR THE AVAILABLE BANDWIDTH OF OVERLAY LINKS

We consider that at each time slot of T seconds a directed overlay link L_{ji} from a peer v_j to a peer v_i is “down” with a probability $P_{ji}(\alpha, W)$ called the overlay link drop probability for a given heterogeneity value α and congestion level value W . $P_{ji}(\alpha, W)$ is defined at start $\forall i, j$ according to following procedure:

1. An indexed list K of pairs (each pair represents a directed overlay link) is produced after a random permutation of all pairs.
2. Probability mass is distributed according to a generalized power law with parameter α and a probability P'_i is assigned to the i th directed link of the list produced in the previous step $\forall i$.
3. The list of produced probabilities in the previous step is traversed from the start and each probability P'_i is multiplied by a weight factor W . If the result is less than 1 then $P_i = P'_i \cdot W$ else $P_i = 1$ and the extra mass $P'_i \cdot W - 1$ is distributed proportionally to the rest of the following link drop probabilities in the list. This step repeats until all items of the list have been visited. In this way the sum of probabilities is the same under the same W and different α values.
4. Probability P_i is assigned to the i th directed link in K .

Steps 2 and 3 are presented in detail in Algorithm 1. By using different values α we model different levels of heterogeneity in terms of expected overlay link rates while the weight W captures the overlay network's congestion level. At each time

slot when an overlay link is “up”, the value of the transmission rate is drawn uniformly at random from the range $[R_L, R_H]$ Kbps.

In Fig.6(a) we draw P_i versus i for various α values and a constant value W while in Fig.6(b) we draw P_i versus i for various W values and a constant value α . The number of peers is 100 and thus the number of distinct directed overlay links that can exist by considering all possible pairs is 9900. In Fig.6(c) we present the percentage of the overlay links which are expected to support a rate above the nominal $R_N = 256$ Kbps versus α and W when link rate is drawn uniformly at random from $[0, 1024]$ when a link is “up”. In Fig.6(d) the average of expected overlay link rates is drawn versus α and W also when link rate is drawn uniformly at random from $[0, 1024]$ when a link is “up”. Notice that according to our algorithm the average is kept constant for the same W and different α values.

Algorithm 1 Generation of the Overlay Link Drop Probabilities

```

1:  $m = n \cdot (n - 1)$  { $m$  is the number of directed overlay links}
2:  $C = (\sum_{i=1}^m \frac{1}{i^\alpha})^{-1}$ 
3: for  $i = 1$  to  $m$  do
4:    $P_i = \frac{C}{i^\alpha}$ 
5:    $P_i = P_i \cdot W$ 
6: end for
7: for  $i = 1$  to  $m$  do
8:    $extramass = \max(0, P_i - 1)$ 
9:    $P_i = \min(1, P_i)$ 
10:  if  $extramass > 0$  then
11:     $s = 0$ 
12:    for  $j = i + 1$  to  $m$  do
13:       $s = s + P_j$ 
14:    end for
15:    for  $k = i + 1$  to  $m$  do
16:       $P_k = P_k + \frac{extramass \cdot P_k}{s}$ 
17:    end for
18:  end if
19: end for

```

REFERENCES

1. N. Magharei, R. Rejaie, Y. Guo, Mesh or multiple-tree: A comparative study of live p2p streaming approaches, in: INFOCOM 2007, Anchorage, Alaska, 6-12 May 2007.
2. N. Laoutaris, I. Stavrakakis, Intrastream synchronization for continuous media streams: A survey of playout schedulers, IEEE Network Magazine 16(3).
3. C. Vassilakis, N. Laoutaris, I. Stavrakakis, The impact of playout policy on the performance of p2p live streaming...or how not to kill your p2p advantage, in: 15th Annual SPIE/ACM Multimedia Computing and Networking (MMCN '08), San Jose, California, January 2008.

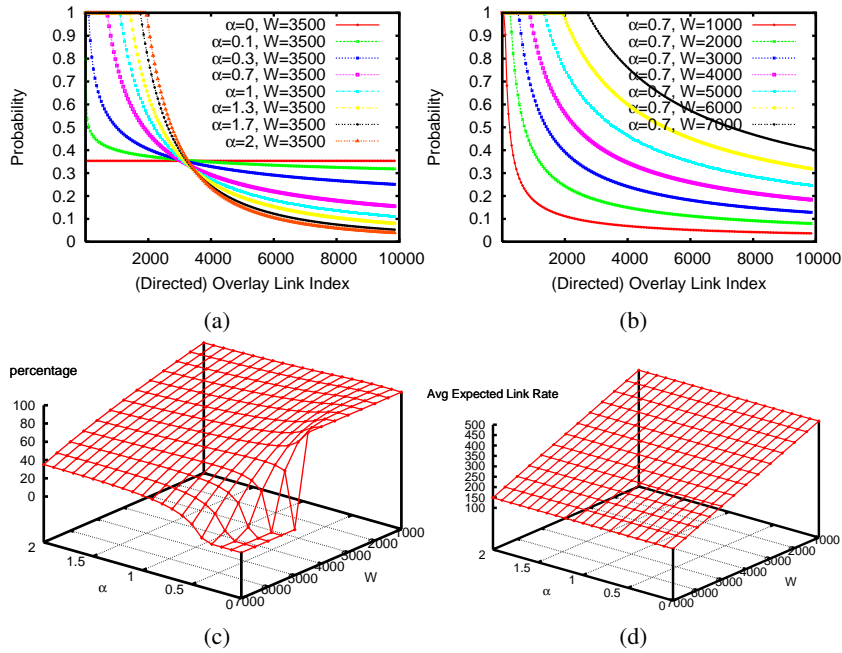


Figure 6. (a) Probability of an overlay link to be “down” versus the overlay link index for a constant W and various α values, (b) Probability of an overlay link to be “down” versus the overlay link index for a constant α and various W values, (c) Percentage of directed overlay links that their expected rate is greater or equal to 256Kbps versus α and W when link rate is drawn uniformly at random from $[0,1024]$ when a link is “up”, (d) Average expected link rate versus α and W when link rate is drawn uniformly at random from $[0,1024]$ when a link is “up”.

4. W.-P. K. Yiu, X. Jin, S.-H. G. Chan, Challenges and approaches in large-scale p2p media streaming, *IEEE MultiMedia* 14(2) (Apr-Jun 2007) pp. 50–59.
5. S. Rao, Establishing the viability of end system multicast using a systems approach to protocol design, Carnegie Mellon University, Phd Thesis, Technical Report CMU-CS-04-168, Oct. 2004.
6. D. L. Hayes, Advanced copyright Issues on the internet, Fenwick and West LLP, 2007.
7. H. Deshpande, M. Bawa, H. Garcia-Molina, Streaming live media over peer-to-peer network, Stanford University, Technical Report, 2001.
8. Y. hua Chu, S. G. Rao, S. Seshan, H. Zhang, A case for end system multicast, *IEEE Journal on Selected Areas in Communication (JSAC)*, Special Issue on Networking Support for Multicast 20(8).
9. S. Banerjee, B. Bhattacharjee, C. Kommareddy, Scalable application layer multicast, in: *ACM Sigcomm 2002*, Pittsburgh, Pennsylvania, August 2002.
10. D. Kostic, A. Rodriguez, J. Albrecht, A. Vahdat, Bullet: high bandwidth data dissemination using an overlay mesh, in: *ACM SOSP '03*, New York, USA, Oct. 2003.
11. D. A. Tran, K. A. Hua, T. T. Do, A peer-to-peer architecture for media streaming, *IEEE Journal on Selected Areas in Communication (JSAC)* 22(1).
12. M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, A. Singh, Splitstream: High-bandwidth multicast in a cooperative environment, in: *ACM SOSP '03*, New York, USA, Oct. 2003.

13. V. N. Padmanabhan, H. J. Wang, P. A. Chou, K. Sripanidkulchai, Distributing streaming media content using cooperative networking, in: ACM NOSSDAV 2002, Miami Beach, FL, USA, May 2002.
14. V. Venkataraman, K. Yoshida, P. Francis, Chunkyspread: Heterogeneous unstructured end system multicast, in: ICNP 2006, Santa Barbara, California, 12-15 November 2006.
15. A. Rowstron, P. Druschel, Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems, in: IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), Heidelberg, Germany, November 2001, pp. 329–350.
16. X. Zhang, J. Liu, B. Li, T.-S. P. Yum, Donet/coolstreaming: A data-driven overlay network for peer-to-peer live media streaming, Vol. 3 of Proc. INFOCOM 2005, Miami, FL, USA, 13-17 March 2005, pp. 2102–2111.
17. N. Magharei, R. Rejaie, Prime: Peer-to-peer receiver-driven mesh-based streaming, in: INFOCOM 2007, Anchorage, Alaska, 6-12 May 2007.
18. Bittorent, <http://www.bittorrent.org>.
19. C.-C. Yeh, L. S. Pui, On the frame forwarding in peer-to-peer multimedia streaming, in: Workshop on Advances in Peer-to-Peer Multimedia Streaming (In conjunction with ACM Multimedia 2005), 11 Nov. 2005, Hilton, Singapore.
20. H. Jiang, S. Jin, Nsync: Network synchronization for peer-to-peer streaming overlay construction, in: ACM NOSS-DAV '06, Newport, Rhode Island, May 2006.
21. N. Laoutaris, B. V. Houdt, I. Stavrakakis, Optimization of a packet video receiver under different levels of delay jitter: An analytical approach, *Performance Evaluation* 55(3-4) (2004) 251–275.
22. V. T. R. Group, <http://trace.eas.asu.edu>.
23. Y. Qi, M. Dai, The effect of frame freezing and frame skipping on video quality, in: 2006 International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP'06), Pasadena, California, Dec 2006.
24. Joost, <http://www.joost.com>.
25. X. Hei, C. Liang, J. Liang, Y. Liu, K. W. Ross, Insights into pplive: A measurement study of a large-scale p2p iptv system, in: Workshop on Internet Protocol TV (IPTV) Services over World Wide Web (in conjunction with WWW2006), Edinburgh, Scotland, May 2006.
26. PPLIVE, <http://www.pplive.com>.