QoS-Competitive Video Buffering

Alexander Kesselman
Tel-Aviv University

Yishay Mansour
Tel-Aviv University

Abstract

Many multimedia applications require transmission of streaming video from a server to a client across an internetwork. In many cases loss may be unavoidable due to congestion or heterogeneous nature of the network. We explore how discard policies can be used in order to maximize the quality of service (QoS) perceived by the client. In our model the QoS of a video stream is measured in terms of a cost function, which takes into account the discarded frames. In this paper we consider online policies for selective frame discard and analyze their performance by means of *competitive analysis*. In competitive analysis the performance of a given online policy is compared with that of an optimal offline policy. In this work we present competitive policies for a wide range of cost functions, describing the QoS of a video stream.

1 Introduction

In a network where resources such as bandwidth and buffering capacity are constrained there is a need for an efficient video delivery system that can achieve high resource utilization and maximize the QoS perceived by the user. In reality both network bandwidth and buffering capacity are likely to be limited. Under such circumstances there may be a situation in which a feasible transmission schedule is impossible, i.e., any transmission necessarily incurs loss of data. In this case instead of being denied service, clients may choose lower quality video streams with occasional frame loss. This is a natural situation in today's Internet where the network may not have sufficient bandwidth to support the peak rate of a video stream [4], or available bandwidth may fall below the requirement at the middle of a video transmission.

When delivering a video stream across a resource-constrained network, a naive approach is to discard frames with no awareness of the video stream properties. As a result the QoS perceived by the user may degrade drastically, even for small amounts of loss (e.g. tail drop of consecutive frames). In this paper we study intelligent selective frame discard policies, which take into account application-specific properties.

Previous work. Recently a number of packet discarding schemes incorporating application-specific information have been proposed. In [7] appears a simple strategy called Frame-Induced Packet Discarding, in which upon detection of loss of a threshold number of packets belonging to a video frame, the network attempts to discard all the remaining packets of that frame. In [3], [10] the problem of optimizing the quality of the transmitted video for a given cost function by jointly selecting the source rate and the channel rate has been considered with leaky bucket constraints. Our work differs from theirs in that we are trying to optimize the QoS perceived by the user, rather than minimizing loss in general.

In [16] *offline* algorithms for optimal selective frame discard have been considered. The notion of *selective frame discard* at the server has been introduced and the optimal selective frame discard problem using a QoS-based cost function has been defined. There appear various algorithms for selective frame discard for stored video delivery across a resource-constrained network. In contrast with [16], we deal with *online* problem, when no information regarding the video stream is known a-priory.

In [6] has been considered the problem of smoothing real-time streams. They determined the optimal tradeoff between buffer space, delay and link bandwidth for lossy smoothing schedules and presented simple generic algorithm. Competitive analysis of jitter control online algorithms appears in [5].

Our work. We measure the performance of our algorithms using competitive analysis [11], [1]. In competitive analysis the performance of an online policy is compared with those of an optimal offline policy, which knows in advance the entire sequence of frame arrivals. Competitive analysis is a natural approach for Internet traffic, which is unpredictable and chaotic [12]. In this paper we study competitive online policies for intelligent selective frame discard at time of buffer overflow. An intelligent frame discard policy should take in account the application-specific properties in its decisions to discard particular frames, minimizing the QoS degradation perceived by the user. We consider video streams without inter-frame dependencies, such as JPEG video streams. In order to measure the *QoS* perceived by the user we define a so called *well-behaved* cost function reflecting the playback discontinuity at the client. The cost of a discarded frame depends either on the distance to the closest discarded frame or on the position within a sequence of consecutive discarded frames it participates in. The QoS level function is the sum of the costs of the discarded frames. We consider various video stream settings, and our main result is that the competitive ratio of the Greedy Policy is bounded by a constant for all of these settings. In addition to the global QoS function we consider local metrics, such as the minimal distance between two discarded frames and maximal length of a sequence of consecutive discarded frames. The Greedy Policy is shown to be competitive with regards to these metrics as well.

The rest of this paper is organized as follows. Section 2 contains the model description. The scheduling policy is defined in Section 3. Section 4 contains analysis of fixed size frames. The conclusion appears in Section 5.

2 Model Description

We consider a system with a single FIFO buffer delivering a video stream. We assume that the buffer can hold exactly M frames when frame size is fixed and at at least M frames of maximal size when frame size may vary. Frames may arrive to the queue at any time and send events are synchronized with time. We divide time into slots so that during each slot one video frame is sent. We assume that differences between transmission times of identical video frames due to variations in frame size are negligible. A queue policy has to decide for each incoming frame whether it should be either rejected or accepted (subject to the buffer capacity constraints). A frame in the buffer may be also *preempted* by the policy. However, all the frames must be sent in FIFO order.

Now we describe the traffic-shaping policy, i.e. the regulation of the rate at which a flow is allowed to inject packets into the QoS network. We consider leaky bucket traffic-shaping mechanism in which only a fixed amount of traffic is admitted to the network. Excess traffic is held in a queue until either it can be accommodated or must be discarded. Wrege et al. [14], [15] investigate providing deterministic performance guarantees to VBR video using multiple leaky buckets traffic constraint function. A (σ, ρ) leaky bucket model with a burst parameter σ and a rate parameter ρ, during time interval of length t, has at most $\rho \cdot t + \sigma$ packet arrivals. Our results vary with the size of σ.

Definition 1 *A (σ, ρ)-source is a source that generates stream that is shaped by a (σ, ρ) leaky bucket policer.*

In our model each frame has the corresponding *cost*. The cost of a frame depends on the positions of previously discarded frames. The goal is to design a policy that minimizes the total cost of the system. Next we give a formal definition of the notation used. We start with a notation for a sequence of frames.

Definition 2 *A sequence of frames is denoted by S. The start and the finish time of S are denoted by t^{start} and t^{final}.*

Next we set the notation to indicate lost frames.

Definition 3 *The loss indicator $X_f(t)$ to whether frame f was discarded at time t, $X_f(t)$ is 1 if f is discarded prior to time t and 0 otherwise.*

Now we define the parameters which determine the cost of a discarded frame.

Definition 4 *The closest discarded frame is said to be* adjacent *discarded frame. The* distance *between two adjacent discarded frames is the number of transmitted frames between them. Two discarded frames are said to be* consecutive *if they are at zero distance. Let $d(f,t)$ be the distance between a frame f and the adjacent discarded frame preceding f at time t. Let $pos(f,t)$ be the position of a frame f within a sequence of consecutive discarded frames at time t, or 0 if f does not participate in such a sequence, i.e., if the previous and the following frames to f are not discarded at time t.*

Once we have the necessary notation we can define the structure of the cost function.

Definition 5 *The cost of a discarded frame f at time t is denoted by $c_{\phi_1,\phi_2}(f,t)$. If $d(f,t) > 0$ then $c_{\phi_1,\phi_2}(f,t) = \phi_1(d(f,t))$, otherwise $c_{\phi_1,\phi_2}(f,t) = \phi_2(pos(f,t))$.*

We define the cost incurred by a policy over a sequence of frames as the sum of the costs of discarded frames.

Definition 6 *The cost of a policy A while scheduling S is the sum of the costs of the discarded frames, i.e., $L_A(S) = \sum_f X_f(t^{final}) c_{\phi_1,\phi_2}(f, t^{final})$.*

Next we define a class of well-behaved cost functions. In order to provide measure of playback discontinuity, a cost function should take two aspects into consideration: the length of a sequence of consecutive discarded frames and the spacing between two adjacent, but non-consecutive discarded frames (see [8]). The motivation behind the definition is as follows. For a single discarded frame the QoS degradation is minor while for sequences of consecutive discarded frames it is significant. Therefore, we would like the cost of a single discarded frame to be bounded by constant and the cost of a discarded frame within a sequence to be bounded by a constant factor of its position within a sequence of discarded frames [1]. The cost of a frame within a sequence must be at least as high as the cost of

[1]Notice that the cost of a lost block would depend quadratically on its length.

a single discarded frame since it is always preferable, with respect to the video stream QoS, to drop a frame whose immediate neighbors are not discarded. The QoS perceived by the user degrades as the distance between two discarded frames decreases or the length of a lost block increases. So we require a cost function to be monotonically increasing with: (1) decreasing distance between two non-consecutive discarded frames (2) increasing position of a discarded frame within a sequence of consecutive discarded frames. The best QoS is achieved when the loss is distributed as evenly as possible over the frames sequence, therefore the cost function should be minimal in this case. We require that the cost function be convex so that its minimum is established when all dropped frames are evenly spaced.

Definition 7 *We say that a cost function c_{ϕ_1,ϕ_2} is well-behaved iff there exist positive constants L, U, L' and U', $0 < L \leq U \leq L' \leq U'$, satisfying the following constraints with respect to ϕ_1:*

1. $\forall d: \; L \leq \phi_1(d) \leq U$ *(bounded cost)*
2. *if* $x < y$ *then* $\phi_1(x) > \phi_1(y)$ *(ϕ_1 is anti-monotone)*
3. $\forall x, y: \; \phi_1((x+y)/2) \leq (\phi_1(x) + \phi_1(y))/2$ *(ϕ_1 is convex)*

and with respect to ϕ_2:

1. $\forall l: \; L' \cdot l \leq \phi_2(l) \leq U' \cdot l$ *(bounded average-cost)*
2. *if* $x > y$ *then* $\phi_2(x) > \phi_2(y)$ *(ϕ_2 is monotone)*
3. $\forall x, y: \; \phi_2((x+y)/2) \leq (\phi_2(x) + \phi_2(y))/2$ *(ϕ_2 is convex).*

Let us consider the following simple well-behaved cost function: $\phi_1(d) = 1 + 1/d$ and $\phi_2(l) = 4l$, where $L = 1$, $U = 2$, $L' = 3$ and $U' = 5$. The functions are presented on Figure 1.

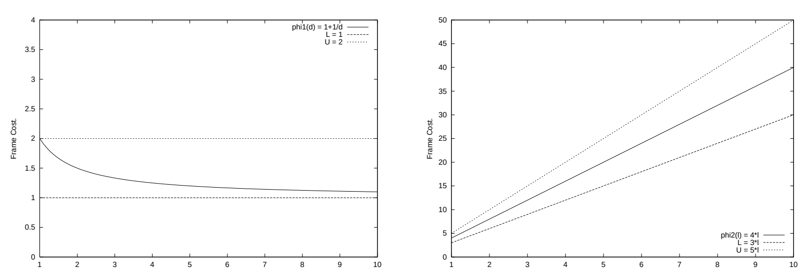

Figure 1: The cost functions $\phi_1(d) = 1 + 1/d$ and $\phi_2(l) = 4l$.

To analyze the performance of the Greedy Policy we introduce the following definitions.

Definition 8 *Let the* current lost set *of a policy at time t be the set of discarded frames that arrived after $t-M-1$. Let the* current set *of a policy at time t be the union of its lost set at time t and the frames in the buffer at time t. Let a* lost block *be a maximal sequence of consecutive discarded frames.*

Now we turn to define competitive analysis. In competitive analysis (see [11]) the performance of an online policy is compared with those of the optimal policy (OPT), which knows in advance the entire sequence of frame arrivals. The competitive ratio is the minimum over all input sequences of the ratio between the cost incurred by OPT and the the cost incurred by the given online policy. More formally,

Definition 9 *A policy A is c-competitive if for every sequence of frames S, $c \cdot L_A(S) \leq L_{OPT}(S)$. (Note that $0 \leq c \leq 1$). In such a case we also say that A is competitive.*

Notice that a competitive online policy must not discard any frame from an input sequence S if OPT does not discard any frame from S.

3 Greedy Policy

We define a natural "Greedy Policy". The state of the policy depends on the previous decisions of the policy to discard particular frames.

Greedy Policy: Each time when the buffer is full and a frame arrives the policy discards a frame that minimizes the sum of the increase in cost of previously discarded frames plus the cost of the discarded frame itself.

3.1 Scheduling Example

Intuitively, when Greedy selects a frame to be discarded, the decision is optimized locally with respect to the current state of the system. On the other hand OPT optimizes its decision globally with respect to the entire schedule. We study the performance of the Greedy policy scheduling moderate and large bursts. Let us consider a system consisting of a buffer that is able to hold 3 frames. We define the *schedule* of a policy to be the sequence of arrive and send events (if any) for a stream of video frames.

Moderate bursts: For moderate bursts Greedy does not drop adjacent frames, as we prove latter. Suppose that at time 0 the buffer is empty and a burst of 3 frames arrives. During the following t time units one frame is sent by the Greedy policy and one frame arrives. Finally, at time t a burst of 4 frames arrives. In this case Greedy would drop 3 frames among the last six frames because its buffer is full. At the same time OPT would evenly distribute the loss over the whole sequence, which can be done since the buffer is not empty throughout this time interval. The resulting video streams appear on Figure 2 (dropped frames are marked by $\times$).

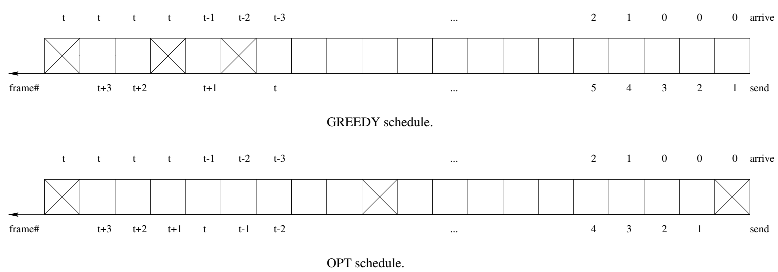

Figure 2: Outcoming video streams of Greedy and OPT (moderate bursts).

Large bursts: In presence of massive bursts both OPT and Greedy are forced to drop consecutive frames. In this case the loss of OPT is uniformly divided between lost blocks while Greedy may create larger lost blocks. However, they are balanced later by future Greedy discards. Suppose that at time 0 the buffer is empty and a burst of 3 frames arrives. During the following t time units one frame is sent by the Greedy policy and one frame arrives. Finally, at time t a burst of 8 frames arrives. Notice that Greedy would have its buffer full and OPT would have its buffer empty at this time. Now Greedy and OPT would discard 7 and 5 frames, respectively. The resulting video streams are shown on Figure 3.

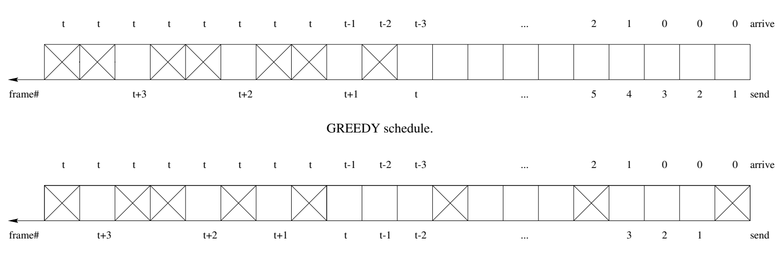

Figure 3: Outcoming video streams of Greedy and OPT (large bursts).

3.2 Implementation Details

The algorithm can be implemented in distributed fashion on routers supporting per-flow queuing, where each video stream is allocated its own buffer space. If a frame consists of multiple packets then loss of a packet belonging to a video frame means loss of the whole frame and causes discard all the remaining packets of that frame.

The main problem is how to exchange information regarding the discarded frames between routers. It can be done by adding to the header of a frame the following *drop info* for both forward and backward directions.

- If the frame is not adjacent to a lost block, it keeps the distance to the closest discarded frame;
- If the frame is adjacent to a lost block, it keeps the length of the lost block.

The maximal distance to the closest discarded frame can be bounded by a reasonable number (e.g. 30 frames). The main advantage of this approach is that the drop policy is stateless, i.e., there is no need to keep additional per-stream information at the router.

Complexity: The Greedy Policy simply scans the drop info stored in the headers of the frames currently in the buffer in order to decide which frame to discard. Thereafter it drops the selected frame and updates the drop info of the relevant frames. Thus, the complexity of the Greedy Policy together with drop info update procedure is *linear* on the number of frames in the buffer, which is typically a small number.

4 Uniform Size Frames

In this section we consider video streams without inter-frame dependencies, such as JPEG video streams, and assume that the buffer can hold exactly M frames. We use the following definitions regarding the local properties of a schedule.

- Let the maximal distance between two adjacent discarded frames produced by any policy when k frames are lost in time interval of length s be denoted by $d_{OPT}(k,s)$.
- Let the maximal length of a lost block produced by any policy when k frames are lost in time interval of length s be denoted by $l_{OPT}(k,s)$.

4.1 Moderate Burst Size

We consider video streams generated by a (σ,ρ) source with $\sigma \leq \frac{5}{4}M$. We show that the competitive ratio of the Greedy Policy is at least L/U.

Although there exist other policies that are able to guarantee the same worst-case performance (e.g. drop each other frame), the Greedy Policy, which tries to optimize the QoS each time a frame is discarded, is a natural one with respect to average-case performance. To establish the competitive ratio of the Greedy Policy we introduce an estimation of the optimal maximal distance between discarded frames. Then we derive the cost of the loss of the Greedy Policy using this estimation. In particular we demonstrate that the distance between two discarded frames using the Greedy Policy is at least half the estimation minus one, and show that it is almost optimal distance for a competitive online policy (up to a factor of 2). First we need some auxiliary lemmas.

Observation 1 *When the number of frames lost in a time interval of s time units is k the optimal maximal distance satisfies $d_{OPT}(k,s) = \lceil \frac{M+s}{k-1} \rceil$.*

The observation holds since during s time units at most $M+s$ frames could be accepted. By the monotonicity and the convexity properties of ϕ_1, the cost function is minimized when the discarded frames are equally distributed, i.e. when the minimal distance between two discarded frames is maximized. A simple combinatorial argument shows that the maximum is achieved when discarded frames divide the $M+s$ accepted frames into $k-1$ balanced parts. The following lemma, compares the performance of the Greedy Policy to d_{OPT}.

Lemma 1 *When frames are scheduled according to the Greedy Policy and the size of the lost set is bounded by k the minimal distance between two adjacent discarded frames is at least $d_{OPT}(k,0)/2-1$.*

Proof. Suppose by way of contradiction that the Greedy Policy discards a frame violating the condition of the lemma. Let the size of the current lost set at this time moment be $m \leq k-1$. We assume that frames in the current lost set are enumerated starting from the end of the sequence. Let d_0 denote the number of accepted frames in the buffer preceding the most recent discarded frame in the current lost set, d_i denote the distance between the i'th and $i+1$'th discarded frames in the current lost set, and d_m denote the number of accepted frames in the buffer succeeding the latest discarded frame in the current lost set or the distance between the latest discarded frame in the current lost set and the previous discarded frame if any.

The Greedy Policy always tries to maximize the resulting minimal distance between two discarded frames by definition of well-behaved cost function. If the condition of the lemma is violated then the following holds. The distance between any pair of adjacent discarded frames is at most

$d_{OPT}(k,0)-2$ and d_0 is not greater than $d_{OPT}(k,0)/2-1$. If it is not the case the Greedy Policy would have been able to discard a frame without having the minimal distance falling below $d_{OPT}(k,0)/2-1$.

Observe that the buffer is full and each of the frames currently in the buffer takes part in exactly one of the above distances. Thus the sum of distances bounds M, and hence,

$$M \le \sum_{i=0}^{m} d_i \le d_0 + (m-1)(d_{OPT}(k,0)-2) + d_m \le (k-1)(d_{OPT}(k,0)-2).$$

The last inequality follows from the fact that if there exists a discarded frame not in the current lost set then $d_m \le d_{OPT}(m,0)-2$ and $m \le k-2$, otherwise $d_m \le d_{OPT}(m,0)/2-1$ and $m \le k-1$. By this way we obtain that

$$d_{OPT}(k,0) \ge \tfrac{M}{k-1} + 2,$$

which contradicts with Observation 1. □

In the next theorem we show that the Greedy Policy is L/U-competitive. Notice that the Greedy Policy schedules a maximum number of frames.

Theorem 1 *For any sequence generated by a (σ,ρ)-source with $\sigma \le \frac{5}{4}M$, the competitive ratio of the Greedy Policy is at least L/U.*

Proof. Notice that OPT discards exactly the same number of frames as Greedy Policy does. Observation 1 and Lemma 1 imply that the minimal distance between two adjacent discarded frames is at least 1. So the Greedy Policy incurs cost of at most U per frame since OPT has to discards frames with cost of at least L. □

Next we show that there exists a sequences of frames such that no online policy could improve the $d_{OPT}(M/4,0)$ distance estimation while OPT is able to discard frames arbitrarily far apart. The intuition behind the proof is that no online policy can discard frames at time of underflow without being 0-competitive. That means that a competitive online policy should be able to receive a burst of maximum size simultaneously, while OPT could smoothly distribute the discarded frames throughout the sequence. Due to the lack of space proof of the following theorem is omitted.

Theorem 2 *There exists an input sequence generated by a (σ,ρ)-source with $\sigma \le \frac{5}{4}M$ so that any online policy discards frames at distance of at most $d_{OPT}(M/4,0)$ while in OPT the distance between any two discarded frames is arbitrarily large.*

The following is an immediate consequence of Theorem 2.

Corollary 1 *The competitive ratio of any online policy is at most L/U.*

4.2 Large Burst Size

We consider video streams generated by a (σ, ρ)-source with $\sigma > \frac{5}{4}M$. This is an interesting case for bursty VBR video streams. Notice that now the Greedy Policy may be forced to discard consecutive frames. In a similar spirit to the distance between discarded frames, we obtain bounds on the lengths of lost blocks. The problem here is that the dependence of a cost of such a sequence on its length is not linear but quadratic. This means that the cost of loss may increase drastically when σ increases.

Nevertheless, it turns out that in this case OPT also suffers large loss. We establish that the competitive ratio of the Greedy Policy is $min(\frac{U'}{11L}, \frac{U'}{15L'})$. In order to show this we introduce an estimation of the optimal minimal length of a lost block. Then we derive the cost of the loss of the Greedy Policy and OPT in terms of this estimation. Moreover, we show that the length of a lost block produced by the Greedy Policy is at most factor of 4 larger than the lower bound for OPT. Before we prove the main theorem we show a few lemmas. The following observation states the value of l_{OPT}.

Observation 2 *When the number of frames lost throughout a time interval of s time units is k, the maximal length of a lost block satisfies $l_{OPT}(k, s) = \lceil \frac{k}{M+s+1} \rceil$.*

The observation holds since during s time units at most $M + s$ frames could be accepted. A simple combinatorial argument shows that any schedule will have a lost block of size at least $\lceil \frac{k}{M+s+1} \rceil$. The proof of the following lemma is omitted.

Lemma 2 *When frames are scheduled according to the Greedy Policy and the size of the lost set is bounded by k the maximal length of a lost block is at most $2l_{OPT}(k, 0) + 1$.*

For simplicity we assume in the sequel that $l_{OPT}(k, s) = \frac{k}{M+s+1}$.

Observation 3 *The cost of k discarded frames divided into lost blocks of length l under a well-behaved function is bounded from below by $L' \cdot \frac{l+1}{2}k$ and bounded from above by $U' \cdot \frac{l+1}{2}k$.*

Next lemma shows that for large bursts OPT unavoidably pays the cost of producing sizable lost blocks.

Lemma 3 *When frames are scheduled according to the Greedy Policy and the size of the current lost set is $k > M$ OPT discards at least $k-M$ frames from the current lost set and pays cost of at least $L'(l_{OPT}(k - M, M) + 1)(k - M)/2$.*

Proof. The current lost set contains frames that arrived during the last M time units. Observe that the number of frames accepted in this time interval by the Greedy Policy and by OPT may differ by at most M frames, i.e. OPT is able to accept M additional frames if its buffer is empty at the beginning of the interval. The lemma follows by Observation 2 and Observation 3. □

The next proof uses a technique of "k-matching" between loss of the Greedy Policy and OPT. In k-matching the loss of the online policy are divided into disjoint sets that are matched to sets of frames lost by OPT so that each frame lost by OPT appears in at most k such sets. The loss of the Greedy Policy are matched to the loss of OPT either by 1-matching or by 3-matching. Evidently, the competitive ratio of the Greedy Policy is at least the minimum of the minimum ratio of 1-matched sets and one third of the minimum ratio among 3-matched sets.

Theorem 3 *For any sequence generated by a (σ, ρ)-source with $\sigma > \frac{5}{4}M$, the competitive ratio of the Greedy Policy is at least $min(\frac{L}{11U'}, \frac{L'}{15U'})$.*

Proof. We divide the schedule of the Greedy Policy into intervals of length M and consider lost sets at the last time moment of every interval. Note that each discarded frame belongs to exactly one of these sets. These lost sets are matched to the loss of OPT either by 1-matching or by 3-matching. First we identify and match the lost sets of the Greedy Policy that participate in 3-matching. Having finished with 3-matching, the remaining lost sets participating in 1-matching are arbitrarily matched to sets of the same cardinality formed from the left unmatched frames lost by OPT. Notice that both Greedy Policy and OPT lose the same number of frames. Thus there remains a sufficient number of free frames lost by OPT for 1-matching provided that 3-matched online lost sets have cardinality smaller than their offline counterparts.

We bound the cost of a lost set of the processed interval by determining the maximal length of a lost block the frames from the interval's lost set could belong to. Suppose that we process the i'th interval $[Mi, M(i+1))$. Let k be the cardinality of the lost set of the interval (i.e. at time M(i + 1) - 1) and let t_{max} be the time moment within the interval $M(i-1) \leq t_{max} < M(i+1)$ in which the cardinality k_{max} of the lost set is maximal. Clearly, the maximal length of a lost block frames from the lost set could participate in is the estimation of the Lemma 2 at time t_{max}, that is $2l_{OPT}(k_{max}, 0)+1$.

Now the value of k_{max} determines the kind of matching in the following way.

(1) If $k_{max} > 10M$ then the lost set of the current interval takes part in 3-matching. Next we show how to find a matching of the interval lost set to a set of frames lost by OPT, whose cost constitutes at least $\frac{L'}{5U'}$ fraction

of its cost. We match the lost set to a set of $k_{max} - M$ frames that were necessarily lost by OPT during time interval $[t_{max} - M - 1, t_{max}]$. This is indeed a 3-matching since each of the offline lost frames could participate in at most three such sets, that is the matched set of the interval and the corresponding matched sets of the adjacent intervals.

The Lemma 3 implies that the cost incurred by OPT is at least

$$L'(l_{OPT}(k_{max} - M, M) + 1)(k_{max} - M)/2.$$

At the same time by Observation 3 the cost incurred by the Greedy Policy is at most

$$U'(l_{OPT}(k_{max}, 0) + 1)k \leq U'(l_{OPT}(k_{max}, 0) + 1)k_{max}.$$

Hence, we can bound the ratio between the cost of the offline and the corresponding online lost sets by $\frac{L'}{5U'}$, since $k_{max} > 10M$.

(2) If $k_{max} \leq 10M$ then the lost set of the current interval takes part in 1-matching. By Observation 2 the optimal maximal length of a lost block is at most 10. Applying Lemma 2 we obtain that the maximal length of a lost block of discarded frames for the Greedy Policy is at most 21. Therefore according to Observation 3 the cost of the lost set is upper bounded by $11k \cdot U'$. The ratio is kept above $\frac{L}{11U'}$ since each lost frame has cost at least L.

The competitive ratio of the Greedy Policy is the minimum of the ratio of 1-matched sets and one third the ratio of 3-matched sets, which establishes the theorem. □

Next we show that the length of a lost block created by the Greedy Policy is at most four times the optimal length of a lost block of OPT plus three.

Lemma 4 *For any input sequence S the length of the maximal lost block satisfies $l_{GREEDY}(S) \leq 4 \cdot l_{OPT}(S) + 3$.*

Proof. By Lemma 2 the maximal length of a lost block produced by the Greedy Policy is at most $2l_{OPT}(k, 0) + 1$ when the size of the lost set is k. At the same time by Lemma 3 the optimal length of a lost block in this case is at least $l_{OPT}(k - M, M)$. The theorem follows since,

$$2l_{OPT}(k, 0) + 1 \leq 4l_{OPT}(k - M, M) + 3.$$

□

We conclude with the following general theorem.

Theorem 4 *For any sequence of frames, the competitive ratio of the Greedy Policy is at least $min(L/U, \frac{L}{11U'}, \frac{L'}{15U'})$.*

5 Conclusion

In this work we study competitive online buffering policies for video transmission across internetwork with leaky bucket constraints. First we consider video encoding schemes with independent frames, such as JPEG and its derivatives. To measure the *QoS* we define a well-behaved cost function reflecting the playback discontinuity at the client. For moderate and large burst parameters we derived the competitive ratio of the Greedy Policy. In addition to the global QoS function we consider local metrics, such as the minimal distance between two discarded frames and maximal length of a sequence of consecutive discarded frames. The Greedy Policy is shown to be competitive with regards to these metrics as well.

The proposed policy may be used for managing current Internet routers that wish to provide QoS. We discuss distributed implementation of the Greedy Policy, which may be easily deployed and operate at very high speeds. Interesting future directions include studying more sophisticated cost functions and performing simulations in which performance of online dropping policies is estimated more properly using global QoS function.

References

[1] A. Borodin and R. El-Yaniv, "Online Computation and Competitive Analysis," *Cambridge University Press*, 1998.

[2] R. L. Cruz, "A Calculus for Network Delay, Part 1: Network Elements in Isolation," *IEEE Transactions on Information Theory*, 37(1):114-131, January 1991.

[3] C.-Y. Hsu, A. Ortega and A. R. Reibman, "Joint Selection of Source and Channel Rate for VBR Video Transmission under ATM Policing Constraints," *IEEE Journal on Sel. Areas in Communications*, vol. 15, pp. 1016-1028, August 1997.

[4] T. V. Lakshman, A. Ortega and A. R. Reibman, "Variable bit rate (VBR) video: Tradeoffs and potentials," *Procedeends of the IEEE*, vol. 86, pp. 741-755, December 1997.

[5] Y. Mansour and B. Patt-Shamir, "Jitter Control in QoS Networks," *39th Annual Symposium on Foundations of Computer Science*, 1998.

[6] Y. Mansour, B. Patt-Shamir and Ofer Lapid, "Optimal Smoothing Schedules for Real-Time Streams," *Proceedings of PODC 2000*, pp. 21-29.

[7] S. Ramanathan, P.V. Rangan, and H.M. Vin, "Frame-Induced Packet Discarding: An Efficient QoS Management Strategy for Video Networking," in *Fourth International Workshop on Network and Operating System Support for Digital Audio and Video*, Lancaster University, November 1993.

[8] J. M. Del Rosario and G. C. Fox, "Constant bit rate network transmission of variable bit rate continuous media in video-on-demand servers," *Multimedia Tools and Applications*, vol. 2, pp. 215-232, May 1996.

[9] J. D. Salehi, Z.-L. Zhang, J. F. Kurose and D. Towsley, "Supporting stored video: Reducing rate variablity and end-to-end resource requirements through optimal smoothing," in *Proceedings of ACM SIGMETRICS*, pp. 222-231, May 1996.

[10] S. D. Servetto, K. Ramchandran, K. Nahrstedt, and A. Ortega, "Optimal Segmentation of a VBR Source for its Parallel Transmission over Multiple ATM Connections," in *Proceedings of the IEEE International Conference on Image Processing*, Santa Barbara, CA, October 1997.

[11] D. Sleator and R. Tarjan, "Amortized Efficiency of List Update and Paging Rules," *CACM 28*, pp. 202-208, 1985.

[12] A. Veres and M. Boda, "The Chaotic Nature of TCP Congestion Control," *Proceedings of INFOCOM 2000*, pp. 1715-1723, March 2000.

[13] D. Wijesekera, J. Srivastava, "Quality of Service (QoS) Metrics for Continuous Media," *Multimedia Tools and Applications*, vol. 2, pp. 127-166, September 1996.

[14] D. E. Wrege, E. W. Knightly, H. Zhang and J. Liebeherr, "Deterministic Delay Bounds for VBR Video in Packet-Switching Networks: Fundamental Limits and Practical Tradeoffs," *IEEE/ACM Transactions on Networking*, 4(3):352-362, June 1996.

[15] D. E. Wrege and J. Liebeherr, "Video Traffic Characterization for Multimedia Networks with a Deterministic Service," *IEEE Infocom '96*, pages 537-544, March 1996, San Francisco, California.

[16] Z. L. Zhang, S. Nelakuditi, R. Aggrawal and R. P. Tsang, "Efficient selective frame discard algorithms for stored video delivery across resource constrained networks," In *Proc. IEEE INFOCOM 98*, 1998.

Alexander Kesselman Department of Computer Science, Tel-Aviv University, *email: alx@math.tau.ac.il.*

Yishay Mansour Department of Computer Science, Tel-Aviv University, *email: mansour@math.tau.ac.il.*