

# Avoiding quality oscillations during adaptive streaming of video

Wassim Ramadan      Eugen Dedu\*      Julien Bourgeois

Laboratoire d'Informatique de l'Université de Franche-Comté  
Numérica, Cours Leprince-Ringuet, BP 21126, 25201 Montbéliard, France  
{Wassim.Ramadan, Eugen.Dedu, Julien.Bourgeois}@pu-pm.  
univ-fcomte.fr

## ABSTRACT

A high number of videos, encoded in several bitrates, are nowadays available on Internet. A high bitrate needs a high and stable bandwidth, so a lower bitrate encoding is usually chosen and transferred, which leads to lower quality too. A solution is to adapt dynamically the current bitrate so that it always matches the network bandwidth, like in a classical congestion control context. When the bitrate is at the upper limit of the bandwidth, the adaptation switches constantly between a lower and a higher bitrate, causing an unpleasant oscillation (zigzag) in quality on the user machine. This paper presents a solution to avoid such oscillations. It uses an EWMA (Exponential Weighted Moving Average) value for each bitrate, which reflects its history. The evaluation of the algorithm shows that loss rate is much smaller, bitrate is more stable, and so received video quality is better.

**Keywords:** Real time content, Video streaming, Rate control, Congestion control

---

\*Corresponding author.

## 1 INTRODUCTION

Nowadays, the number of videos encoded in several bitrates and accessible for everyone increases significantly day by day. Their contents are generally delivered to final user using streaming services over Internet. These services as well as the demand for high video quality (e.g. HD and 3D videos) are in constant progression. They require more and more bandwidth, hence available bandwidth variation must be taken into account to shorten buffering time at the receiver.

Currently, one video bitrate is chosen at the beginning of a video streaming; the transmission is controlled at the network layer (TCP or UDP) and application is not involved at all. Hence, two choices exist when playing streamed video content. The first is to choose a low video bitrate, and the video is played directly, without interruption. The second is to choose a bitrate higher than average bandwidth, buffer mul-

timedia data at the user and play the video when the buffer has sufficient data; because the buffer empties, the user will be confronted to many play/pause during the streaming. Both cases are unpleasant to the user eye: the first has the advantage to watch a fluid video but with low quality, while the latter allows to have a good quality but with either short or long waiting time and with frequent interruptions. Users wishing to have an instant access without delay to multimedia content and with the best possible quality can use a new kind of multimedia streaming services, *multimedia content adaptation*.

For video streaming in a network with highly variable bandwidth, a *content adaptation* is a way to adapt the video bitrate to the network characteristics, hence to improve the video quality perceptible by the final user. A cooperative approach between application layer and network layer can be used. Transport protocol handles the congestion control on the network side, while the application handles the video bitrate control on the server side. Bitrate can be controlled by changing the quantisation parameters, or the FPS (frames per second) or other parameter.

An undesirable effect which appears when multiple video qualities are available and when bandwidth changes, is that the bitrate control leads to constant switching between two qualities (bitrates); one is smaller than the bandwidth and the second is greater. For example, when a user is connected to Internet through a 2.5Mb/s link and the video is available in 2 and 3Mb/s qualities, then an adaptive streaming algorithm will constantly switch between these two qualities; obviously, the best solution would be to stay with 2Mb/s much more time before retry-

ing 3Mb/s quality. We call this problem *zigzag quality switching*.

Few papers treat this issue, and they do not solve it completely. This paper presents a solution, called ZAAL (**Z**igzag **A**voidance **A**lgorithm), to this problem<sup>1</sup>. It uses a successfulness value of each bitrate, and its average is constantly updated. This average is used to take the decision whether a next higher quality can be chosen or not, thus preventing the zigzag.

This paper is organized as follows. Section 2 formulates the problem which we try to solve, and section 3 compares it to similar existing problems. Section 4 presents our ZAAL algorithm as a solution to the problem, and its performance is evaluated through real experiments in section 5. Finally, section 6 concludes this article and presents some perspectives.

## 2 PROBLEM FORMULATION

*The zigzag quality switching is the fact that an adaptive video streaming sender application keeps switching the video quality between two values, especially when one value is greater and the other is smaller than the bandwidth capacity.*

In a network with highly variable bandwidth, video adaptation is very important. It aims to adapt the video bitrate to the network characteristics and thus to improve the video quality perceptible by the final user. The adaptation can be done at several OSI layers, but we take the appli-

---

<sup>1</sup>This article is an extension of one of our previous papers [1].

cation layer to exemplify. To adapt video to network conditions, the sender application controls some video parameters, such as the bitrate, FPS (Frames per second) or image size; for example, if more bandwidth is available, application increases the video bitrate, and if less bandwidth is available, application decreases it. To reach this goal, application should constantly retrieve the bandwidth (each 100ms, each 2 sec. or each I image for MPEG video for example) and adapt its bitrate accordingly. To minimise the quality fluctuation, it is not preferable, in our opinion, to change quality each tiny period (at each sent packet for example). Sometimes the application has the choice among multiple available video qualities (e.g. many videos on Internet are encoded with several qualities). This video adaptation is known in the literature as “rate adaptive video control”.

Fig. 1 presents an example which allows to better understand how video adaptation works. The transport protocol has a congestion control which gives the rate at which the packets leave the socket buffer (and afterwards the machine, and enter the network). When current bandwidth is smaller than video bitrate, the socket buffer fills with packets. When the buffer becomes full, exceeded packets generated by application will fail to be written. We call these packets “failed packets”. The goal of an adaptive video content application is to readjust video bitrate in this case by choosing a lower video quality. When new bitrate does not cause failed packets, application will retry a higher quality to match again the available bandwidth and to enhance the video quality.

Our previous paper VAAL (Video Adaptation at Application Layer) [2] is an example of




Figure 1: Video data flow on the sender side.

video adaptation method. It uses transport protocol buffer overflow as a method to find out the available bandwidth and to adapt video content bitrate to the discovered bandwidth. Each 2 seconds, the server application computes the number of failed packets. This number is used to control the video bitrate afterwards. A high number means smaller bandwidth and smaller bitrate. Zero error indicates either a stable or more bandwidth, so bitrate of sent video can be increased. The experiments presented in this paper use VAAL.

An adaptive application, such as the one presented above, keeps switching between two qualities: one with bitrate higher than bandwidth and with failed packets, and another one with bitrate smaller than bandwidth and with no error. We noticed this phenomenon in our experiments (see Fig. 4(a) for a clear example, where during the first minute the video bitrate is continuously toggling between 0.5 and 1 Mb/s). This is what we called “zigzag quality switching” problem; it has two undesirable effects:

- Numerous bitrate changes, leading to unstable video quality. It is preferable to keep the video quality stable as much as possible while maximising it.
- Many lost packets (when bitrate is higher than bandwidth), leading to lower video quality and wasted resources (CPU and

network) for these lost packets. Indeed, videos are generally more affected by packet losses than lower bitrate.

Note that *the problem to be solved is not to reduce the number of quality changes*, for example when such changes use much resources (CPU, disk etc.) *It is neither to use the highest available quality at a certain moment* (which can potentially lead to zigzag, for example when right afterwards the bandwidth decreases). The problem to be solved is the *zigzag oscillations*, i.e. avoiding to reuse a quality which has recently been used (one or more times) and has proved unsuccessful. As such, a method which does not take into account the recent history somehow cannot solve this problem.

### 3 POSITIONING COMPARING TO RELATED WORK

We found that our problem has similarities with other problems, and with two techniques already found in the literature.

#### 3.1 Similar Methods

One method which can be used to solve zigzag quality switching is presented in [3]. It is proposed for multicast streaming video but can be adapted to unicast transmissions too. The sender sends several layers (base and enhancement), and the receiver subscribes dynamically to one or several of them. The receiver uses a timer for each level of subscription (layer). At the beginning the timer is short. When the timer expires

and no loss was experienced, the receiver subscribes to the next new level and the timer for that level is started. If on the contrary the level led to lost packets, the receiver goes back to previous level and the timer for the level with losses is multiplicatively increased.

In this method, there is no superior limit for the timer, which means that the quality increasing can be forbidden for a very long time, even if bandwidth becomes greater in the meantime. Also, each time, only the timer of the current bitrate is updated, which means that good conditions for higher levels do not ameliorate the timer of lower levels. Both these characteristics are unrealistic. On the contrary, our EWMA-based algorithm does not have these drawbacks.

A metric to evaluate the jerkiness of a video, given by the number of quality changes, is given in [4]. It uses a formula to calculate the Effective Frame Rate (EFR) of a video:

$$EFR = \sum_{i=1}^{N-1} \left( \frac{fps_i - P * qchanges(max(i - W, 0), i)}{N} \right) \quad (1)$$

where  $fps_i$  is the number of frames delivered in the  $i$ th second of the video,  $W$  is the window size,  $P$  is a weighting factor for variation in frame rate, and  $qchanges(max(i - W, 0), i)$  is the number of quality changes in the range  $i - W$  to  $i$ .

This formula is used to limit the number of quality changes during each window ( $W$ ). As such, it counts the number of quality changes, no matter *where* they appear inside the window, which is however an important parameter of the jerkiness. For example, 10 quality changes in the first 2 seconds of a video of 1 minute

are worse visually than 10 quality changes dispersed through the whole video. The goal of EWMA, used in our solution, is exactly to take into account the time of the change too. Moreover, in [4] the window is static (coarse-grained moving) and fixed (same size), while in our case we use a sliding and dynamic window.

Finally, our goal is not to limit the number of quality changes, but to avoid quality increases which lead to quality decreases right afterwards. For all these remarks, the formula in this paper is not interesting in avoiding the zigzag problem we want to tackle.

The adaptation video method described in [5] does not cope directly with zigzag problem but presents a way for smoothing sent video bitrate and reducing the frequency of video quality changes. The main idea is that the application does not switch to a higher video quality (or video layer in a hierarchical video encoding) until the sender is certain that the video will continue playing even after a reduction of the congestion window.

This paper presents results for AIMD and SQRT congestion controls for video streaming. To guarantee the continuous video playing, the sent video quality should be the highest quality which verify:  $C < R - \beta R^l$ , where  $C$  is the bitrate at which the video is encoded,  $R$  is the transmission rate and  $\beta = 1/2$  and  $l = 1$  for AIMD based congestion control and  $l = 1/2$  for SQRT based one. This formula implies that the highest video quality does not exceed half of the transmission rate for AIMD, and it is  $\beta\sqrt{R}$  smaller than the bandwidth for SQRT.

Results given in [5] show that this formula

works well for AIMD based congestion control streaming application but the video bitrate is most of the time very low compared to available bandwidth. On the other hand, it has a little impact on SQRT based algorithm, i.e. the quality changes are not significantly reduced. Finally and most importantly, this formula does not really solve the zigzag quality switching, but merely reduces the bitrate artificially.

### 3.2 Comparison With Bandwidth Estimation Techniques

These techniques aim to estimate the available bandwidth at the time of the measurement.

A well-known technique is packet pair [6]. The basic idea is that the sender sends a pair of packets of the same size to a destination (it could be a specialised server or simply an aware receiver). The destination host then responds by sending an echo for each received packet. By measuring the changes in the time spacing between the two packets, the sender can estimate the available bandwidth of the network path as following:

$$bw = \frac{s}{t_2 - t_1} \quad (2)$$

where:

- $bw$  is the bandwidth of the bottleneck link
- $s$  is the packet size
- $t_1$  and  $t_2$  are the arrival time of the first and second packet respectively.

The accuracy of this technique depends on several assumptions. The most important one is

that the two packets should be enqueued *one after the other* at the bottleneck link, which is not guaranteed when a router has a non FIFO queue or when another packet is inserted between the two packets. Other variants of packet pair technique are developed [7, 8, 9] to mitigate this effect.

Another known technique is packet train [10]. A burst of packets is sent between source and destination. Only packets in the same train are used for measurement. A train is formed by packets for which the spacing between two consecutive received packets does not exceed some inter-packet (inter-train) gap. Like in packet pair, inter-arrival times between packets in the same train are used for estimating the available bandwidth.

All these methods have several drawbacks:

1. Sending/receiving additional data packets is needed to estimate the available bandwidth. If video data packets themselves are used for this, receiver must distinguish them from other data, for example by adding a new field in the packet header.
2. Probing packets should be specifically dispersed, i.e. they need a specific timing when they are sent. This makes the use of the video data packets themselves difficult for probing purposes.
3. If a new connection is used for probing packets, sender/receiver should send/listen to a different socket (port).
4. Not only the sender application but also the receiver application must be modified to be able to respond to those probing packets.

Changing both endpoints is known to be difficult to realise in practice.

Finally, and the most important, these techniques have a different goal than ours. As they do not take into account the bitrates used in the past, they cannot avoid the zigzag.

The previous cited reasons show that current bandwidth estimation techniques are not appropriate to solve the zigzag quality switching problem.

### 3.3 Comparison With Network Congestion Control

Our problem is similar to classical network congestion control problem, but not identical. We consider both window-based (written as TCP [11] in the following) and equation-based (written as TFRC [12] in the following) congestion controls.

In fact, both TCP/TFRC and our ZAAL method try to solve a multi-criterion optimisation problem, but the criteria involved to achieve this goal are not really identical.

*Zigzag:* ZAAL aims to avoid the zigzag between two consecutive qualities, which occurs when (1) the inferior quality is lower than bandwidth, hence no loss, hence the quality is increased, *and* (2) the superior quality is higher than bandwidth, hence a few losses, hence the quality is decreased to the inferior quality again. TCP/TFRC have fine-grained sending rate; TCP, being a *data* transport protocol, does not pay attention to such zigzag, while TFRC smooths the sending rate, without trying to avoid the zigzag.

*Losses:* TCP/TFRC aim to send at the maximum rate, even if this yields a few losses. On the contrary, ZAAL aims to avoid losses as much as possible. In reality, losses cannot be completely avoided unless a very low bitrate is used. So for ZAAL it is better to maintain a quality without or with few losses, instead of the next higher quality with high loss rate. In fact, ITU.T G.1070 [13] recommends that the end-to-end IP packet loss rate in video streaming should be less than 10%. As an example, during a video transmission ZAAL prefers a 2Mb/s sending rate without losses to a 3Mb/s sending rate (50% more packets) with 20% losses.

*Sending rate:* TCP/TFRC aim to send at the highest possible rate. This is why they *constantly increase* sending rate (using various laws, which differ from one TCP variant to another). A video adaptation method aims that too, but does not increase quality if ZAAL considers, for zigzag/loss avoidance as shown above, that the next higher quality should not be tried at this time. Thus an adaptation method can be seen as quality increasing/decreasing proposition and ZAAL as quality increasing-only blocker.

This comparison shows that classical congestion control algorithms have different goals, hence they are not an appropriate to solve the zigzag problem.

## 4 ZIGZAG-AVOIDING ALGORITHM OVERVIEW

To solve the zigzag quality switching generated by a video adaptation method, an additional specific algorithm can be used, such as ZAAL.

ZAAL algorithm works by avoiding constantly using bitrates higher than the available bandwidth. For that, it maintains an average value for each bitrate, called *successfulness* in the following. When the adaptive algorithm considers to increase bitrate (and only in this case), ZAAL checks if the successfulness of the higher bitrate is lower than a threshold, called  $\beta$ ; if this is the case, a higher bitrate *cannot* be chosen. Otherwise said, application uses a higher bitrate  $i$  only if its successfulness  $S_i > \beta$ . After this process, the average successfulness is updated.

Note that ZAAL is not an adaptation method. It is used only to prevent an adaptation method from frequently switching the video quality. The only information ZAAL needs comes from the adaptation method, whether some bitrate causes lost packets or not.

### 4.1 Algorithm

ZAAL algorithm uses the *successfulness* value each time an adaptation period ends, which depends on the adaptation algorithm (e.g. 2 sec. in case of our VAAL adaptation method). Average successfulness value is calculated separately for each bitrate (with different weights), denoted by  $S_i$ , which indicates if bitrate of index  $i$  can be used for the next period or not. As such, this average value expresses the application last attempts to use the corresponding bitrate. In brief, when a bitrate generates failed packets to the transport protocol buffer, corresponding successfulness value is greatly reduced; when a bitrate was successful, the successfulness value is greatly increased; finally, when a bitrate is used for a long time, the successfulness values corresponding to higher bitrates are slowly increased.

As a general rule, the smaller the successfulness average value, the more the corresponding bitrate causes failed packets and application must avoid using it.

The average successfulness  $S_i$  (where  $i$  is bitrate index) of each bitrate, which changes according to the history of that bitrate, is calculated using an EWMA algorithm (Exponential Weighted Moving Average). Using EWMA allows to give greater weight to recent history compared to older history, since obviously current bandwidth is better expressed by recent bitrate usage than by older ones. Additionally, different weights are used, based on the bitrate involved.

At the beginning of a video transmission, all  $S$  values are set to 1. Then they are calculated each time the application wants to adapt the video bitrate to the available bandwidth, using the following general formula:

$$S_i = (1 - \alpha/d)S_i + s(\alpha/d) \quad (3)$$

where:

- $s$  is the successfulness at the time of measurement (the current “observation”) and can be either 0 or 1. 0 value is used when the bitrate did not give good results (hence its successfulness average value will decrease). 1 value is used when the bitrate gave good results, either because the bitrate did not cause failed written packets, or because the bitrate was not used recently (hence its successfulness average value will increase).
- $\alpha$  is the degree of weighting increase/decrease, a constant smoothing

factor between 0 and 1. A higher  $\alpha$  discounts older successfulness values faster.

- $d$  is a division factor allowing to speed up or slow down the average value increasing depending on the bitrate involved. In our algorithm,  $d$  has three values: 1, 2 and 4. For  $s = 1$ , the greater the  $d$ , the slower the increasing of the average  $S_i$  value. They will be explained more later.

$S$  values are arithmetic reals between 0 and 1. They can change in three cases:

1. First, when the application increases the video bitrate, i.e. the new bitrate  $k$  is higher than the actual one  $j$ . This appears when the application does not sense lost packets for a while with actual video quality. In this case  $S$  should increase ( $s = 1$ ) for all bitrates  $i$  lower than or equal to the current bitrate. Increasing speed must be high ( $d = 1$ ). So:

$$S_{i|i \leq j} = (1 - \alpha)S_i + \alpha \quad (4)$$

2. Second, when the application maintains the bitrate. This happens either when some packet losses occur but their rate is acceptable to maintain the current quality  $j$ , or successfulness of the higher bitrate  $k$  ( $S_k$ ) is low and application avoids using it.  $S$  value increases for all qualities but with different division factor values (different speeds). So:

- $d = 1$  and  $s = 1$  (big increase) for all bitrates  $i$  lower than the current one  $j$ :

$$S_{i|i < j} = (1 - \alpha)S_i + \alpha \quad (5)$$



- $d = 2$  and  $s = 1$  (small increase) for current bitrate  $j$ :

$$S_j = (1 - \alpha/2)S_j + \alpha/2 \quad (6)$$

- $d = 4$  and  $s = 1$  (very small increase) for the bitrate  $k = j + 1$  higher than the current one  $j$ :

$$S_{k(k=j+1)} = (1 - \alpha/4)S_k + \alpha/4 \quad (7)$$

3. Third, application decreases the video bitrate. This appears when the available bandwidth is not enough anymore for the actual bitrate (i.e. bandwidth decreases during the current period), or when the application tries a bitrate higher than the bandwidth. In this case, average value must be reduced for the current bitrate  $j$ , other bitrates average values do not change. Hence,  $s = 0$  and  $d = 1$  (big decrease). So:

$$S_j = (1 - \alpha)S_j \quad (8)$$

In our experiments we intuitively set  $\alpha = 0.3$  and  $\beta = 1 - \alpha = 0.7$ . Other values were analysed but either they lead to high number of adaptation iterations for the algorithm to converge, or they minimize its effects. Determining the best values of  $\alpha$  and  $\beta$  and how they affect the stability and the adaptivity of ZAAL is still future work.

## 4.2 Discussion

We present in this section some characteristics of zigzag-avoiding algorithm. To simplify results, we use  $\alpha = 0.3$  and  $\beta = 0.7$ . We can notice the following:

1. At the beginning, average values are set to 1. They are greater than  $\beta = 0.7$ , so there is no restriction using any bitrate, i.e. application is authorised to use all bitrates.

2. Studying the decreasing phase of the algorithm we can notice:

- As mentioned previously, when a bitrate is greater than the available bandwidth, application avoids to use it for a while by decreasing its successfulness average value. Based on equation 8, the first time this occurs,  $S_i = 1 - 0.3 * 1 = 0.7$ . Hence, a bitrate which causes failed writing packets cannot be used two consecutive times.

- The smallest  $S$  value for a bitrate  $i$  can appear when it is used with  $S_i = 0.7 + \epsilon$  and it leads to failed packets. So, its new  $S$  value is equal to  $S_i = 0.7 * (0.7 + \epsilon) = 0.49 + \epsilon$ .

- A higher bitrate does not necessarily mean a smaller  $S$  value. When reducing an  $S$  value for a bitrate  $i$ , higher bitrates do not change.

3. Studying the increasing phase of the algorithm we can notice:

- The maximum number of times ZAAL prevents an application to increase bitrate occurs when the higher bitrate  $k$  has the minimal  $S$  value of  $S_k = 0.49 + \epsilon$  (see previous point) and ends when  $S_k$  becomes  $> 0.7$ . Using equation 7, this corresponds to

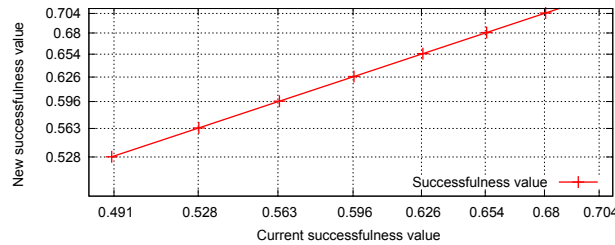


Figure 2: Successfulness average increasing during the biggest period of time.

7 unsuccessful attempts by the application for the higher bitrate  $k$ , followed by an 8th successful attempt. Fig. 2 shows the evolution of this value: there are 7 unsuccessful attempts between  $0.49 + \epsilon$  and a value greater than 0.7.

- $S$  values between 0.7 and 1 are only possible when the bitrate does not cause failed packets in two cases: it increases slowly beyond 0.7 when the quality is stable (see equation 6), or it increases quickly beyond 0.7 when higher qualities are possible (see equation 5).
4. Finally, using an exponential average rather than a linear average allows to remember past values for a longer time when increasing average values  $S$ , and to give more weight to recent history than to old history.

### 4.3 ZAAL Complexity

- ZAAL is a simple and easy to implement algorithm. Indeed, ZAAL consists of 3 if

clauses with a loop over all bitrates inside each clause.

- ZAAL is scalable, since it has a linear complexity in number of bitrates (a loop on all bitrates), and also in number of clients (ZAAL is executed independently for each of them). Moreover, it has a negligible execution time. Indeed, it needs to calculate just one value per available bitrate during each adaptation period (e.g. each 2 sec.), using a simple formula, as shown above. Hence, using ZAAL does not really affect the server processing capacity when the number of clients increases. The small additional impact of ZAAL can be easily compensated by its other positive side (i.e. when using ZAAL, the server avoids processing packets which would otherwise be lost on the network).

## 5 EXPERIMENTAL RESULTS

Fig. 3 shows the real network used to realise experiments with the program presented in the previous section. A video streaming connection is made between a sender and a receiver with wired interface for both. An intermediate machine (called *shaping machine* in the following) is added between the sender and the receiver<sup>2</sup>.

<sup>2</sup>We preferred to add an intermediate machine since making the traffic shaping on DCCP sender is not working currently (see thread “DCCP.BUG called” on DCCP mailing list, available at <http://www.spinics.net/lists/dccp/msg04264.html>), making it on the access point is not interesting either, because it has an




Figure 3: Network topology used for experiments.

This shaping machine has two interface cards: a wired interface connected to the sender and a wireless one (54Mb/s) connected to an access point (AP). The receiver is connected to the same AP via its wired interface. We chose a wireless network in order to be nearer the reality, since wireless networks are commonly used today to access Internet and such networks are more challenging for the algorithms involved in video adaptation. The video streaming uses a real video, available in four qualities 3Mb/s, 2Mb/s, 1Mb/s and 512kb/s, as mentioned before. The video has 180s. Detailed information about experiment parameters are given in appendix A.

Two series of tests are done. For the first series (traffic shaping series), just one flow is present at any moment during the video transmission, which can thus use all the available bandwidth. On the other hand, the output bandwidth of the shaping machine is changed three times to simulate a variable bandwidth: 600kb/s for the first minute, 2300kb/s for the second minute, and traffic shaping was stopped for the last minute (hence the output bandwidth is the original wireless bandwidth). This series allows to see the effect of ZAAL in a network where bandwidth changes over time.

older linux kernel, and making it on the receiver machine does not give accurate results.

For the second series, two numbers of concurrent flows are present during the whole transmission: five and ten flows. In fact, five flows at maximum bitrate are comparable to the bandwidth of the network ( $5 \times 3\text{Mb/s} = 15\text{Mb/s}$ , which is about the bandwidth provided by a classical 54Mb/s wireless network), while 10 flows clearly exceed the bandwidth. This allows to see what happens when multiples flows sense the available bandwidth, especially to check if this leads to a wide oscillation in performance. In fact, as the number of flows increases from five to ten, the number of video adaptation decisions will increase, allowing thus to better verify the usefulness of ZAAL.

All the flows use VAAL [2] as adaptation algorithm. Additionally, either all the flows use ZAAL, or none of them.

The series above consist in executing the same experiments five times. In fact, depending on the network conditions one method can take advantage over the others just because during its experiment the network conditions were better; with five repetitions, the chance for something to happen again and again is very small. The overall average of all experiment repetitions is presented here. Note that there is no retransmission for lost packets in all our tests (as given by DCCP).

This section aims to check if ZAAL avoids zigzag quality switching, and using ZAAL leads to worst, similar or better quantitative results.

## 5.1 Zigzag Avoidance

In this section we present the quality variation with and without ZAAL. In the following figures, the x-axis represents the time from 0 to

180s, the duration of a video transmission, and the y-axis shows the video bitrate (with or without ZAAL).

### 5.1.1 One flow in case of traffic shaping

As expected, ZAAL minimises the zigzag effect: during the first minute in Fig. 4(a) (where ZAAL is not used), the video bitrate is continuously toggling between 0.5 and 1Mb/s (the bandwidth is 600kb/s), while when ZAAL is used, in Fig. 4(b), application uses 0.5Mb/s bitrate most of the time. The same conclusions can be drawn from the second minute. During the last minute, bandwidth is wide enough to support 3Mb/s, and ZAAL finally allows this bitrate.

Further analysis of Fig. 4(b) (with ZAAL) confirms the useful properties of ZAAL algorithm. First, application waits for at least 2 periods of time before trying a bitrate which has recently caused losses (e.g. at the beginning, 1Mb/s did not work between 0s and 2s, hence it was retried not at 4s, but at 6s). Second, when a bitrate causes losses for several consecutive times, application waits more and more time to retry it (e.g. 1Mb/s at 32s, then at 46s, and finally at 64s). Third, the maximum period during which the video quality was prevented to increase is 14s, as shown in section 4.2, point 3 (e.g. the first unsuccessful attempt at 50s followed by the next successful attempt at 64s); during that time there were no losses (bandwidth much higher than bitrate), however ZAAL correctly prevented the bitrate increasing.

More precisely, Tab. 1 presents the number of zigzags during the first and the second minute for the five repetitions done with and with-


out ZAAL (there is no difference during the third minute). It is clear that ZAAL leads to much fewer zigzags in each of the repetitions, for example in the first repetition the number of zigzags is 4 for ZAAL against 13 without ZAAL in first minute, and 0 against 12 in second minute, so 84% less in total. In average for of all the five repetitions, ZAAL leads to between 75% (the third repetition) and 86% (the fifth repetition) fewer zigzags number during the transmission. Naturally, this has a very big impact on the video quality perceived.

### 5.1.2 Second series: five and ten concurrent flows

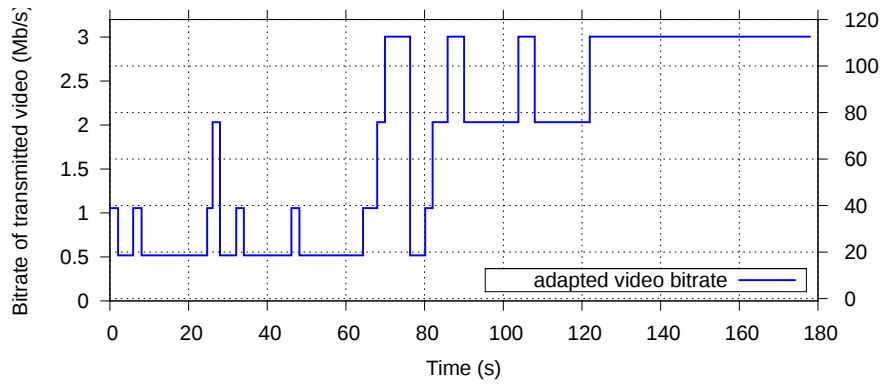
In the second series of tests the bandwidth is not limited in any manner. Results of the five repetitions of these experiments are presented in Tab. 2 for five concurrent flows and Tab. 3 for ten concurrent flows. These tables give the total number of zigzags (the sum for all the flows) during the first, second and third minute of the video transmission. Dividing the time in this way allows to confirm that the zigzags are distributed and not concentrated in one interval of time.

#### Five flows:

In these experiments the server is transmitting five concurrent flows at the same time, and the available bandwidth is theoretically higher than the data rate needed for all flows with the highest available bitrate of our experiments. As the video adaptation method senses the available bandwidth, the highest bitrate is used most of the time (the bitrate is rarely changed) and



(a) without ZAAL: many zigzags occur



(b) with ZAAL: few zigzag occur

Figure 4: Quality adaptation for one flow in case of traffic shaping, under the *same* network conditions.

Table 1: Number of zigzags with and without ZAAL for the five repetitions in case of traffic shaping.

<i>Repetition number</i>	<i>Method</i>	<i>First minute</i>	<i>Second minute</i>	<i>Total</i>	<i>Reduction</i>
1	Without ZAAL	13	12	25	
	With ZAAL	4	0	4	84%
2	Without ZAAL	13	11	24	
	With ZAAL	4	0	4	83%
3	Without ZAAL	13	7	20	
	With ZAAL	4	1	5	75%
4	Without ZAAL	13	12	25	
	With ZAAL	4	1	5	80%
5	Without ZAAL	13	9	22	
	With ZAAL	3	0	3	86%
Total	Without ZAAL	65	51	116	
	With ZAAL	19	2	21	82%

the received video does not suffer from quality switching. Tab. 2 confirms this idea where the total number of zigzags with and without ZAAL is between 0 and 2 for nearly all flows in all repetitions except for flows without using ZAAL in the first repetition where the total number of zigzags reaches 31. This means that when the available bandwidth is higher than the data rate, using a zigzag avoiding algorithm such as ZAAL is not really necessary; it is however useful especially if the algorithm does not consume CPU or network resources, which is the case for ZAAL.

#### Ten flows:

This test is similar to the previous one but with ten concurrent flows; the available bandwidth is smaller than the data rate of all flows with the highest available bitrate. Comparison of zigzags number with and without using ZAAL is presented in Tab. 3. This table shows first that video quality is not stable anymore and the number of zigzags is much greater than the

tests with five flows. Second, using ZAAL leads (as for the first series with traffic shaping) to between 83% and 87% fewer zigzags and about 85% in average when ZAAL is used. As a conclusion, ZAAL minimises the unnecessary oscillation in video quality.

In these experiments with ten flows we noticed that all flows have the same tendency when choosing video bitrates. Fig. 5 shows the tendency of one flow out of the ten concurrent flows. We can see that the bitrate changes often between 1Mb/s and 2Mb/s, which indicates that one flow does not stay all the time at a high bitrate (e.g. 3Mb/s); if this happened, it would reduce the available bandwidth for other flows. Also, like in the previous test with traffic shaping, application adapts often video quality while avoiding bitrates causing lost packets (e.g. in Fig. 5, 3Mb/s bitrate causes lost packets at 26s, then it is not used until 68s). At the same time, it avoids frequent changes in video bitrate during transmission while improving the use of the bandwidth.

Table 2: Number of zigzags (total for all flows) with and without ZAAL for the five repetitions in case of five concurrent flows.

<i>Repetition num</i>	<i>Method</i>	<i>First minute</i>	<i>Second minute</i>	<i>Third minute</i>	<i>Total</i>
1	Without ZAAL	7	7	17	31
	With ZAAL	0	0	0	0
2	Without ZAAL	0	0	0	0
	With ZAAL	1	0	0	1
3	Without ZAAL	0	0	0	0
	With ZAAL	0	0	0	0
4	Without ZAAL	0	0	0	0
	With ZAAL	2	0	0	2
5	Without ZAAL	0	1	1	2
	With ZAAL	0	0	0	0
Avg	Without ZAAL	7	8	18	33
	With ZAAL	3	1	0	3

Table 3: Number of zigzags (total for all flows) with and without ZAAL for the five repetitions in case of ten concurrent flows.

<i>Repetition</i>	<i>Method</i>	<i>First minute</i>	<i>Second minute</i>	<i>Third minute</i>	<i>Total</i>	<i>Reduction</i>
1	Without ZAAL	89	102	94	285	
	With ZAAL	18	6	11	35	87.7%
2	Without ZAAL	83	95	91	269	
	With ZAAL	16	12	11	39	85.5%
3	Without ZAAL	83	91	102	276	
	With ZAAL	20	10	11	41	85.5%
4	Without ZAAL	75	111	103	289	
	With ZAAL	22	7	14	43	85.1%
5	Without ZAAL	101	107	101	309	
	With ZAAL	18	18	15	51	83.5%
Total	Without ZAAL	431	506	491	1428	
	With ZAAL	94	53	62	209	85.4%

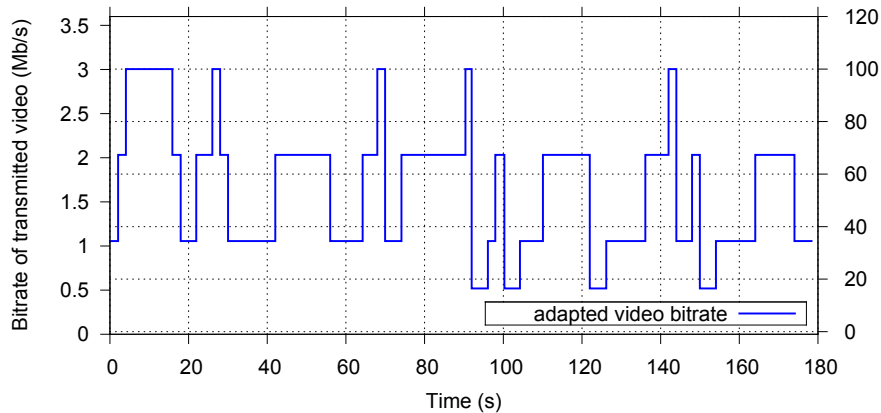


Figure 5: Quality adaptation with ZAAL for flow number 1 in ten concurrent flows test.




Figure 6: Percentage of sent packets by each flow at application layer using ZAAL: the percentages are nearly equal.

This test also checks the fairness of ZAAL. Fig. 6 shows the percentage of sent packets by each flow at application layer when using ZAAL. It shows that ZAAL maintains the fairness among concurrent flows on server (i.e. all flows have nearly equal percentage of sent packets). On the other hand, the fairness on the network is guaranteed by using a TCP-friendly network congestion control.

## 5.2 Performance Comparison

Even if reducing packet loss rate is not the main goal of ZAAL, we investigate how ZAAL affects it. For that, we compare again the same application with ZAAL and without ZAAL. We consider that if a new transmission method, like ours, is able to maximise the number of received packets while minimising the difference between the number of sent and received packets, it will ameliorate the received video quality. Thus, we compare on the one hand the number of received packets, and on the other hand the number of lost packets, in average. It should be noted that, as already known, video quality over network transmission is more affected by packet losses rather than video bitrate value.

In the following figures, the x-axis represents the repetition number, followed by the average of all repetitions. The number of sent and received packets is put on the y-axis (and they use the same point type to distinguish them more easily; of course, the curve of received packets is always lower than the curve for sent packets).



Note that, even if all the curves are put on the same graph, the execution is done at different times. Also, even if the curves use lines for better visualisation, the flows are independent.

### 5.2.1 One flow in case of traffic shaping

Fig. 7 presents results of this series of experiments. It is clear that the number of received packets is smaller when ZAAL is used (by about 6% in average, 40460 received packets for ZAAL compared to 42961 without ZAAL), because of ZAAL preventing high bitrate for some period. On the other hand, when ZAAL is used the flow has in average 38% fewer lost packets (6104 with ZAAL compared to 3819 without ZAAL, under the same network conditions). In general, it is much better to receive 6% fewer packets and have 38% fewer losses, so we can consider that using ZAAL improves the video quality perceived too.

### 5.2.2 Five flows

Results of this test are shown in Fig. 8. This figure shows that ZAAL is clearly better in some experiments (e.g. 1 and 5), while without ZAAL is clearly better in others (e.g. 2 and 4). Looking at the average, ZAAL is slightly better in terms of the number of received packets and packets loss rate. As a conclusion of this test, we can say that when the available bandwidth is sufficient for all flows, the quantitative results are nearly equal with and without ZAAL.

### 5.2.3 Ten flows

The average number of sent and received packets of the ten concurrent flows gives yet better results for ZAAL (see Fig. 9). Except for repetition 2, all repetitions have higher number of received packets when ZAAL is used. Moreover, packet loss rate is smaller for ZAAL in all repetitions, even for 2. In average, number of received packets are about 7% greater in case of ZAAL (32477 compared to 30280), and loss rate is about 57% smaller too (4236 compared to 9825). As a conclusion, ZAAL is better in terms of sent and received packets, avoiding the zigzag in the same time.

From the results of this section we can conclude that ZAAL is better and, even if sometimes its number of received packets is smaller, it reduces the rate of lost packets while maximizing the use of the bandwidth.

## 6 CONCLUSIONS AND PERSPECTIVES

This article has presented a simple brand new method to avoid the undesirable constant switching in quality occurring during a video streaming using content adaptation. It is a general solution, since it can be integrated to any adaptation method, and no matter if the video is encoded in multiple qualities or in multi-layers. The proposed solution uses a history of quality successfulness to infer if a quality should be chosen at a given time. It is a non intrusive method, i.e. it does not change the video transmission and does not need feedbacks from the network. Experiments confirm that with our so-

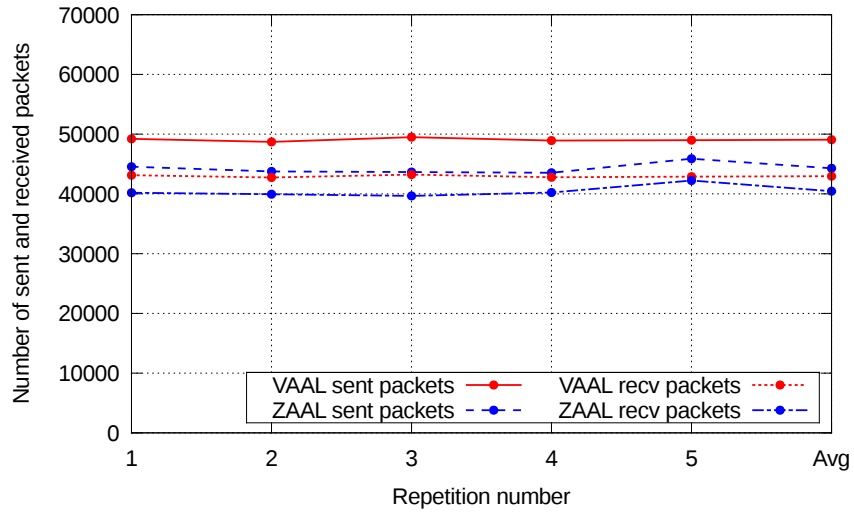


Figure 7: VAAL vs ZAAL in traffic shaping test.




Figure 8: VAAL vs ZAAL in five concurrent flows test.




Figure 9: VAAL vs ZAAL in ten concurrent flows test.

lution the zigzag quality switching appears very rarely, and that it has positive results on video quality.

A perspective of our work is to consider a hybrid solution, which uses a bandwidth estimation method (to decide whether to increase or not the quality) when the maximum period of non-increasing with our method is reached. However, a very interesting future work is to analyse a more general algorithm with VAAL/ZAAL constraints, but applicable to the whole class of network congestion control methods.

## References

[1] Wassim Ramadan, Eugen Dedu, and Julien Bourgeois. Avoiding zigzag quality switching in real content adaptive video streaming. In *International Conference on Digital Information and Communication Technology and Its Appli-*

*cations (DICTAP)*, 1, pages 421–435, Dijon, France, June 2011. CCIS 167.

- [2] Wassim Ramadan, Eugen Dedu, and Julien Bourgeois. VAAL, video adaptation at application layer and experiments using DCCP. In *WPMC, 13th Int. Symposium on Wireless Personal Multimedia Communications*, pages 1–5, Recife, Brazil, October 2010.
- [3] Steven McCanne, Van Jacobson, and Martin Vetterli. Receiver-driven layered multicast. *SIGCOMM Computer Communication Review*, 26:117–130, August 1996.
- [4] Wu-chi Feng. On the efficacy of quality, frame rate, and buffer management for video streaming across best-effort networks. *Journal of High Speed Networks*, 11:199–214, 2002.
- [5] Nick Feamster, Deepak Bansal, and Hari Balakrishnan. On the interactions between layered quality adaptation and congestion control for

streaming video. In *11th International Packet Video Workshop*, Kyongju, Korea, April 2001.

- [6] V. Jacobson. Congestion avoidance and control. *SIGCOMM Computer Communication Review*, 18:314–329, August 1988.
- [7] Robert L. Carter and Mark E. Crovella. Measuring bottleneck link speed in packet-switched networks. *Performance Evaluation*, 27-28:297–318, October 1996.
- [8] Srinivasan Keshav. Packet-pair flow control. *IEEE/ACM Transactions on Networking*, February 1995.
- [9] Constantinos Dovrolis, Parameswaran Ramanathan, and David Moore. Packet-dispersion techniques and a capacity-estimation methodology. *IEEE/ACM Transactions on Networking*, 12:963–977, December 2004.
- [10] Raj Jain and Shawn A. Routhier. Packet trains: Measurements and a new model for computer network traffic. *IEEE Journal on Selected Areas in Communications*, 4:986–995, 1986.
- [11] Jon Postel. Transmission control protocol, September 1991. RFC 793.
- [12] Sally Floyd, Mark Handley, Jitendra Padhye, and Joerg Widmer. TCP Friendly Rate Control (TFRC): Protocol specification, September 2008. RFC 5348.
- [13] ITU-T. Opinion model for video-telephony applications, April 2007.

Parameter name	Parameter value
Experiments place	In building
Packet size	1024 bytes
PC1 (sender): Wired card Product Technology Vendor	82567LM-3 Gigabit Connection Intel Corporation
PC2 (shaper machine): Wireless card Product Technology Vendor	BCM4312 802.11b/g Broadcom Corporation
PC2 (shaper machine): Wired card Product Technology Vendor	RTL8111/8168B Gigabit Connection Realtek
PC3 (receiver): Wired card Product Technology Vendor	BCM5755M Gigabit Connection Broadcom Corporation
PC1,2&3 Wired bandwidth	100Mb/s
PC1,2&3 OS	Linux (Ubuntu 64bits)
PC1,2&3 OS kernel	2.6.35 generic
DCCP	Included in the kernel
Wireless bandwidth	54Mb/s
Distance (AP ↔ PC2)	50cm
Shared with other users	no

Table 4: Experiment parameters.

## A Experiment parameters

The real network parameters used for the experiments are presented in table 4.