

Interactive Multimedia Streams in Distributed Applications

Edouard Lamboray[†], Aaron Zollinger[†], Oliver G. Staadt[‡], Markus Gross[†]

[†]Computer Science Department, ETH Zurich, Switzerland

[‡]Computer Science Department, University of California, Davis

{lamboray, zollinger, grossm}@inf.ethz.ch staadt@cs.ucdavis.edu

Abstract

Distributed multimedia applications typically handle two different types of communication: request/reply interaction for control information as well as real-time streaming data. The CORBA *Audio/Video Streaming Service* provides a promising framework for the efficient development of such applications. In this paper, we discuss the CORBA-based design and implementation of *Campus TV*, a distributed television studio architecture. We analyze the performance of our test application with respect to different configurations. We especially investigate interaction delays, i.e., the latencies that occur between issuing a CORBA request and receiving the first video frame corresponding to the new mode. Our analysis shows that the interaction delay can be reasonably bounded for UDP and RTP. In order to provide results which are independent from coding schemes, we do not take into account any media specific compression issues. Hence, our results help to make essential design decisions while developing interactive multimedia applications in general, involving e.g. distributed synthetic image data, or augmented and virtual reality.

Keywords: Distributed Systems; Augmented, and virtual realities; Computer conferencing, teleconferencing, and videoconferencing; Performance of systems

1. Introduction

Distributed component architectures leverage the development of networked applications since they provide a powerful programming model for remote method invocation. However, distributed multimedia applications require additional features which allow for the efficient transmission of real-time streaming data.

During the last years, CORBA has been used in many research and industrial projects and has become one of the foremost standards in distributed computing [5]. Stable implementations are widely available, both as open-source and commercial software. However, CORBA was not suitable for time- and performance-critical applications, until the *Real-Time CORBA* specification was released [21].

This paper analyzes the TAO/ACE framework [22] and its CORBA *Audio/Video Streaming Service* implementation in the context of dynamically changing real-time transmissions [6]. We built an A/V streaming application as a test bed for the TAO/ACE toolkit and we benchmarked TAO's CORBA implementation using our test application. The major contribution of this work consists in the measurement of the latencies that occur between the invocation of a CORBA request and the reception of the first corresponding

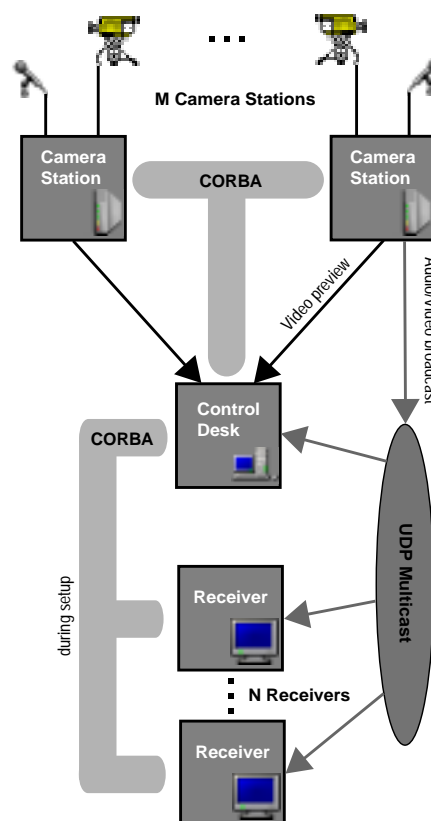


Figure 1: Campus TV, a distributed TV studio architecture.

video frame. Our quantitative results confirm that these latencies are bounded when the streaming data is transmitted by unreliable transport protocols, such as UDP and RTP. The absolute value of the latency depends on the average network load. Furthermore, we assess the scalability of our test application with respect to the number of simultaneous video streams. Unlike existing applications, the Campus TV test bed handles a large number of high-quality videostreams with little processing overhead due to coding. Hence, it allows for the efficient testing of the raw network and data transmission aspects of the CORBA A/V Streaming Service. Our study of CORBA's suitability for distributed and interactive multimedia streaming is part of the *blue-c* project, whose aim is the development of a novel platform for highly-immersive collaborative virtual environments [24, <http://blue-c.ethz.ch>]. It comprises, among other things, real-time acquisition of several video images of a human user and reconstruction of a 3-D graphical representa-

tion. This *geometry-enhanced video* will be transmitted via a high-speed network to several blue-c portals, where it will be seamlessly integrated into a distributed virtual environment. The blue-c system will eventually lead to a distributed virtual reality platform on which collaboration is possible using the most natural ways of interhuman communication and interaction, enhanced through the feeling of total immersion. Several research groups at ETH Zurich currently develop the various blue-c hard- and software components.

The blue-c environment will be a heterogeneous system on different platforms, including SGI Irix, Linux and Microsoft Windows. The increasing complexity of this system demands a middleware with an advanced programming model for distributed computing, supporting portability and real-time features. The chosen middleware needs to provide an efficient way of transmitting both control information, allowing for dynamic configuration of the system's setup and the exchange of less time-critical information, as well as latency-critical real-time streams with bandwidth requirements up to several megabits per second.

From a practical point-of-view, Campus TV, our test bed, implements a distributed television studio which is depicted in Figure 1. The following components are connected in a local area network:

- A variable number of *camera stations* that acquire audio and video signals and continuously transmit them to a control desk.
- A *control desk*, providing a graphical user interface which allows an operator to preview and configure all available streams.
- A theoretically unlimited number of *receivers* that listen to the multicast address of the TV program.

The remainder of this paper is organized as follows. After reviewing CORBA and its use in multimedia applications in Section 2 and related work in Section 3, we describe the overall architecture of our application in Section 4. Section 5 discusses some implementation details of the different components and Section 6 presents our performance analysis.

2. CORBA and Multimedia

A lot of effort has recently been put into the development of general-purpose communication middleware. A middleware is an intermediate software layer between low-level application programming interfaces and application code. The use of middleware toolkits leads generally to a better structured software architecture and allows for easier portability from one platform to another. However, not every middleware standard is appropriate for building distributed multimedia applications. The special requirements of this type of applications can be summarized as follows:

- Handling of continuous media streams.
- Support for multi-party communication.
- Quality of Service management.
- Real-time synchronization.

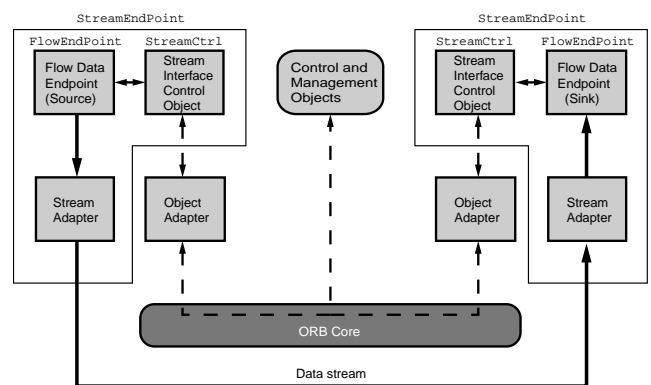


Figure 2: Data transmission using CORBA's Audio/Video Streaming Service.

Classical environments for distributed computing often do not fulfill the requirements of distributed multimedia applications [2]. Their initial focus lies in remote request/reply object interactions and their capabilities in continuous media modeling are limited [10]. Also the *Common Object Request Broker Architecture* (CORBA), supported by the Object Management Group (OMG) [5], does not naturally support the transmission of real-time streaming data. The CORBA layer introduces a large overhead through its marshalling operations and through the retransmission of lost packets [17, 22].

CORBA's *Audio/Video Streaming Service*, specified in [6], uses a promising concept for the transmission of time-critical streaming data. It makes a distinction between control information (i.e., connection setup, device configuration) and the real-time payload data. The control information can perfectly be handled by CORBA's client-server programming model, but the real-time data is directly streamed using classical transport protocols. On one hand, this design profits from advantages of the CORBA programming environment, on the other hand, it produces no overhead for the time-critical data. In order to guarantee interoperability between different A/V streaming applications, the service defines common interfaces for control and management of streams. Moreover, each stream can have an associated media controller, which implements application specific functionality. The A/V Streaming Service specification does not prescribe interface details of media controllers, such that they can be flexibly designed by the application developers.

A schematic description of the A/V Service is depicted in Figure 2. A stream contains one or many data flows and connects as many data sources with their corresponding data sinks. An endpoint, represented by the `StreamEndPoint` interface, comprises:

- A flow endpoint, which is either a data source or a data sink, represented by the `FlowEndPoint` interface.
- A stream interface control object, represented by the `StreamCtrl` interface, providing an IDL defined interface for controlling and managing the stream.
- A stream adapter, that receives or transmits data frames over a network.

3. Related Work

The Center for Distributed Object Computing at Washington University, St. Louis, provides with the TAO/ACE framework an advanced CORBA implementation that includes real-time features and additional services of CORBA 2.x, including the A/V Streaming Service [17]. ACE, the *Adaptive Communication Environment*, is an open-source object-oriented C++ framework that implements many core design patterns for concurrent communication software [20]. Furthermore, it provides an operating system abstraction layer and therefore improves the portability of applications that are built upon ACE. *The ACE ORB* (TAO) implements the standard CORBA reference model with most of the enhancements for real-time applications. Further information about TAO/ACE, as well as the sources of the current version can be downloaded at <http://www.cs.wustl.edu/~schmidt/TAO.html>. An MPEG video server is included as an A/V Streaming Service example application in the TAO/ACE distribution. It is based on the *Distributed A/V MPEG Player* from the Oregon Graduate Institute [4]. The Mbone video-conferencing application *vic* [15] was also ported to TAO's A/V Streaming Service [17]. Recent work includes a video distribution application, combining the A/V Streaming Service with the QuO Quality-of-Service framework [13]. None of these applications allows a simple adaptation for a performance analysis similar to the tests we present in Section 6.

In the field of collaborative virtual reality and distributed simulation, design suggestions for CORBA-based frameworks have been made [1, 7, 18] and on-going projects exist [9, 11], but no complete framework has been implemented until today.

ISO published through the *Reference Model for Open Distributed Processing* (RM-ODP) a meta-standard for distributed processing [2, 19]. CORBA, as well as other concrete standards, like PREMO (PRe-sentation Environments for Multimedia Objects) [8], implement the concepts of RM-ODP up to various extents. Further attempts to integrate CORBA with multimedia processing and transmission can be found in [3] and [26]. At the time of those research projects, the CORBA A/V Streaming Service was not yet fully specified.

The *Real-time Transport Protocol* (RTP) is an appropriate transport protocol for multimedia data delivery [23], it will be taken into account in Section 6 of this paper.

McCanne et al. propose a common infrastructure for multimedia networking middleware, the *MASH* project [14]. Their research effort generated a large number of applications, among which a video conferencing application that can be remotely controlled [12], and a control system for live webcasts [25]. The webcast application, which has a more developed functionality than Campus TV, currently still requires special hardware for video switching and is implemented using ad-hoc remote method invocation.

4. System Description

This section gives an overview of the Campus TV test application we built for the evaluation of the CORBA A/V Streaming Service. The supported data classes are described, together with some important design decisions.

4.1 Overview

In our application, several camera stations grab live video images, potentially at a low resolution and at a low frame-rate. This information is continuously streamed to the control desk. A user can select the camera that should send its image as the current image to the receivers. In the following, we will refer to the current camera station as the broadcast camera station. We use this term from the area of television transmission, but it should not be mixed up with its meaning in the context of computer networks.

The choice of the broadcast station is communicated to the camera station using CORBA requests, whereas the video and audio data are distributed using configurable IP transport protocols. The camera station which gets the broadcast token from the control desk starts sending its input in high resolution and at highest achievable frame-rate to a multicast address. All control information, like image dimensions and frame-rate, can be dynamically configured during a session. Changes in the setup are also propagated using the ORB. Figure 1 shows a system overview of our test bed.

The receivers first ask the control desk for information about the currently distributed program. Then they add themselves to the multicast session using the A/V Streaming Service. The control desk knows about all the distributed camera stations delivering possible input. But it does not keep track of the receivers and does not exchange any information with them, except for the setup data.

Since the high quality stream of the broadcast camera station is distributed using UDP multicasting, the system scales very well to a high number of receivers. Furthermore, there is no a priori limit on M , the number of camera stations, even though every new camera station introduces additional network traffic and increases the workload at the control desk. In fact, we can trade-off M against the quality of the preview images, while dynamically configuring the image resolution and the frame-rate of the preview video streams. Finally, since there is no centralized sender of the broadcast stream, the A/V data is transmitted with minimal overhead and latency.

4.2 Supported Data

4.2.1 Control Information

The control information, which is different from the audio and video samples, includes commands for managing the streams, i.e., start/stop, as well as configuration parameters for the data acquisition. It can also be used for monitoring the system. The control desk configures for example:

- The dimensions of the video images and the acquisition frame-rate at the camera station.
- The sample frequency of the audio acquisition.
- The permission to send to the multicast address.

The action of switching from camera station A to camera station B as the broadcast station can thus be summarized with the following actions:

- Send current broadcast settings to B .

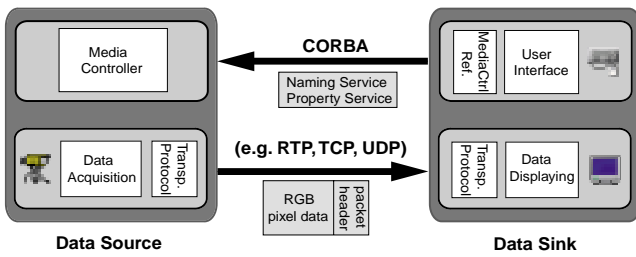


Figure 3: Streaming video images from a data source to a data sink using a point-to-point connection.

- Instruct *A* to stop sending to the multicast address.
- Instruct *B* to start sending to the multicast address.

The A/V Streaming specification suggests that a media control object is associated to each `FlowEndPoint`. The media controller implements device specific operations, which are not general enough to be handled by `StreamCtrl` objects. Since the media controller is completely specified by the application developer, any type of information can be communicated and the requirements of different applications can be met. Figure 3 depicts a typical stream setup in a point-to-point connection and shows the different paths for control and real-time streaming data

4.2.2 Audio/Video Transmission

The video frames are grabbed as RGB images with 24 bits per pixel color information. As camera stations, we use SGI O2 workstations with video option. In this paper, we focus on dynamically configurable A/V streams and we assess the performance of TAO's A/V Streaming Service implementation. Hence, we did not yet integrate any compression techniques in Campus TV. Of course, this would be necessary for deploying Campus TV in a large scale environment.

TAO's A/V Streaming Service implementation provides a pluggable protocol framework, in which the most common IP-based transport protocols, such as TCP, UDP (unicast and multicast), and RTP, as well as ATM, are already included. A factory object encapsulates the concrete protocol objects and simplifies the addition of future data transfer protocols. Furthermore, it allows us to rapidly integrate different protocols into our application. Since the connection setup is completely handled by the A/V Streaming Service specification, the application programmer only needs to adapt his payload data to the various protocols. In our case, we use for example sequence numbers for identifying the fragmented image packets. TCP and RTP have already sequence numbers in their protocol headers, UDP has not. Hence, we implemented an additional software layer, that handles the protocol-dependent packetization of our payload data.

The CORBA A/V specification distinguishes between flows, carrying data in one direction, and streams, comprising a set of flows. There are two ways of specifying a connection, either via the *light* or via the *full profile*. In the light profile, the application is granted access only to stream endpoints, as well as to virtual and multimedia devices. According to [6], virtual and multimedia devices abstract physical or logical devices consuming or producing the

samples of the multimedia stream. If access to the flow endpoints or the flow connection is required, the full profile must be used. The TAO/ACE framework implements both configurations. We used the full profile, since the separate access to the video and audio flows makes our design more flexible for further investigations.

A critical issue for the efficient distribution of data from one to many users is the availability of point-to-multipoint communication. Of course, a point-to-multipoint distribution can be realized using many point-to-point connections, but this strategy does not make efficient use of the available resources. The CORBA A/V Service specification includes point-to-multipoint binding, but does not define multipoint-to-multipoint binding, where several sources communicate with several sinks. The TAO framework, however, supports multipoint-to-multipoint communication, which finally is an important feature of our test application.

Additionally, in classical UDP multicasting, the data source is not aware of whoever is listening to its data transmission. In the CORBA A/V Service, however, each data sink needs to be explicitly added to the data sources' stream controllers. In our application, the control desk defines naturally the stream configuration and candidate camera stations as well as receivers must judge during connection establishment if they are able to handle the available streams.

5. Implementation

In the following, some implementation details of the different components of our Campus TV application are presented.

5.1 Control Desk

The control desk is used for configuring a broadcast session and is the interactive part in our application. A screenshot representing the control desk's graphical user interface is shown in Figure 4. The left part of the user interface is dedicated to the thumbnail preview images, the right part deals with the broadcast stream.

The control desk first registers itself at the CORBA *Naming Service*, where the clients can later retrieve a reference to the control desk object.

Upon the clients' requests, the control desk distributes setup information to both camera stations and receivers. From a practical point of view, a separate service could offer this functionality to the receivers.

Furthermore, the control desk initializes the multicast transmission, since it is the first listener to the attributed multicast address. It also creates the first data source, which is actually a dummy source, but which is required by TAO's implementation of the A/V Streaming specification for multipoint connection setup.

The application can be implemented according to two different strategies: *reactor-based* or *process-based*. For the Campus TV control desk, the process-based strategy would lead to a new process for every connecting camera station. Currently, we implemented the reactor-based strategy, i.e. all input from the camera



Figure 4: The graphical user interface of the control desk.

stations is handled by the same process, and events indicate that new data has arrived. Since the update rates of the thumbnail preview images are not very high, the reactor-based control desk's performance is sufficient. An analysis of the control desk's scalability with respect to the number of camera stations can be found in Section 6.3.

Apart from the *Naming Service*, the control desk also uses the CORBA *Property Service* for retrieving a reference of the media controller of a given camera station.

All graphical user interfaces have been implemented using Trolltech's Qt cross-platform GUI framework (<http://www.trolltech.com>). Since our goal is to build a portable communication layer, we also need a GUI toolkit available for all platforms of interest. Another significant advantage of the Qt library consists in its OpenGL support (<http://www.opengl.org>). The Qt classes `QGLWidget`, `QGLContext` and `QGLFormat` allow respectively to create an OpenGL rendering context and display format and to render an OpenGL scene integrated into a Qt GUI. Although OpenGL rendering is not necessary for Campus TV, we will need it for our future research, which deals with geometry-enhanced video and its integration into virtual scenes.

Rendering video images in OpenGL is possible with two different strategies: The OpenGL `glDrawPixels` function allows us to write a block of pixels directly to the frame buffer. The second possibility consists in a texture map of the video image on a polygon. On modern graphics processing units, texture mapping is generally faster and less CPU-consuming than the pixel drawing strategy.

Note that the Simple DirectMedia Layer (<http://www.libdsl.org>) provides a similar functionality and could be used for an alternative implementation.

5.2 Camera Station

The video images are grabbed at regular time intervals. The frame-rate of the video acquisition can be controlled by changing the time-out value. The audio acquisition is implemented in an analogous way to the video acquisition.

When a camera station is started, it asks for a reference to the control desk at the *Naming Service* and then connects itself with the control desk. It starts transmitting video data according to the current configuration. The camera station joins the multipoint-to-multipoint connection already during setup. Hence, it is ready to start sending both audio and video data to the multicast address as soon as it becomes the broadcasting station.

5.3 Receiver

The receiver first retrieves information about the current session from the control desk. Then it connects to the multicast address and presents the received data to the user.

Receiver and control desk only interact during the initial connection. The control desk cannot directly configure the receivers and is unaware of the number or type of current receivers. Currently, our implementation does not include an explicit synchronization for audio and video play out.

6. Performance Analysis

This section summarizes the performance analysis we conducted with the Campus TV test bed. After describing the experimental setup, we discuss our results concerning interaction delays, scalability and audio/video jitter.

6.1 Experimental Setup

For the following experiments, we used up to ten SGI O2 workstations with video option as possible camera stations and an SGI Octane as the control desk's host. The workstations were all running SGI Irix 6.5 and connected by a 100 Mbps Fast Ethernet switch. We ran the tests when the average load on our local area network was low.

Note that all experiments in Section 6.2 and 6.3 were done without audio transmission. Moreover, some of the following configurations are only useful for testing purposes. Running Campus TV in "application mode" requires low resolution preview images at 5-10 Hz and high resolution broadcast images, using RTP for data transmission.

6.2 Interaction Delay

We call *interaction delay* the duration between issuing a CORBA request at the control desk and the moment of reception of the first video frame according to the new mode. In practice, our camera station's media controller implements a method for changing the color of a given pixel after frame acquisition. At the control desk, we check the test pixel's RGB values in the received video frames, and hence we can measure the delay between issuing the `set-`

Color command at the control desk and the reception of the corresponding frame. The IDL interface of the function is as follows:

```

struct Color {
    short r;
    short g;
    short b;
};
void setColor(in Color col);

```

We measured the interaction delay by generating automatically a series of `setColor` commands at the control desk. We fixed a range for the duration in between two `setColor` invocations and generated uniformly distributed time-outs in that range. The test results presented in this paper were obtained with time-out values between 250 and 750 ms. They are based on 1024 samples respectively, i.e. every experiment ran for about 8 minutes and 30 seconds. We repeated the same experiments for various time-out ranges, but we did not observe significant differences in the resulting interaction delays.

Three image resolutions for both preview and broadcast images can be configured:

- low: 64 x 64 pixels.
- medium: 128 x 128 pixels.
- high: 256 x 256 pixels.

Table 1: Interaction delay measurement statistics for different protocols and preview image resolutions, using a broadcast resolution of 256x256 pixels.

Preview/ Broadcast	Mean [ms]	StDev [ms]	Min [ms]	Max [ms]	Loss [%]
UDP-lo/UDP	135	20	104	179	0
RTP-lo/RTP	136	20	104	182	0
TCP-lo/RTP	134	19	104	183	0
UDP-md/RTP	148	21	113	199	0
RTP-md/RTP	146	21	101	212	0
TCP-md/RTP	184	45	111	577	0
UDP-hi/UDP	258	32	203	316	32
RTP-hi/RTP	257	32	204	320	36
TCP-hi/RTP	425	146	250	1074	3

Table 1 shows the mean value and the standard deviation for different preview image resolutions and for UDP, RTP and TCP point-to-point connections respectively. The resolution of the broadcast image was 256 x 256 pixels in each case. The interaction delay was measured on the broadcast connection. With the resolution of the preview images, we can influence the network traffic. In the low and medium resolution configuration, we used ten camera stations, sending preview images at 10 Hz. This leads to a network traffic of approximately 34 Mbps or 62 Mbps respectively. The broadcast video is always streamed at the maximum possible frame-rate. For the low and medium sized preview images, a broadcast frame-rate

of 14-17 Hz is achieved. In the high resolution case, we only used five camera stations at 10 Hz, which leads to an approximate network traffic of 100 Mbps, i.e. a certain amount of packets is certainly lost. The broadcast frame-rate drops down to 5-7 Hz. The loss indicated in Table 1 refers to the broadcast stream.

Note that in the UDP test case, we also used UDP multicasting for the multipoint-to-multipoint broadcast communication, in the RTP and TCP test cases, we used RTP multicasting for the broadcast image. A sample is regarded as lost if a corresponding broadcast image frame is never received.

In Figure 5 and 6, the vertical axis indicates, for a time interval Δt , the percentage of samples, where the interaction delay was larger than Δt . Hence, Figure 5 and 6 represent an approximation of a probability distribution $P(t)$, where $P(t = \Delta t)$ is the probability that the interaction delay is larger than Δt .

In Figure 5, we observe that in the configuration with low sized preview images, 99% of the interaction delays were shorter than 175 ms. The choice of the transport protocol did not significantly influence the probability distributions. In the case of medium sized preview images, the probability distributions for UDP and RTP are only slightly different from the previous case. But for TCP, we already observe a significant increase of the interaction delay.

In the case of preview images at a high resolution, see Figure 6, the average interaction delay increases by more than 100 ms for UDP and RTP. Also the standard deviation of the UDP and RTP data sets increases by approximately 50%. We realize that in this configuration, TCP is not competitive anymore, the interaction delay values of the TCP samples are much higher than for the two unreliable protocols. Note that in this test case, where the network traffic is very high, a certain amount of multicast frames from the broadcasting camera station get lost, and hence not every sample generates a valid interaction delay. The amount of lost frames is indicated in Table 1.

Furthermore, part of the observed interaction delay is due to the low frame-rate at the high resolution configuration. If both resolutions for preview and broadcast images are 256 x 256 pixels, the frame-rate of the broadcast stream drops down to 5-7 Hz. The expected delay simply due to this low frame-rate is at least 70 ms, and hence cannot be totally neglected.

The probability distribution curves of the UDP and RTP experiments lead to the conclusion that the interaction delay can be bounded. For each setting, the interaction delays occurred in a given time frame, whose mean and width were depending on the average network load. The interaction delay is much less predictable for TCP because of its inherent flow control algorithms.

Finally, we observe that the probability distributions for RTP and UDP are in general very similar. On one hand, this can be expected since RTP uses UDP as its transport protocol. On the other hand, it confirms that TAO's RTP implementation does not introduce a significant amount of overhead.




Figure 5: Semi-logarithmic plot of the interaction delay distribution for configurations with low and medium resolution preview images from ten camera stations, the broadcast image has high resolution.

6.3 Scalability

Figure 7 shows the average CPU load at a camera station for various configurations. The camera station which is also the broadcast station needs to transmit both low quality preview images and high quality broadcast images. We observe that the CPU load increases significantly with the size of the video frames. The acquisition frequency is only relevant for the preview images, the broadcast images are always streamed at the maximum possible frame-rate. Note that in the configuration with high-resolution broadcast images and low-resolution preview images, acquisition frequencies higher than 20 Hz are not achievable. The two first data sets in Figure 7 were obtained with a preview stream only.

In Figure 8, it appears that the CPU load at the control desk is much more influenced by the number of camera stations than by their frame-rate. If we increase the number of camera stations from one to eight, the CPU load only grows from 30% to 60%. For high acquisition frequencies, the network load does not allow a very regular frame-rate for the broadcast stream. This explains why the variation in some test cases is not monotonic. In fact, the update rate of the broadcast image at the control desk lies between 14 and 17 Hz.

The scalability experiments were all done with RTP and RTP multicast for the preview and the broadcast images respectively.

6.4 Audio/Video Jitter

Additionally, we investigated the phase jitter between audio and video data. Looking at the video and audio packet transmission, we observed no delays between audio and video packets at the broadcast camera station and at the receivers. However, as can be seen in Figure 7, the high quality video transmission uses already 80% of the camera station's CPU. In this case, a constant delay of about 0.5 seconds between audio and video play out can be perceived. It can be explained by our non-optimized audio acquisition schedul-

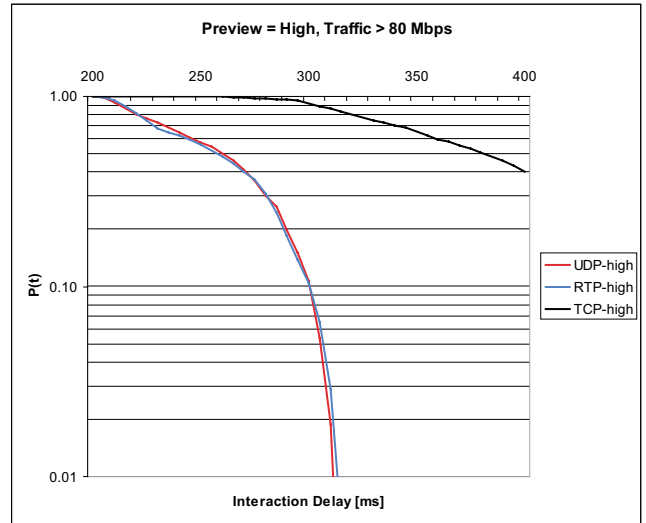


Figure 6: Semi-logarithmic plot of the interaction delay distribution for configurations with high resolution preview images from five camera stations, the broadcast image has high resolution.

ing and play out algorithms. An optimization of the audio transmission was beyond the focus of this paper, since, in the blue-c system, speech transmission will be handled by a different subsystem than the geometry-enhanced video transmission.

7. Conclusions and outlook

The performance analysis presented in this paper shows that the CORBA A/V Streaming Service introduces no critical overhead for multimedia data transmission. The concept of media controllers allows the flexible integration of application specific interactions within the streaming framework. Our measurements show that the interaction delay is bounded when using appropriate transport protocols, like UDP and RTP. The absolute value of the interaction delay depends on the average network load. Furthermore, the Campus TV control desk scaled well with an increasing number of camera stations. In the future, the pluggable protocol framework of the TAO service implementation will allow us to run similar tests on different network configurations. Finally, the results of this study will influence important design decisions in our future distributed virtual reality platform, i.e. we will use our findings for deciding what type of requests can be handled by the ORB, and what type of information, apart from the multimedia data, needs to be streamed in real-time.

However, concerning the TAO/ACE framework, we encountered some problems during the implementation of Campus TV. The connection setup is not completely transparent and we had to take special care with RTP transmission. This is especially due to the fact that TAO's RTP implementation is still a bit unwieldy. But improved TAO versions appear on a regular basis and its open source character allows for bypassing local problems.

Note that in our target application, the blue-c system, the data for controlling and monitoring will be much more complex and diverse than in the Campus TV test application. At that point, we

will take fully advantage of the standardized stream interfaces of the A/V Streaming Service as well as of the flexibility provided through the concept of media controllers. Possible optimizations with respect to acquisition and compression will also be taken into account for the geometry-enhanced video streams in the upcoming blue-c prototype. A first prototype is currently under development and uses extensively the functionality provided by CORBA and its A/V Streaming Service. In the future, we also want to further investigate the possibilities offered by the MPEG-4 standard, which supports the integration of multimedia streams and traditional A/V formats with synthetic images as well [16].




Figure 7: CPU load for different RTP Preview / Broadcast configurations at the camera station.

Finally, we envision to extend the Campus TV application such that special effects, e.g., image warps and fade-outs, can be performed on the video images. These special effects can be easily and efficiently implemented at the receiver, using OpenGL commands on the polygons on which the live video texture is mapped.

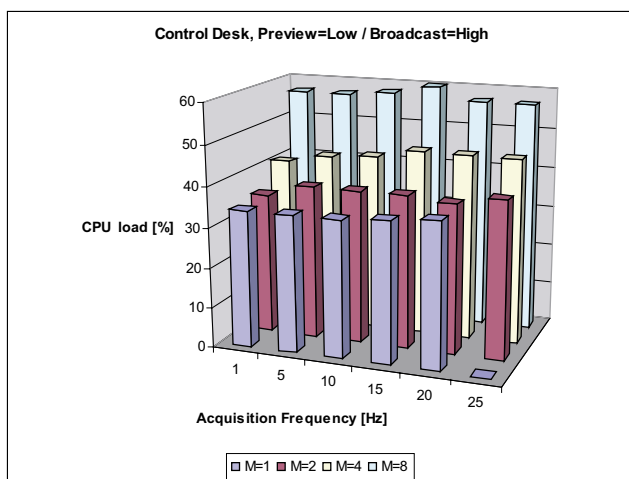


Figure 8: CPU load at the control desk with respect to the number of preview images and preview acquisition frequency with low preview size and high resolution broadcast using RTP.

8. Acknowledgements

We would like to thank the numerous researchers who developed the TAO/ACE framework as an open-source project. Many thanks to the Perceptual Computing and Computer Vision Group at ETH Zurich for letting us use their hardware equipment. The blue-c project is funded by ETH Zurich as a “Polyprojekt”.

9. References

- [1] T. A. Au. “Performance Issues of HLA Run Time Infrastructure based on CORBA.” In *Proceedings of The Simulation Technology and Training (SimTecT2000) Conference*, Sidney, Australia, 2000.
- [2] G. Blair and J.-B. Stefani. *Open Distributed Processing and Multimedia*. Addison Wesley Longman, 1997.
- [3] C. Blum and R. Molva. “A CORBA-based platform for distributed multimedia applications.” In *Proceedings of Multimedia Computing and Networking MMCN’97*, San Jose, CA, February 1997.
- [4] S. Cen, C. Pu, R. Staehli, and J. Walpole. “A Distributed Real-Time MPEG Video Audio Player.” In *Proceedings of the Fifth International Workshop on Network and Operating System Support of Digital Audio and Video (NOSSDAV ’95)*, pages 151–162, Durham, New Hampshire, April 1995.
- [5] “The Common Object Request Broker: Architecture and Specification.” Object Management Group, Revision 2.5, September 2001.
- [6] “Audio/Video Stream Specification.” Object Management Group, January 2000.
- [7] F. V. Deriggi Jr., M. M. Kubo, A. C. Sementille, J. R. F. Brega, S. G. dos Santos, and C. Kirner. “CORBA Platform as Support for Distributed Virtual Environments.” In *Proceedings of the IEEE Virtual Reality 1999 conference*, pages 8–13, March 1999.
- [8] D. J. Duke and I. Herman. “A standard for multimedia middleware.” In *Proceedings of the 6th ACM international conference on Multimedia*, pages 381–390, Bristol, United Kingdom, September 1998.
- [9] J. Gallop, C. Cooper, I. Johnson, D. Duce, G. Blair, G. Coulson, and T. Fitzpatrick. “Structuring for Extensibility - Adapting the Past to Fit the Future.” In *Proceedings of CBG2000, the CSCW2000 workshop on Component-Based Groupware*, December 2000.
- [10] D. G. Waddington and G. Coulson. “A Distributed Multimedia Component Architecture.” In *Proceedings of the 1st International Workshop on Enterprise Distributed Object Computing, Gold Coast, Australia*, pages 334–347, October 1997.
- [11] “Standard for Modeling and Simulation (M&S) High Level Architecture (HLA) - Framework and Rules.” IEEE Standard 1516, September 2000.
- [12] T. Hodes, M. Newman, S. McCanne, J. Landay, and R. Katz. “Shared Remote Control of a Video Conferencing Application.” *SPIE Multimedia Computing and Networking*, pages 17–28, January 1999.
- [13] D. A. Karr, C. Rodrigues, Y. Krishnamurthy, I. Pyarali, and D. C. Schmidt. “Application of the QuO Quality-of-Service Framework to a Distributed Video Application.” In *Proceed-*

ings of the 3rd International Symposium on Distributed Objects and Applications, September 2001.

- [14] S. McCanne, E. Brewer, R. Katz, L. Rowe, E. Amir, Y. Chawathe, A. Coopersmith, K. Mayer-Patel, S. Raman, A. Schuett, D. Simpson, A. Swan, T.-L. Tung, D. Wu, and B. Smith. "Toward a Common Infrastructure for Multimedia-Networking Middleware." In *Proceedings of 7th Intl. Workshop on Network and Operating Systems for Digital Audio and Video (NOSSDAV'97)*, pages 41–51, May 1997.
- [15] S. McCanne and V. Jacobson. "vic: A Flexible Framework for Packet Video." In *Proceedings of the ACM Multimedia'95*, pages 511–522, San Francisco, California, November 1995.
- [16] "Overview of the MPEG-4 Standard." ISO/IEC JTC1/SC29/WG11 N3444, May/June 2000.
- [17] S. Mungee, N. Surendran, Y. Krishnamurthy, and D. C. Schmidt. *The Design and Performance of a CORBA Audio/Video Streaming Service*, chapter in *Design and Management of Multimedia Information Systems: Opportunities and Challenges*. Idea Group Publishing, Hershey, 2000.
- [18] C. O’Ryan, D. L. Levine, D. C. Schmidt, and J. R. Noseworthy. "Applying a Scalable CORBA Events Service to Large-scale Distributed Interactive Simulations." In *Proceedings of the 5th Workshop on Object-oriented Real-time Dependable Systems*, Monterey, CA, November 1999.
- [19] "Open Distributed Processing - Reference Model." Standard ISO/IEC 10746, 1995.
- [20] D. Schmidt, M. Stal, H. Rohnert, and F. Buschmann. *Pattern-Oriented Software Architecture: Patterns for Concurrent and Networked Objects*, volume 2. John Wiley & Sons, 2000.
- [21] D. C. Schmidt and F. Kuhns. "An Overview of the Real-time CORBA Specification." *IEEE Computer, special issue on Object-Oriented Real-time Distributed Computing*, 33(6):56–63, June 2000.
- [22] D. C. Schmidt, D. L. Levine, and C. Cleeland. *Architectures and Patterns for High-performance, Real-time CORBA Object Request Brokers*, volume 48 of *Advances in Computers*, pages 1–118. Academic Press, 1999.
- [23] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. "RTP: A Transport Protocol for Real-Time Applications." RFC 1889, January 1996.
- [24] O. G. Staadt, M. H. Gross, A. Kunz, and M. Meier. "The Blue-C: Integrating Real Humans into a Networked Immersive Environment." In *Proceedings of ACM Collaborative Virtual Environments 2000*, pages 201–202, San Francisco, September 2000. ACM Press.
- [25] T.-P. Yu, D. Wu, K. Meyer-Patel, and L. Rowe. "dc: A Live Webcast Control System." In *Proc. IS & T/SPIE Symposium on Electronic Imaging: Science & Technology*, Multimedia Computing and Networking, San Jose, CA, January 2001.
- [26] T.-H. Yun, J.-Y. Kong, and J. W. Hong. "A CORBA-based Distributed Multimedia System." In *Proceedings of the Fourth Workshop on Distributed Multimedia Systems*, pages 1–8, Vancouver, Canada, July 1997.