Managing Video Data in a Mobile Environment

Rafael Alonso
Yuh-Lin Chang
Liviu Iftode
V. S. Mani
Matsushita Information Technology Laboratory
2 Research Way
Princeton, NJ 08540-6628
{alonso, yuhlin, liv, mani}@research.panasonic.com

Abstract

Two key technological trends of the last few years have been the emergence of handheld computational elements and the implementation of practical wireless communication networks. These two changes have made mobile computer systems feasible. While there has been much research interest devoted to mobile computer issues, such systems have not yet been commercially successful. This has been ascribed to the lack of a *killer mobile app*. We believe that the support of video on mobile systems will indeed make possible many new interesting applications.

However, providing mobile video is a non-trivial task, and much work needs to be done before practical systems are widely available. In this short note we address the issue of mobile multimedia from a practitioner's perspective. We note what software and hardware are currently available in the market in support of mobile multimedia, and point out some of their deficiencies. We also discuss some of the communication and data management research issues that need to be tackled in order to address said deficiencies. Exploring these research issues is the focus of our project.

1 Introduction

Recent improvements in a number of technologies have made it possible to consider the design and implementation of practical mobile computer systems. Such systems have generated much research interest (witness the many recent workshops on the topic and the large numbers of papers published in the area). This interest has spanned a number of research areas, from databases (e.g., the MOBIDATA workshop in 1994) to operating systems and communications (e.g., the MobiCom'95 conference). But all this academic interest has not resulted in commercial success. For example, the earliest widely marketed personal digital assistant (PDA), the Apple Newton, saw its sales slow down markedly once early adopters' sales dwindled. Many believe that this lack of market success is due in great part to the lack of a *killer mobile app* [1] that will entice users to invest in the technology. We believe that the support of video on mobile systems will indeed lead to a number of potentially interesting applications.

In this short note we do not have the space to motivate further the importance of mobile video, but instead we will focus on describing our current approach to the problem. We start by providing some background information on a couple of related areas. In particular, we present a short overview of compression protocols, and discuss various aspects of wireless communications. Compression is of interest in this area because of the very large amount of data involved. Thus, in the first section below we provide a brief survey of compression techniques; of particular importance for mobile work is the MPEG4 effort which should address low bit rate environments commonly found on mobile networks. In Section 3 we present a sampling of communication related issues. Specifically, we focus on some of the communication requirements of mobile data.

Our research approach is described in Section 4. We start by giving an overview of our experimental

[1] By *killer app* we mean the application that will motivate users to invest in a technology, much in the manner that spreadsheet applications fostered sales of personal computers

setup. This setup is composed of (mostly) commercially available items. Although other researchers in mobile video have developed highly customized approaches (notably the INFOPAD project at U.C. Berkeley), we have attempted to put together a mobile video system using only off-the-shelf components (both hardware and software). Our approach is to carefully examine such a system to determine where it fails to provide adequate multimedia service. We will then focus our efforts in developing customized solutions to address the failings we detect. We conclude by presenting some of the initial observations we have made on our multimedia setup.

2 Video Compression Standards

In this section we would like to give a brief survey of video compression standards that are of importance to mobile multimedia data management. (For further information, we suggest the following references: [11, 12, 13].) The most important standards effort is MPEG (Moving Pictures Expert Group), which is the name of the ISO committee working on digital video and audio compression. So far two sets of standards, MPEG-1 and MPEG-2, have been finalized. Both standards contain three parts: video encoding, audio encoding, and "system", which includes information about the synchronization of the audio and video streams. In a nutshell, the MPEG encoding algorithm uses two basic methods to compress video data: block-based motion compensation to explore temporal redundancy, and DCT (Discrete Cosine Transform) to explore spatial redundancy. MPEG-1 can deliver VHS-quality digital videos with data rates around 1.5 Mbps (mega bits per second) and up to 4 Mbps. This corresponds to the data retrieval speed from CD ROM and DAT (digital audio tape). In a 1.5 Mbps stream, the video stream takes about 1.15 Mbps, and the remaining bandwidth is used by the audio and system data streams. MPEG-2 uses similar compression algorithms as MPEG-1. However, MPEG-2 is designed to offer higher quality at a bandwidth of between 4 and 20 and can go up to 80 Mbps. The target applications are HDTV and broadcast TV. MPEG-1 can be considered as a subset of the MPEG-2 standards.

Of particular interests to the mobile computing community is the upcoming MPEG-4 standard, which targets the very low bit rate (below MPEG-1 bit rate) applications such as videophone, multimedia database retrieval, and video games. The major differences between MPEG-4 and all existing compression standards are its emphasis on content-based manipulation of audio/video data and its open toolkit-based approach. The focus of MPEG-4 will be on content-based interactivity, high compression, and universal accessibility. To provide these new capabilities, MPEG-4 will develop a flexible syntactic descriptive language, and a number of audio-visual coding tools. This approach will provide the ability for decoders to use a rich set of algorithms according to the content or applications. The standard is still under development, and it is scheduled to be finalized in 1998. New compression techniques, such as wavelet, model-based algorithms, etc., are expected to be employed.

One last standard of interest is H.261, which a standard recommended by the International Telecommunications Union (ITU). It describes the video coding and decoding methods at video rates of p x 64 kbit/s, where p is in the range 1 to 30. The standard has three parts, the video source coder, the video multiplex coder and the transmission coder. H.261 is intended for videophone applications and for video conferencing. Videophone is less demanding of image quality, and can be achieved for p=1 or 2. For a videoconferencing application (where there are more than one person in the field of view) higher picture quality is required and p must be at least 6. The actual encoding algorithm is similar to (but not compatible with) that of MPEG. Another difference is that H.261 needs substantially less CPU power for real-time encoding than MPEG. Currently H.261 is supported by most major PC-based videoconferencing manufacturers.

There are many other compression standards, but their main significance as far as data manipulation is concerned is the high computational overhead they impose on any system. Given the relatively low computing power of mobile machines, hardware support for compression algorithms is a necessity for mobile multimedia. At the current time no handheld device of which we are aware has hardware to aid in the compression and decompression process. Determining the proper hardware and software support for compression algorithms is one of the aspects of our work.

3 Wireless communication

As the transmission capacity of wireless communication moves closer to the video transmission rate, multimedia services are expected to use wireless connections to provide limited-quality video and audio streams. However the requirements for high quality multimedia connections are still difficult to fulfill in a wireless medium which suffers from severe shortcomings in several areas (e.g., low bandwidth and high bit error rate). Moreover, these limitations aggravate the impact that the real-time transmission constraints of multimedia communications imposes on network protocols.

Although wireless products have already appeared on the market, adequate protocols are still a debated and research-focused area. Wireless communications is closely associated with mobility and mobility in turn introduces new protocol demands such as routing in a dynamic network and temporary disconnection handling which were not supported (or necessary) in existing internetwork protocols. Several mobile-host network-layer protocols have been proposed in recent years ([2], [3], [4]). To implement wireless communication as an extension of the existing Internet, compatibility with TCP/IP must be ensured. The main idea is to provide two identification paths for a mobile host at the network layer: one of them, called home, with a fixed location, the other bound to the current location of the mobile. In addition, a forwarding mechanism with various degrees of optimization is provided in order to allow communication to follow the mobile's motion in the network.

Recent studies have shown that mobility affects the transport layer as well ([5], [6]). In the absence of additional information the transport protocol cannot distinguish between packet losses due to motion from those due to congestion. In order to provide a specific reaction the transport layer must also be aware of mobility.

The performance improvement opportunities for wireless communication do not stop here. For instance, studies in various application areas, particularly in databases ([7]), showed that substantial benefits both in functionality and performance, can be obtained when mobility is considered in designing the application layer protocols.

Supporting multimedia applications over wireless communications entails additional technical difficulties. For example, multimedia communication is particularly demanding in terms of delay constraints since end-to-end delay has to meet playback deadline. Additionally, video compression introduces variable transmission delays (called packet delay jitter) which must be compensated. In general, bandwidth reservation, stream synchronization, buffering and prefetching are the main issues in designing specific protocols for multimedia communication ([8]).

Furthermore, a multimedia stream has some additional characteristics that differentiates it from other data types. Consider that usually a multimedia stream contains data of various importance to the quality of the presentation ([9]). For instance, audio performance is critical to the perceived quality of a playback as well as to the overall understanding of the information which is sent. In many cases the video information may be treated as secondary, serving only to enrich the audio to which it is subordinated as in a speech recording. Therefore, misses in the video stream can be tolerated more easily than misses in the audio stream. It is then reasonable that, under the pressure of a bandwidth shortage, one would handle an audio stream with higher priority than a video stream. However, if one uses a general-purpose protocol that does not take information content into account, the protocol is unable to make intelligent decisions to provide the maximum possible quality.

Multimedia dedicated protocols, used in wired local-area networks between the video server and clients, can be extended over wireless connections as long as there are sufficient resources. Unfortunately, a wireless mobile host has limited CPU power, buffering capabilities, and communication bandwidth. Therefore, multimedia communication requires specific techniques to deliver acceptable quality under wireless links to mobile hosts.

Acceptable multimedia playback quality imposes strict delay constraints. Under existing wireless network interfaces, these time constraints rule out the usage of a reliable transport protocol such as TCP which guarantees data delivery but under no real-time pressure. Alternative video-ignorant non-reliable protocols may cause arbitrary loss of data which may non-uniformly damage the quality of the multimedia playback. Audio over video priority can be employed to trade quality for continuous playback when time is the critical. Since audio data needs only a relatively small fraction of the bandwidth required for video data we believe that the two streams deserve distinct protocols in the limited-bandwidth environment of mobile systems.

One of the key aspects of our project is the investigation of the effects of wireless communication on multimedia quality using existing dedicated and internet protocols. We are particularly interested in exploring the benefits of using knowledge about multimedia content to improve the overall performance of communication over wireless links.

4 Experimental Approach

The main goal of our project is to explore research issues in mobile multimedia. In an area such as this one where there is a great deal of commercial interest it is very easy for a research project to develop solutions for non-existing problems or for problems whose answers are already available in the marketplace. To avoid these pitfalls, our research plan has been to compose the best possible mobile multimedia system crafted from off-the-shelf hardware and software components. Such a system will not be able to support well or fully the type of multimedia services we envision, but by determining the shortfalls of the system we will guide our research towards their solutions.

In this section we will first describe the environment we have put together, and then point out some of the problems we have measured or observed. We also briefly sketch our plans for dealing with those difficulties.

4.1 Experimental Environment

One of the difficulties with mobile multimedia is that current PDAs do not have sufficient processing capacity to fully process video information. Thus, since we are targeting the mobile environment of a few years hence we substitute laptops and desktop PCs to emulate PDAs. We employ mostly 486 PCs, running Windows for Workgroups (unless otherwise noted).

The machines in our system are connected by a two-piece network. A few PCs are connected together via a standard Ethernet connection. A second set of PCs and laptops with WaveLAN adapters are connected together in a wireless Ethernet network. This wireless network supports a throughput of about 1.5 Mbps. All the computers on this wireless network use WaveLAN adapters (ISA bus version on the desktop PCs and PCMCIA versions on laptop computers). In addition, we have a wireless bridge (a Solectek Airlan Access Point) that forwards packets between the wired and wireless networks. In this fashion, the computers on both sides of the bridge appear to be connected together in one Ethernet. Of course, the wired side of the network has a bandwidth of about 10Mbps while the wireless network has a bandwidth of only about 1.5 to 2 Mbps. (We will refer to the wired Ethernet as *VideoNet* and to the wireless Ethernet as *WaveLAN* network.)

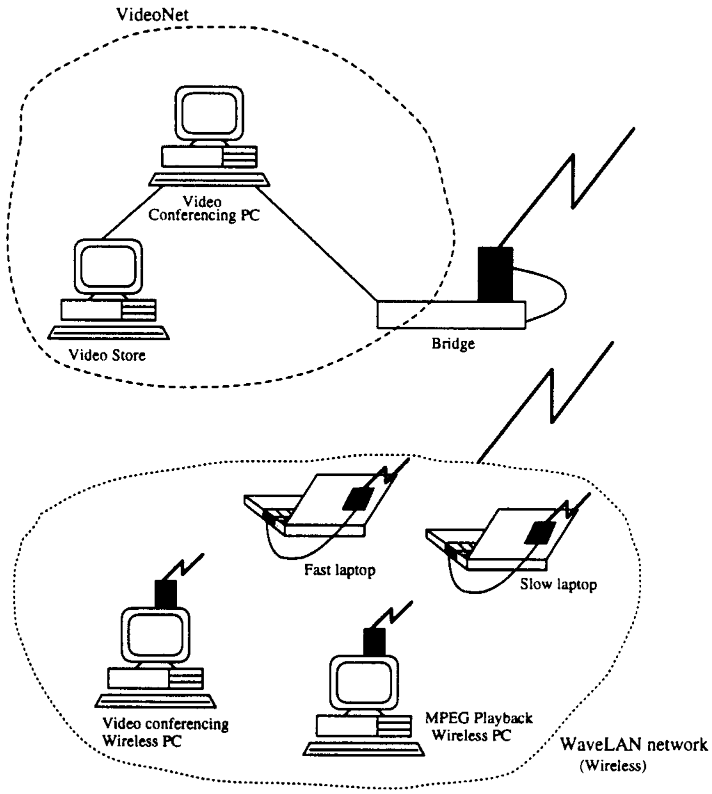

Figure 1: Wireless Video Network Infrastructure

The VideoNet computers do not use standard Ethernet protocols, but rather use the StarWorks multimedia streaming protocol which is a protocol designed to deliver multimedia data smoothly over Ethernet networks. (StarWorks is a commercial video server software system.) All the computers on the VideoNet and the WaveLAN network run Microsoft's TCP/IP stack as well as IPX.

The main software multimedia component in our system is the StarWorks video server. A video server is the item in charge of delivering several video streams in multiple formats to a number of clients on the network. The design of a video server is quite complex and indeed it is one of the main items in our research agenda. For the purposes of this note we will only describe our work with StarWorks and omit mention of our research in this area.

Figure 1 shows a schematic picture of our facilities. The *Video Store* is a Dell 450 (i486) computer containing two large disks storing a variety of video data. This machine is on the VideoNet and runs the StarWorks Video Server software. There are several *Video Conferencing PCs* in our network. One in particular contains an Intel Smart Recorder capture card and a camera to enable real time capture of video. This video capture card captures feed from the camera and converts it into an AVI stream using the Intel Indeo R3.2 compression codec. (Indeo is yet another compression scheme; it is different from the MPEG schemes described previously and it is somewhat less CPU-intensive to process.) The InVision video conferencing software installed on this computer allows video conferencing between this computer and another similarly equipped computer. In addition we have a *Video Conferencing Wireless PC* which is a Dell 466 MXV computer with a ISA WaveLAN adapter (desktop version) and an Intel Smart Recorder card with a camera. This machine is currently used to simulate a powerful PDA of the future that will have two way video conferencing capabilities. Such a PDA would require a small camera as well as some hardware for real time capture.

Since processing power considerations are very important to us we have in our network both a *Slow* and a ***Fast laptop***. The Slow laptop is a 386 based Dell 320 laptop computer with a PCMCIA WaveLAN adapter. This machine is currently used to study effects of a slow CPU on networking protocols. (With our current setup, for the most part, the wireless bandwidth is the limiting factor; however, we would like to be able to study the effects on protocols when the CPU becomes the limiting factor.) The Fast laptop is a 486 NEC Versa laptop with a PCMCIA WaveLAN adapter. This machine also runs the InVision software which means it can receive video feed from a PC running the InVision system with a camera. It functions, therefore, as a one-way video phone.

Finally, there is an MPEG playback Wireless PC which is a Dell 466 MXV computer also has an ISA WaveLAN adapter and a Sigma Designs MPEG playback card installed in it. This machine is currently used to simulate a PDA of the future that will have MPEG video playback facility in hardware. (Note that in the very near future, we plan to use a Panasonic Pentium laptop with an MPEG playback module to replace this machine; this laptop plays MPEG back with hardware support using an add-on module that slips into the slot occupied normally by the floppy drive.)

4.2 Initial Experiences

We experimented sending video clips from the Video Store PC to the wireless PCs and laptops over the wireless network. In this experiment we used the StarWorks multimedia streaming protocol which tries to ensure smooth delivery of data isochronously. Notice that the wireless network bandwidth of 1.5Mbps is more than the bandwidth afforded by many single speed CD-ROM drives (150KBs). We played back both AVI and MPEG movie clips to see what kind of performance we were likely to obtain.

AVI files with a resolution of 240x180 required a bandwidth of about 1.5 Mbps. Using the software-only Indeo 3.2 codec, we noticed that there were frequent audio and video packet losses. The reason for packets being dropped is that inherently the codec does not compress data as much as other formats (such as MPEG) thus requiring a larger bandwidth.

MPEG files with a resolution of 320x240 required a smaller bandwidth, about 1.2 MBps. Using the hardware-assisted playback board, we were able to obtain fairly smooth playback. Even with this hardware support, occasionally (depending on scene changes and the amount of action involved) there were frames where some visual disturbances could be noticed or audio disturbances heard. Of course, software-only decompression of MPEG files yielded very poor performance with 486 computers. We plan to use Pentium PCs to do software MPEG decompression in the near future.

In addition, we also experimented with the InVision system to attempt video conferencing over the wireless link between two PCs. Notice that since this stream of data is only in AVI format, we had to drop the bandwidth down to about 500Kbps to receive data reliably. Of course, with such bit rates, we had to use small window sizes (160x120) and a frame rate of about 5 per second. In spite of this, we noticed very poor quality video, and (worse still) poor quality audio as well. When the laptop was used as the receiver of the video feed from the wired PC, we noticed that the reception was even worse; only about 2 frames per second were effectively being displayed on the receiver side.

Our overall experience was that, with limited network bandwidth (about 1.5 or 2Mbps as with

WaveLAN for instance), and well designed network protocols ensuring isochrony, reasonable quality video reception was possible provided we had hardware support for decompression. And, if hardware support is available, MPEG is the format that provides superior quality video because of its higher compression rates.

However, our current experiments raise a number of other issues. First, we only experimented with a single video stream. For an environment such as our research laboratory we would expect a need to support many additional streams. In addition, changes to current communication protocols are needed. Our current plans call for research in this latter area. Furthermore, there are two alternatives to achieve suitable decompression. One is to add additional hardware to a mobile machine. The other is to use higher capacity CPUs to do software decoding. We are in the process of studying the tradeoffs involved in this decision. Lastly, although our our current video server was able to deliver a single stream well, we are not sure if it will scale to a large number of users (mobile and otherwise). We are presently designing algorithms and developing software for a larger scale video server.

5 Conclusions and Current Work

In this note we have discussed a number of issues related to mobile multimedia. Our research approach in this area has been experimental. Our first step has been to construct a mobile multimedia facility with off-the-shelf components. We have conducted some experiments that have shed some light on what technological issues we need to address to develop a practical mobile multimedia system. While we have only done some initial work in this area, we are convinced that this is an exciting topic of research that encompasses numerous and interesting research problems.

References

[1] D. Molta and J. Linder, *Wireless Act: Wireless LANs Making Big Sense, Network Computing, 6 (8), 82-102 (1995)*

[2] J. Ioannidis, D. Duchamp and G. Q. Maguire Jr., *IP-based Protocols for Mobile Internetworking, Proc. SIGCOMM'91 Conf.: Commun. Architectures, Protocols, Sept 1991, 235-245*

[3] D.B. Johnson, A. Myles, and C. Perkins *The Internet Mobile Host Protocol (IMHP), Internet Draft, Feb. 1994*

[4] F. Teraoka, K. Claffy, M. Tokoro *Design, Implementation and Evaluation of Virtual Internet Protocol, Proc. 12th Int. Conf. Distrib. Computing Syst., June 1992, 170-177*

[5] Ramon Caceres, Liviu Iftode *Improving the Performance of Reliable Transport Protocols in Mobile Computing Environment, IEEE J. Select Areas Commun, 13 (5), 850-858, (1995)*

[6] P. Manzoni, D. Ghosal, G. Serazzi, *Impact of Mobility on TCP/IP: An Integrated Performance Study, J. Select Areas Commun, 13 (5)858-868, (1995)*

[7] B. R. Badrinath, T. Imielinski, *Querying in a Highly Mobile Distributed Environments, Proc 18th Conf. Very Large Databases' 92, Vancouver, B.C. Canada, Aug 1992*

[8] G. Karlson, M Vetterli *Packet Video and Its Integration into the Network Architecture, IEEE J. Select Areas Commun, 7 (5), 739-751, (1989)*

[9] A. Jones and A. Hopper *Handling Audio and Video Streams in a Distributed Environment, Proc. 14th Symposium Operating Systems Principles, 231-244, Ashenville, Dec. 1993*

[10] M. Woo, N. Prabhu, A Ghafoor *Dynamic Resource Allocation for Multimedia Services in Mobile Communication Environments, J. Select. Areas Commun, 13 (5), 913-923, (1995)*

[11] L. Filippini, *MPEG-FAQ, http://www.crs4.it/HTML/-LUIGI/MPEG/mpegfaq.html, July, 1995*

[12] D. Le Gall, *MPEG: A Video Compression Standard for Multimedia Applications, Communications of the ACM, 34 (4), 46-58, (1991)*

[13] Jan Ozer, *Video Compression for Multimedia, Academic Press, (1995)*