

# An Online Meeting Tool for Low Bandwidth Environments

Tresor Mvumbi  
Dept of Computer Science  
University of Cape Town  
Private Bag X3,  
Rondebosch, 7701  
South Africa  
tresormvumbi@gmail.com

Flora Kundaeli  
Dept of Computer Science  
University of Cape Town  
Private Bag X3,  
Rondebosch, 7701  
South Africa  
kndflo001@uct.ac.za

Zafika Manzi  
Dept of Computer Science  
University of Cape Town  
Private Bag X3,  
Rondebosch, 7701  
South Africa  
zafika.manzi@gmail.com

Kyle Williams  
Dept of Computer Science  
University of Cape Town  
Private Bag X3,  
Rondebosch, 7701  
South Africa  
kwilliams@cs.uct.ac.za

Hussein Suleman  
Dept of Computer Science  
University of Cape Town  
Private Bag X3,  
Rondebosch, 7701  
South Africa  
hussein@cs.uct.ac.za

## ABSTRACT

Online meetings allow for remote conferencing and collaborative work among geographically dispersed participants and can save time and expenses that an ordinary face-to-face meeting would require. However, carrying real-time communication within the packet-switched Internet is a challenging task, especially in an African context, which is characterized by low bandwidth and unstable Internet connections. This paper presents and evaluates a tool that was designed to enhance the user experience for Web-based conferencing, given the constraints of Internet conditions typical of Africa. Approaches used to achieve this goal included: reprioritisation of multimedia streams, image differentiation, half duplex communication mode and stream compression. It was found that less than 56 kbps of bandwidth was required in order to: transmit audio; use video to convey presence; share slides and screen; and support text-based chat and floor control. Furthermore, users were largely satisfied with the tool and felt that it created a good user experience.

## Categories and Subject Descriptors

H.4.3 [Information Systems Applications]: Communications Applications; H.5.2 [Information Interfaces and Presentation]: User Interfaces

## General Terms

Design, Experimentation, Performance, Reliability

## Keywords

Web-based meeting, video conferencing, low bandwidth

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAICSIT '12, October 1-3, Centurion, Tshwane, South Africa  
Copyright 2012 ACM 978-1-60558-950-3 ...\$10.00.

## 1. INTRODUCTION

Internet technologies have developed rapidly in recent decades and are currently mature enough to support low-cost, real-time communication services. These developments have literally changed the way people meet and collaborate to make decisions [3]. New collaborative and conferencing environments based on the Internet are now common tools for many people and organizations around the world.

A Web meeting system is an Internet-based tool offering a virtual environment for remote meeting and collaborative work among geographically dispersed participants and can be used to avoid travel expenses and time required for face-to-face meetings [11]. Some of the common features offered by Web meeting systems include: audio and video communication; slide show presentation sharing; screen sharing; and text-based chat [13].

The Internet provides a public packet-switched network with a relatively high probability of loss and random delay in packet delivery [9]. These transmission problems directly affect any service relying on the Internet for communication. Another important factor affecting Internet services is the amount of bandwidth available, where bandwidth is a measurement of data quantity that a link can transmit per unit of time. This factor is particularly important in Africa and developing countries where the amount of available bandwidth is relatively low. Despite the growth in use of ICTs on a global scale, Internet access is still limited in most African countries [1]. It was estimated that, in 2007, there was a total of 43 Gbps of international bandwidth in Africa, 80% of which was devoted to North African countries and South Africa [20]. Putting this into perspective, the total international connectivity in Sub-Saharan Africa in 2007 was less than one third of that in India [20]. In recent years, however, Africa has realised substantial growth in international connectivity, which is currently estimated to be approaching 1 Tbps [17]. However, while Africa has experienced an increase in international bandwidth in recent years, the amount of bandwidth still remains far below that in the rest of the world. Furthermore, in regions where there is poor fixed line telecommunications infrastructure, people may rely

on mobile networks for Internet connectivity, for instance, using the General Packet Radio Service (GPRS) and Enhanced Data Rates for GSM Evolution (EDGE), which provide data rates of between 56-114 kbps and 236.8-473.6 kbps respectively. As a service relying on the Internet, a Web meeting tool is directly impacted by underlying networking problems such as those that occur as a result of limited bandwidth.

Web meeting tools offer several features that are affected by underlying network problems in different ways, such as:

- **Audio conferencing:** unpleasant or even unintelligible sound playback.
- **Video conferencing:** blocking and jerky video playback.
- **Screen and image sharing:** poor image quality.
- **Text chat:** delayed delivery of messages.

Under low bandwidth conditions these problems can seriously degrade the communication quality, making an Internet conferencing solution practically useless. This paper presents and tests a set of design choices that can notably enhance the quality of communication of Web conferencing tools in low bandwidth environments, such as those that occur in many parts of Africa and in developing countries. The features or services offered by Internet meeting tools have different needs in terms of bandwidth usage and, consequently, they are not all affected the same way by networking problems. Since the primary objective of a Web meeting tool is to support human communication, special emphasis needs to be put on user experience and satisfaction. This paper thus focuses on how the user experience can be positively enhanced despite networking problems, with a specific focus on the priority that is assigned to services and features in order to offer the best tradeoff between quality and utility.

The rest of this paper is structured as follows. Section 2 presents the background to this study and discusses some related work. Section 3 presents the design of a tool for Web-based meetings in low bandwidth environments, including a description of the approaches used for stream recording, compression, transmissions and for congestion control. Section 4 presents the evaluation of the system from a bandwidth usage perspective as well as from a user perspective and, lastly, Section 5 concludes the paper and discusses possibilities for future work.

## 2. BACKGROUND

There are a number of commercial Web meeting conferencing systems available, such as GoToMeeting<sup>1</sup>, Adobe Connect<sup>2</sup>, BeamYourScreen<sup>3</sup>, WebEx Meeting Center<sup>4</sup> and GoMeetNow<sup>5</sup>. While these commercial tools may work well in environments with large amounts of available bandwidth, they are not necessarily suitable for the type of Internet connectivity that typically exists in Africa and developing countries.

<sup>1</sup><http://www.gotomeeting.co.uk/fec/>

<sup>2</sup><http://www.adobe.com/products/adobeconnect.html>

<sup>3</sup><http://www.beamyourscreen.com/>

<sup>4</sup><http://www.webex.com/products/web-conferencing.html>

<sup>5</sup><http://www.gomeetnow.com/>

## 2.1 Transmission Rate Adaptation

Several approaches can be taken in order to support real-time Web-based conferencing in low bandwidth environments or under poor network conditions. A common approach is to use transmission rate adaptation, in which the stream rate is adapted to suit the bandwidth capacity in order to ensure the best Quality of Service (QoS) [16]. The main challenge in doing this is supporting heterogeneous environments where meeting participants have different bandwidth capacities [12]. The simplest approach to transmission rate adaptation is source-based adaptation, in which a uniform representation of the signal is sent from the source at a fixed rate [8]. However, this approach does not work well in heterogeneous environments since areas with low bandwidth capacity suffer from congestion, whereas high capacity areas are underutilized [16]. An alternative to source-based adaptation is receiver-based adaptation in which several signals are broadcast and the receiver, depending on its available bandwidth, will receive an appropriate stream [2]. Two approaches can be used for receiver-based adaptation: the simulcast model in which the receiver selects a single flow channel based on its available capacity [12]; or a multilayer model in which multiple flows can be incrementally combined in order to provide progressive refinement of the signal [16].

## 2.2 Stream Compression

Signal compression can also be used in low bandwidth environments in order to reduce the amount of bandwidth used and thus improve the user experience. Many of the components that make up the signal in a Web-based meeting could benefit from compression, such as video, audio, screen sharing and slide sharing.

Video often accounts for a large percentage of the bandwidth required to transmit a signal with minimal interruption and often benefits the most from compression. A number of video compression schemes exist, such as the MPEG-4 face animation standard, which achieves high compression rates by transmitting only face model parameters [6] and H.26x, which uses waveform-based compression [7]. Shortcomings of these approaches are that it is difficult to make synthesised faces look natural when using the MPEG-4 face animation model and H.26x encoding is not very efficient in producing low bitrate video [6]. These two techniques were combined in order to create a low bitrate face video streaming system where prior knowledge about faces was incorporated into the waveform-based compression in order to improve compression [21]. In a solution proposed by Cohen et al [6] for low bitrate face video transmission, the encoder selected only a few good quality faces, which were then compressed and transmitted, while the decoder used image morphing-based rendering in order to generate a normal video rate. Additional standards for low bandwidth video compression exist, such as the Discrete Cosine Transform (DCT) [6].

Lower video frame rates can also be beneficial in low bandwidth environments. For instance, if video is only used to provide presence, then one frame every five seconds could suffice [19]. However, if complex emotions need to be portrayed then one frame every five seconds is insufficient. Thus, in limited bandwidth environments it may be suitable to make use of a video stream only to convey presence if there is insufficient bandwidth to transmit video at a higher frame rate.

### 2.3 Duplexing Mode and Floor Control

The duplexing mode could also potentially play a large role in the performance of an online meeting system in low bandwidth environments. In full duplex mode, data is transmitted in both directions and thus parties can both receive and send data simultaneously. This is in contrast to half duplex mode where data can be transmitted in both directions, however, only in one direction at a time, effectively halving the bandwidth use. Floor control often plays an important role in online conferences where participants may be given exclusive or non-exclusive permission to transmit and receive data. When in full duplex mode, floor control could be used to prevent all meeting participants from speaking at the same time or to allow certain participants to share their screen with other participants. Floor control is, arguably, more important in half duplex mode where participants may need exclusive speaking rights. In both cases, however, floor control is likely to lead to a decrease in bandwidth use since transmission is controlled to some extent. There are two types of floor control mechanisms available: receiver-based floor control and sender-based floor control [10]. When there is no floor control, all participants in a meeting transmit their signal and receive and process the signal of all other participants. In receiver-based floor control, all participants in a meeting transmit their signal, however, receivers only process the signal of the participant who has the floor and, in sender-based floor control, only the signal of the participant who has the floor is transmitted [10]. Management of the floor can either be by a moderator or automatically using a mechanism such as a turn taking protocol [5]. Furthermore, mechanisms like hand raising can be used in order to request access to the floor [15].

### 2.4 Additional Resources for Communication

In addition to the expected audio and video channels, additional resources can be used in video conferencing. For instance, slideshows often play an important role in a presentation. There are two common cases for retrieving slides in a presentation: retrieving all slides before a presentation begins or retrieving slides on demand as they are needed [22]. Retrieving slides on demand in an environment with limited bandwidth may not be suitable since it may congest the network and result in delay. Thus, pre-retrieving slides may be a more appropriate approach. An alternative approach was proposed by Yang [22] where a just-in-time retrieval policy was specified that required the transmission of an object to be completed just before the object was displayed and required that the time for retrieving an object be estimated in order to ensure timely transmission [22].

In low bandwidth environments, text-based chat provides an alternative means of communication. For instance, it could be useful to combine video, audio and text-based chat so that, when one medium is unavailable due to low bandwidth, mediums with lower bandwidth requirements can be prioritised in order to continue communication. Scholl et al [19] made use of text-based chat to complement the video (video-chat) as opposed to the common approach of using audio and video and it was found that most users found the application useful. Furthermore, in another study, Scholl et al [18] showed that, even when bandwidth is not an issue, text-based chat can have benefits over audio chat, such as: lowering the cost of interrupting others; making it easier to communicate in a second language; eliminating the effect

of background noise, especially in public places; making it easier to communicate in larger groups; and enabling asynchronous communication.

Lastly, screen sharing can also be useful in Web-based meetings to complement presentations, i.e., to further demonstrate an idea or show an example of a product. However, most screen sharing applications do not work well with low bandwidth as the presenter's screen image has to be streamed across the network [14], thus motivating the need for screen sharing techniques that are more appropriate for low bandwidth environments.

This section has discussed some of the components that make up a Web-based meeting system and also discussed some of the difficulties that exist in low bandwidth environments and some common solutions that are used to overcome them. In the next section, a tool that was designed for Web-based meetings, with a specific focus on it being usable in low bandwidth environments, will be described.

## 3. SYSTEM DESIGN

The design choices presented in this section focus on how to provide a good user experience within limited bandwidth conditions that are typical of Africa. Web meeting tools offer a broad range of features that are affected differently by networking problems. As a general rule of thumb, features with high real-time bandwidth requirements are the most affected. These features include: audio streaming, video conferencing, screen and presentation slide sharing. The rest of the section will focus on techniques and design choices that minimise bandwidth usage for these features, while maintaining an acceptable user experience. Features like text chat, polling or floor control mainly rely on exchange of text messages and, thus, are less affected by bandwidth problems. The specific objective is to design a meeting tool that delivers:

- An audio stream conveying clear speech, while using the smallest amount of bandwidth possible;
- A video stream that provides a good sense of presence and improves the user experience; and
- Desktop and presentation sharing features that minimise bandwidth usage.

### 3.1 Overall Architecture

The system design is based on the Client-Server model, where clients initiate communication by requesting services from the server and the server provides services to one or more clients. The roles of the clients and server are summarised below:

- Clients: record streams (audio, video and screen), compress streams, send packets to the server, receive packets from the server, decompress and play back the streams.
- Server: receives and buffers packets coming from a specific client and broadcasts packets to the rest of clients.

A key design decision in order to minimise bandwidth usage is to only allow a single participant to present at a time, while the others follow. This enables half duplex communication, which effectively reduces bandwidth usage by a factor of two. Figure 1 shows how this mode of communication works, with one client machine sending a

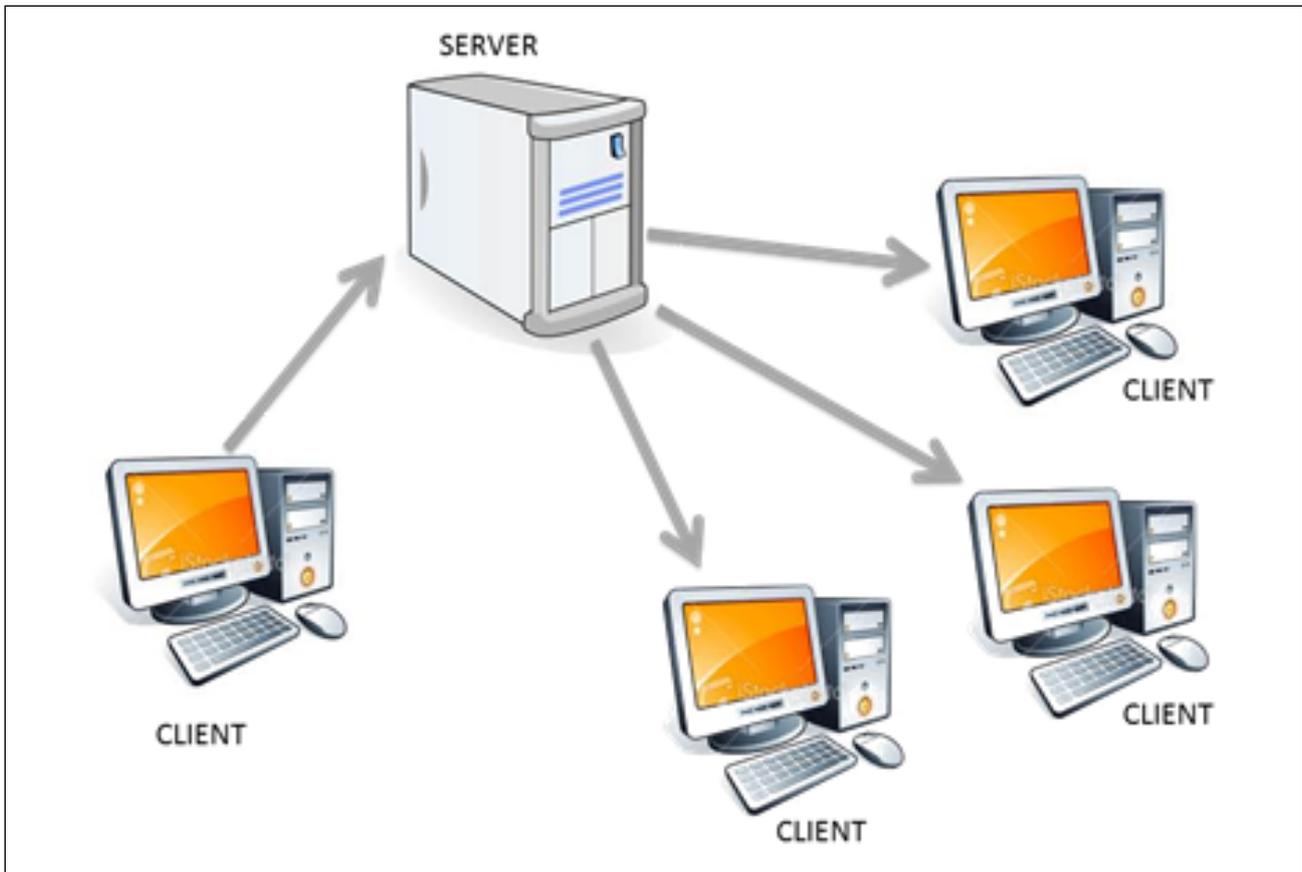


Figure 1: Half duplex client-server communication

signal to the server, which then broadcasts to the other clients.

In the next sections, the design of the components that make up the system will be discussed, while emphasising the design decisions related to minimising bandwidth usage.

### 3.2 Audio Conferencing

The audio conferencing component was designed to deliver an audio stream that is clear enough to convey intelligible speech, while using the smallest amount of bandwidth possible. Audio communication involves recording, compression, transmission, decompression, buffering and playback. These are described below.

The conversion of analog sound into a digital signal requires the analog sound to be sampled, which involves dividing the time axis into a number of discrete blocks called samples, and quantization, which involves dividing the vertical axis, which represents the signal strength, into several discrete levels. The number of samples and quantization of each sound sample directly determine how close a digital signal is to its analog equivalent. The bit rate is a measure of the number of bits that are transmitted per unit of time and provides an indication of the quality of a digital recording. Sound bit rate is given by:

$$\text{Bit Rate} = C \times SR \times SS, \quad (1)$$

where  $C$  is the number of channels,  $SR$  is the sample rate and  $SS$  is the sample size. In this study, audio was recorded for a single channel at a sample rate of 8000 samples per second and a sample size of 8 bits, resulting in a sound bit rate of 64 kilo bits per second (kbps).

Thus, bandwidth of at least 64 kbps is needed for transmission of the uncompressed audio stream, which is, arguably, too much for a low bandwidth context. The solution to this problem is stream compression, which reduces the quantity of data needed. Different algorithms and formats have been proposed for audio compression, including: A-Law, M-Law, MP3 and Groupe Speciale Mobile (GSM) [9]. These formats were tested during prototyping, and it was found that they substantially degraded the sound quality for bit rates below 32 kbps. Therefore, the sound packets were compressed using the ZIP format. This format can compress without data loss, resulting in a clear and sharp sound quality. The compressed stream uses 16 kbps on average, which represents a 75% decrease in the stream size. In addition, ZIP produces different packet sizes based on the stream's content, such that, when there is silence, the bit rate drops below 8 kbps. Figure 2 summarises the main steps from audio recording to streaming with the parameters used in this study.

The participant that is presenting transmits the compressed audio signal to the server, which then broadcasts the compressed signal to the listening participants. The audio packets are decompressed and buffered before playback and a dynamic buffer size is used. Figure 3 illustrates the process, from reception to playback.

### 3.3 Video Conferencing

The raw video stream from a camera is not compressed and usually has a natively high bit rate (around 18 Mbps for a 320X240 video at 15 frames per sec). A bit rate this high is obviously not feasible in a low bandwidth environment. Similarly, even when the video stream from a cam-

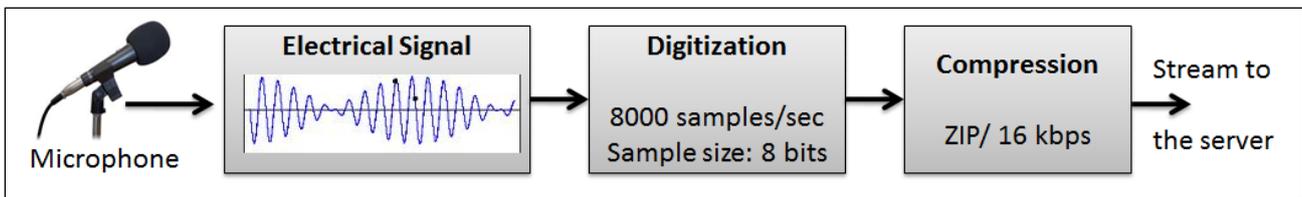


Figure 2: Audio recording, compression and transmission

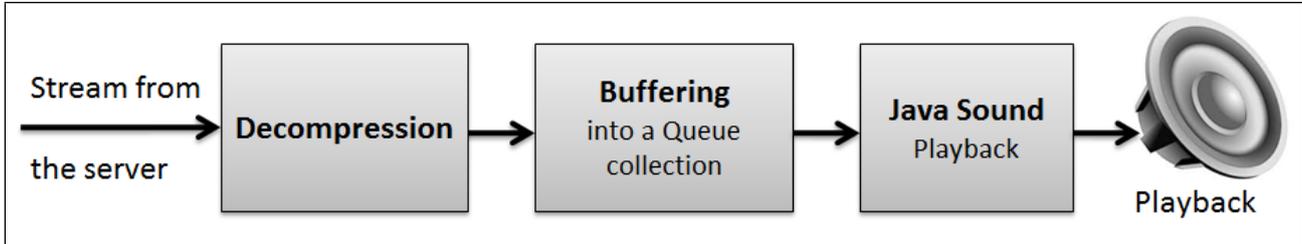


Figure 3: Audio reception, decompression and playback

era is compressed, the bandwidth requirements often remain too high for use in limited bandwidth environments. Thus, in this study, the decision was taken to use video only to convey presence by transmitting single images intermittently. To have full control over the compression and stream process and to achieve a greater compression rate, each image is recorded individually. This technique allows for the delivery of even very low frame rates (far below 1 frame per second) and allows for at least a certain sense of presence to be conveyed when the bandwidth available is very limited.

In a meeting context, images coming from the camera are often very similar due to a speaker usually being still for most of the meeting. This fact was exploited in order to achieve better compression. Instead of transmitting a completely different image each time, the decision was made to calculate, compress and transmit only the differences between two successive images, where the difference is calculated as the difference in the RGB values for corresponding pixels. JPEG compression is used and, once the compressed difference between images reaches the destination, the original image is reconstructed by aggregating differences. This process eventually degrades the image quality and, to avoid complete degradation, a key frame is sent after a certain number of iterations, as illustrated in Figure 4.

### 3.4 Slide Sharing

Slide shows often play an important role in presentation to clarify meaning or provide additional information. As such, the sharing of slides is often an important feature in a Web-based conferencing system. As mentioned in Section 2, there are two common approaches that can be used for slide sharing: loading of slides on demand or pre-loading of slides. In a limited bandwidth environment, the loading of slides on demand may be infeasible as it may lead to network congestion. Thus, in this study the latter approach was taken where all slides are downloaded before an online meeting starts. In order to synchronise the slideshow position for all meeting participants, whenever the presenter changes a slide, the index of the new slide is broadcast to all participants, thus ensuring that the slide each participant is viewing remains current. This approach almost nullifies bandwidth usage

during the meeting for slide sharing.

### 3.5 Screen Sharing

Screen sharing allows the presenter to broadcast their desktop to other participants while presenting. This feature is very bandwidth intensive, as it requires regularly sending across several relatively large screenshots. The design aim was thus to find the best tradeoff between image quality and bandwidth required. During a meeting, only one participant can present at a time. Therefore, in order to reduce bandwidth usage, only the presenter is able to broadcast their desktop. To further reduce bandwidth usage, the image differentiation technique is once again used. This approach takes advantage of the fact that consecutive screen images are quite similar and therefore only the parts of the image that have changed are transmitted across the network. Screen images are often made up of large empty areas with the exact same colour, thus motivating the choice of GIF as a compression scheme that achieves good compression with acceptable quality.

### 3.6 Floor Control

Floor control could play an important role in a meeting, especially when the meeting is conducted in half duplex mode. In the tool developed, access to the floor is controlled by the meeting host. Furthermore, the system also supports hand raising as well as the use of emoticons in order to convey a participant's opinion on an issue being discussed, as shown in Figure 5. Lastly, the tool also supports polling.

### 3.7 Bandwidth Control and Mitigation Techniques

Network problems lead to interference in packet delivery. For instance, these problems may include: connection interruptions; delays; and a decrease in the available bandwidth. Detecting packet delivery problems and evaluating their level can help to trigger appropriate counteractions and allow for streams to adapt to network conditions. Evaluating the actual bandwidth available between two nodes on the Internet is a complex task [4]. The bandwidth is often evaluated by uploading and/or downloading a file with a given size and determining the time it takes. However, in the context of low bandwidth, this approach

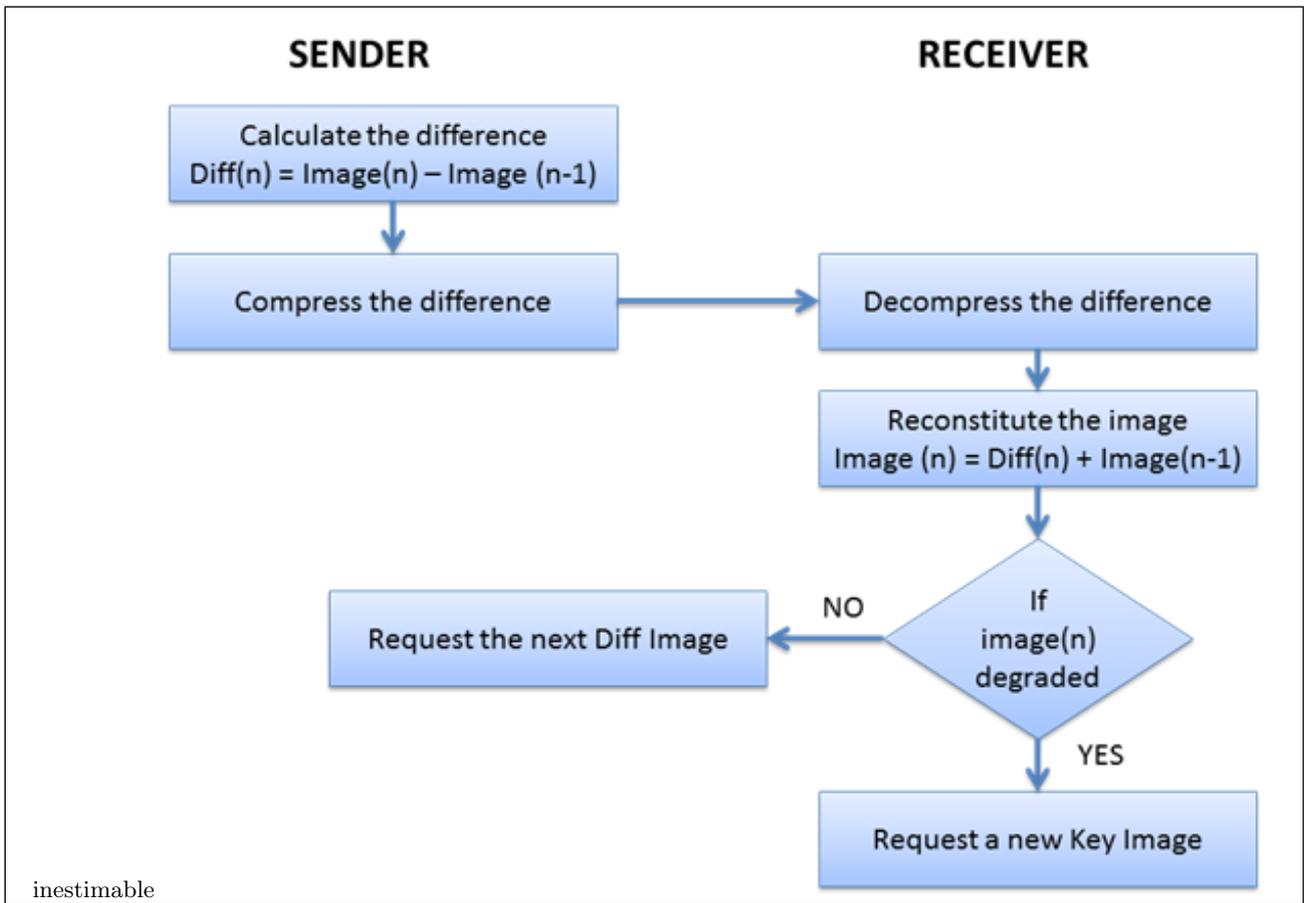


Figure 4: Image differentiation algorithm

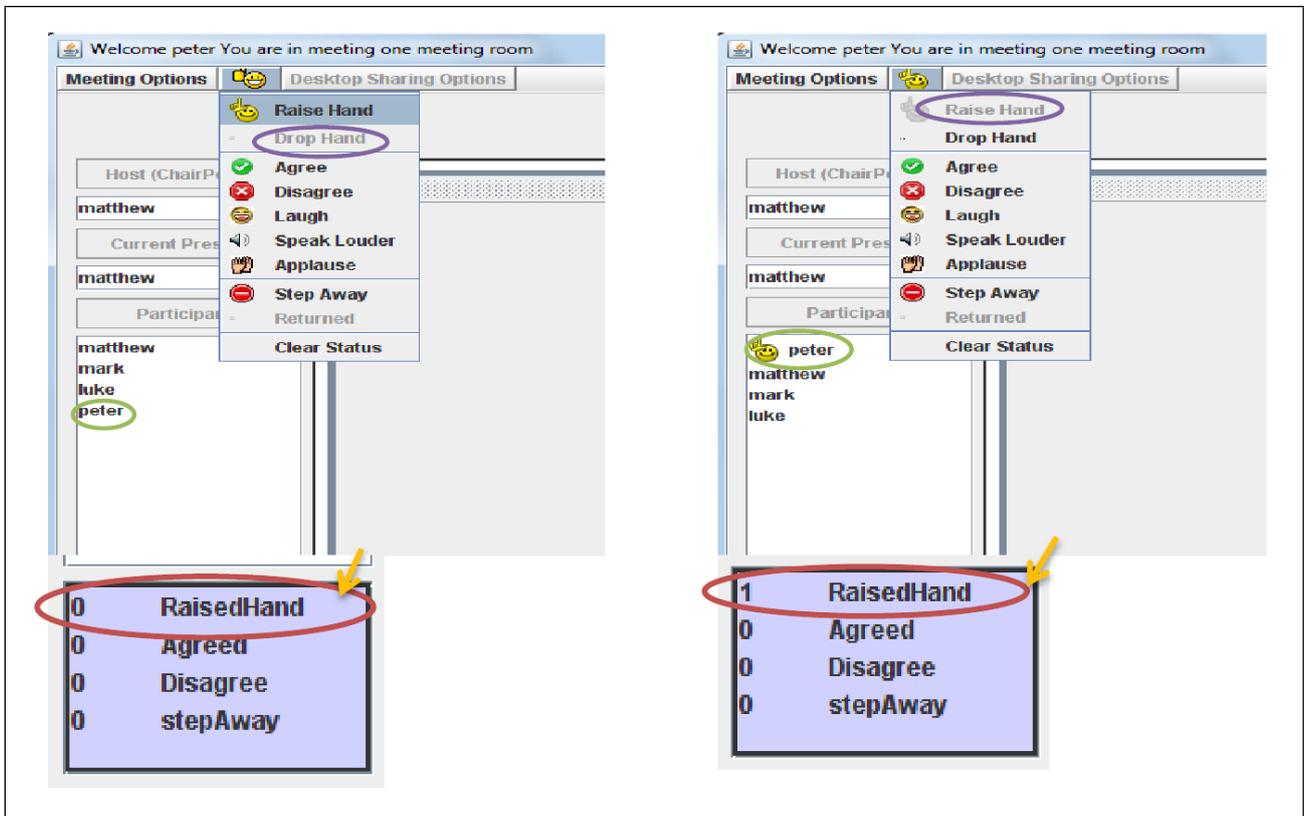


Figure 5: Hand raising

is inefficient as it consumes the limited network resources.

The bandwidth and congestion control proposed in this paper is based on monitoring sound data packet delivery. The sound stream is delivered at a constant pace of one packet per second. The receiver records the time when each packet is received. In the ideal network, the time span between consecutive deliveries should constantly be one second and, when this time increases far above one second, it is likely that something is affecting packet delivery, such as network congestion. There is no perfect network and, even within a Local Area Network (LAN), a small variation in delay at any time is normal. Thus, in order to make the congestion detection system more robust, the average delay over the last 10 packet transmissions is calculated.

When the average delay goes above 3 seconds, the system automatically stops the stream with the lowest priority. The default priority order is (from highest to lowest): floor control, text chat, audio, video, slides sharing and desktop sharing. When the audio stream itself is stopped, the system will try restarting it after a given time and check for congestion. Should no congestion be detected for a given time, the system automatically adds streams with lower priorities and continues to monitor for congestion.

This section has described the design of an online meeting tool for low bandwidth environments. The design considerations for each component that the system is made up of were discussed, with a specific focus on how they have been designed in order to minimise bandwidth use. In the next section, the evaluation of the system is described, including an evaluation of bandwidth usage, as well as a user study.

## 4. EVALUATION

The tool described in the previous section was evaluated in terms of its bandwidth usage. In addition to this, the effect that the bandwidth mitigation strategies had on the user experience was also evaluated by means of a user study. These are presented below.

### 4.1 Bandwidth Usage

The purpose of the bandwidth usage evaluation was to determine the amount of bandwidth used by each of the components that make up the meeting tool in order to determine if they were feasible for use in a limited bandwidth environment. Bandwidth usage was calculated by logging all transmission between the server and each of the clients.

#### 4.1.1 Audio Conferencing

Audio is often one of the most important aspects of a meeting and thus, after the features that use a negligible amount of bandwidth, such as text chat and floor control, was given the highest priority. The bandwidth usage of the audio sub-system was evaluated by simulating a meeting with multiple participants. The meeting, which took place in half duplex mode, began with two participants and, every minute, a new participant joined the meeting until there were a total of eight meeting participants. Figure 6 shows the average amount of bandwidth used by each participant in kbps.

The average bandwidth usage per user was about 16 kbps. Since half duplex mode was used, each participant can only either be sending or receiving an audio stream at any given time. Thus, new participants joining the

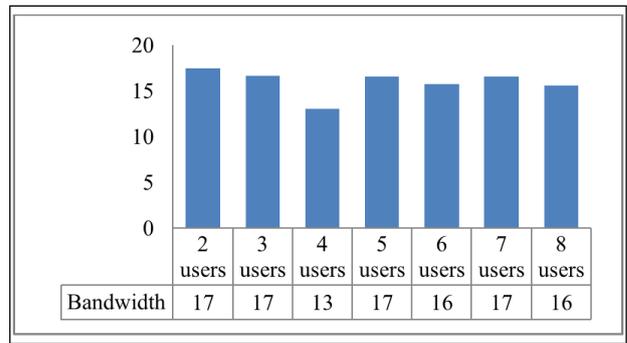


Figure 6: Average bandwidth used per user (in kbps) for a meeting where a new participant joined every minute in half duplex mode

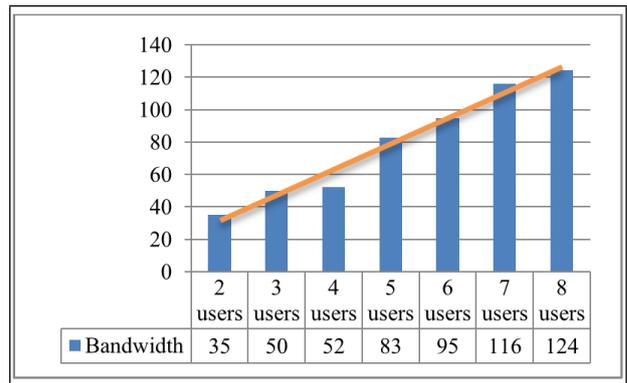


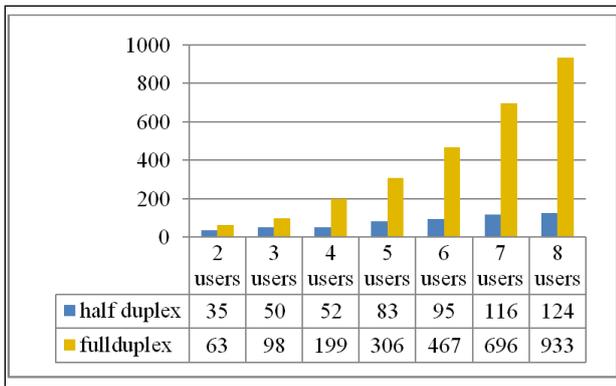
Figure 7: Average bandwidth used on the server (in kbps) during an audio meeting where a new participant joined every minute in half duplex mode

meeting has no effect on the bandwidth requirements of the other participant. However, since the server receives the audio stream from the speaker and then broadcasts it to the remaining meeting participants, its bandwidth requirements scale linearly with the number of participants, as is shown in Figure 7.

In order to minimise bandwidth usage, the system was designed to use half duplex mode. Figure 8 compares half duplex mode to full duplex mode and shows why this is a beneficial approach. In the figure, half duplexing clearly leads to a large reduction in bandwidth use. Furthermore, it also shows that while additional participants in full duplex mode leads to a large increase in bandwidth usage for participants, the increase is only very slight for half duplex mode. The benefits of half duplexing are clear and it is felt that this is a reasonable approach since, in a meeting, there is usually only one speaker speaking at a time and the use of floor control and other mechanisms can be used in order to support the flow of discussion.

#### 4.1.2 Video Conferencing

As mentioned in Section 3.3, it was decided that, in order to minimise bandwidth usage, video would only be used to convey presence through the use of still images whose frequency is dependent on the frame rate. Furthermore, in order to further decrease bandwidth usage, a compression scheme was devised that only transmits the differences between successive images. This compression scheme was evaluated in order to determine the effect that



**Figure 8: Comparison of server bandwidth (in kbps) usage between half and full duplex mode for an audio meeting where a new participant joined every minute**

**Table 1: Bandwidth use in kbps with and without compression**

	No Compression	Compression
Minimum	95	98
Maximum	93	61
<b>Mean</b>	<b>94</b>	<b>68</b>

it had on the bandwidth usage. Table 1 shows the maximum, minimum and mean bandwidth usage with and without compression at a frame rate of 1 frame per second, where the mean represents the average of all packets.

As can be seen from Table 1, the compression scheme is beneficial, leading to a 28% decrease in bandwidth usage.

To evaluate the video stream bandwidth requirements, an experiment was conducted under 3 different frame rate conditions: low frame rate (1 frame every 5 seconds), medium frame rate (1 frame per second) and high frame rate (3 frames per second). The results of this experiment are shown in Table 2.

As can be seen from Table 2, for the lowest frame rate, the resulting video stream uses only 15 kbps, which is smaller than the average audio stream (16 kbps). Such a stream can help to convey a certain sense of presence for meetings where the available bandwidth is very limited.

#### 4.1.3 Slide Sharing

Bandwidth usage evaluation for slide sharing is straightforward: the presentation slides are uploaded to the server then downloaded onto other clients before the meeting starts. Therefore, the total bandwidth usage for each client corresponds to the size of presentation slides. The server needs to broadcast received slides to the rest of the participants, therefore the bandwidth used at the server is equal to the slide size times the number of meeting participants. During the actual meeting, slide sharing uses less than 1 kbps for sharing the ID number of the slide

**Table 2: Maximum, minimum and mean bandwidth (in kbps) usage for 3 different frame rates**

	0.2 FPS	1 FPS	3 FPS
Minimum	107	98	294
Maximum	0	61	64
<b>Mean</b>	<b>15</b>	<b>68</b>	<b>181</b>

currently presented.

#### 4.1.4 Screen Sharing

To determine the bandwidth requirements for screen sharing, an experiment was run with 3 meeting participants (1 presenter and 2 listeners). It was found that, in order to transmit a screen shot of the presenter’s screen at 1 frame per second, an average of 22 kbps was required per participant. Without using the compression algorithm devised as part of this study, a total of 96 kbps was required, clearly demonstrating that the compression algorithm was beneficial. Furthermore, as was the case for video streaming, the bandwidth requirements could significantly be reduced by reducing the frame rate. However, this was not done since, whereas the purpose of the video stream was only to convey presence, screen sharing is often used to demonstrate ideas and thus a higher frame rate is often required.

## 4.2 User Evaluation

It has already been discussed how limited bandwidth can negatively impact a Web-based meeting tool. Thus, design choices were made in order to minimise the negative effects of poor underlying network conditions, while still making an online meeting possible. However, these design choices could potentially impact the user experience when using the Web-based meeting tool. Thus, a user evaluation was conducted in order to gain insight into the effect that the design choices had on user experience and satisfaction.

#### 4.2.1 Audio Conferencing

In order to evaluate the effect that the design choices had on audio conferencing, 13 users were recruited to participate in the user study. Audio was delayed by 1.5 seconds to allow for buffering and the users were asked to rank the sound quality on a scale of 1 to 5, with 1 being poor and 5 being excellent. The average ranking for the sound quality was 4. Furthermore, 6 of the 13 users said that they noticed the delay due to buffering; however, only 1 of these users reported it to have had a negative impact on the meeting experience. Overall, the audio experience was reported to be either good or excellent by 84% of users.

#### 4.2.2 Video Conferencing

The effect of the design choices for video conferencing on the user experience was evaluated using the same 13 users as for the audio conferencing. The video stream was evaluated at 3 different frame rates: a low frame rate (0.2 FPS), an average frame rate (1 FPS) and a relatively high frame rate (3 FPS). Using the same scale of 1-5, users ranked the video quality as fair when the low frame rate was used. However, interestingly, 75% of the users agreed that the video conveyed either a good or strong sense of presence to the virtual meeting. Naturally, the overall appreciation increased for higher frame rates. At 3 FPS, all of the users reported the video stream to be either good or excellent.

#### 4.2.3 Slide and Screen Sharing

For slide sharing, the slides are uploaded beforehand by the meeting administrator. This task can be done several minutes before the meeting starts, allowing plenty of time for the upload and download of presentation slides for even slow connections. During an actual meeting, it took less than a second for the slide change to be reflected on all

participant computers. Therefore, as expected, none of the users reported a delay or any other problems on slide changing during presentation.

Six users evaluated screen sharing and the presenter's desktop was streamed at 1 image per second. This setting required an average bandwidth usage of around 20 kbps per user. Most of the users (more than 80%) who evaluated the system appreciated the responsiveness and screen image quality produced by GIF compression. The use of handraising and emotions were also very well received by the users.

### 4.3 Discussion

The analysis of bandwidth usage showed that half duplex communication is more efficient than full duplex communication. This fact is particularly true for a meeting context, where communications are less interactive compared to a phone call, for example. Image differentiation helped to substantially reduce bandwidth usage for video streaming and desktop sharing by sending across only the differences between successive images. An experiment indicated an average reduction of 28% in bandwidth usage when image differentiation is used. This reduction increases when higher frame rates are used, as successive images are more similar. Combining image differentiation with low frame rates can result in a very light video stream. When the image is updated every 5 seconds, the resulting video stream requires only 15 kbps ( 2 KB/s), which is about equal to the audio stream bandwidth. Pre-loading of slides avoids streaming the presentation in real time and, thus, the bandwidth required for slide sharing during the meeting is almost null. In addition, changing the slides during the meeting is very fast and it takes less than a second to synchronize clients. This result is possible because only the slide ID number is sent across the network and the actual content is loaded locally.

Putting this into perspective, the features can all be combined as follows:

- A clear audio stream (radio quality) at 16 kbps;
- A presence video stream at a low frame rate video (0.2 FPS) at 15 kbps;
- A desktop stream at 22 kbps (at 1 FPS);
- Slide sharing at less than 1 kbps;
- Chat and floor control at around 1 kbps;

Thus, the total bandwidth requirement for a single meeting participant could be estimated to be around 55 kbps, thereby allowing someone to participate in a meeting while on a slow 56 kbps dial-up connection. Similarly, services can be dropped when the amount of bandwidth available decreases, thus still allowing for the meeting to take place.

The user study showed that users were satisfied with the system, rating the audio as being good on average and also noting that the video aspect was successful at conveying a sense of presence. Furthermore, users liked the mechanisms that had been developed for floor control and that supported the use of half duplex communication.

## 5. CONCLUSIONS

There is currently a global demand and need for real time communication and collaborative tools via the Internet. This need of Internet real-time and multimedia communication is also true for Africa and developing countries

where the amount of available bandwidth may be limited. This paper has described an investigation into the feasibility of an Internet-based meeting tool that can provide a satisfactory user experience with limited bandwidth. Approaches used to address this challenge included:

- Prioritisation of features: floor control, text chat, audio, video, slides sharing, desktop sharing (respectively from the highest to lowest priority).
- Pre-loading of presentation slides.
- Implementation of an image differentiation algorithm to reduce the size of images transmitted.
- Use of half duplex mode instead of full duplex mode.
- Compression of the audio and video streams.

The above approaches guided the development of an experimental Web meeting prototype. The objective of the prototype was to deliver a good user experience while using the lowest bandwidth possible and coping with networking problems. The system was shown to allow someone to participate in a meeting with less than 56 kbps of bandwidth available, making it suitable for use on a dial-up Internet connection using a 56 kbps modem, or on a GPRS or EDGE mobile data network. Furthermore, the ability to prioritise streams meant that, even when participants had less bandwidth, they could still participate due to feature prioritisation.

This study has shown that it is possible to provide a fairly good Internet meeting experience within limited bandwidth environments. Future work seeks to investigate other ways to further minimise the amount of bandwidth that is required while continuing to enhance the user experience.

## 6. REFERENCES

- [1] G. A. Ajuwon and L. Rhine. The level of Internet access and ICT training for health information professionals in sub-saharan Africa. *Health Information & Libraries Journal*, 25(3):175–185, 2008.
- [2] E. Amir, S. McCanne, and R. Katz. Receiver-driven bandwidth adaptation for light-weight sessions. In *Proceedings of the fifth ACM international conference on Multimedia*, MULTIMEDIA '97, pages 415–426, New York, NY, USA, 1997. ACM.
- [3] S. Ana-Maria, M. Bizoi, and F. G. Filip. Exploring multimedia web conferencing. *Informatica Economica*, 13(3):5–17, 2009.
- [4] L. S. Brakmo, S. W. O'Malley, and L. L. Peterson. TCP Vegas: new techniques for congestion detection and avoidance. *SIGCOMM Comput. Commun. Rev.*, 24(4):24–35, Oct. 1994.
- [5] A. Chan. *Designing Haptic Icons to Support an Urgency-Based Turn-Taking Protocol*. PhD thesis, 2004.
- [6] M. Cohen and J. Wang. System and method for very low frame rate video streaming for face-to-face video conferencing, 2009. Patent US 7583287.
- [7] G. Côté, B. Erol, F. Kossentini. H.263+: Video coding at low bit rates. *IEEE Transactions on Circuits and Systems for Video Technology*, 8:849–866, 1998.
- [8] S. E. Deering and D. R. Cheriton. Multicast routing in datagram internetworks and extended LANS. *ACM Trans. Comput. Syst.*, 8(2):85–110, 1990.

- [9] S. Foo, S. C. Hui, and S. W. Yip. Enhancing the quality of low bit-rate real-time internet communication services. *Internet Research*, 9(3):212–224, 1999.
- [10] J. Garcia-Luna-Aceves, P. Mantey, and S. Potireddy. Floor control alternatives for distributed videoconferencing over IP networks. In *Collaborative Computing: Networking, Applications and Worksharing, 2005 International Conference on*, page 10 pp., 0-0 2005.
- [11] G. Fox, G. Gunduz, and A. Uyar. Audio video conferencing. Technical report, 2001. Florida State University Information Technology Research Group.
- [12] P. Gill, L. Shi, A. Mahanti, Z. Li, and D. L. Eager. Scalable on-demand media streaming for heterogeneous clients. *ACM Trans. Multimedia Comput. Commun. Appl.*, 5(1):8:1–8:24, Oct. 2008.
- [13] R. Jain, P. Kim, and Z. Li. Experiential meeting system. In *Proceedings of the 2003 ACM SIGMM workshop on Experiential telepresence*, ETP '03, pages 1–12, New York, NY, USA, 2003. ACM.
- [14] A. M. Lieb and J. L. Bentom. Method and system for browser-based screen sharing, 05 2011. Patent US 2011/0126130 A1.
- [15] R. Malpani and L. A. Rowe. Floor control for large-scale MBONE seminars. In *Proceedings of the fifth ACM international conference on Multimedia*, MULTIMEDIA '97, pages 155–163, New York, NY, USA, 1997. ACM.
- [16] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven layered multicast. *SIGCOMM Comput. Commun. Rev.*, 26(4):117–130, 1996.
- [17] H. Research. Africa’s international bandwidth approaches 1 tbps mark, 2012. Africa Bandwidth Maps, <http://www.africabandwidthmaps.com/?p=2787> (Accessed: 16 May 2012).
- [18] J. Scholl, J. McCarthy, and R. Harr. A comparison of chat and audio in media rich environments. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, CSCW '06, pages 323–332, New York, NY, USA, 2006. ACM.
- [19] J. Scholl, P. Parnes, J. D. McCarthy, and A. Sasse. Designing a large-scale video chat application. In *Proceedings of the 13th annual ACM international conference on Multimedia*, MULTIMEDIA '05, pages 71–80, New York, NY, USA, 2005. ACM.
- [20] International Telecommunication Union. *African Telecommunication/ICT Indicators 2008: At a Crossroads*. International Telecommunication Union, 2008.
- [21] Z. Wen, Z. Liu, M. F. Cohen, J. Li, K. C. Zheng, and T. S. Huang. Low bit-rate video streaming for face-to-face teleconference. In *Proceedings of the 2004 IEEE International Conference on Multimedia and Expo*, pages 1631–1634, 2004.
- [22] C. C. Yang. User-interaction supported data-retrieving engine for distributed multimedia presentations. In *Communications, 2001. ICC 2001. IEEE International Conference on*, volume 10, pages 3244 –3250 vol.10, 2001.