Introduction

Video-on-Demand, teleconferencing, digital libraries and remote learning and tutoring systems are just a few of the many distributed multimedia applications that have emerged in recent years. Widespread use of such applications will require networks that can offer *integrated* transport services to both *discrete media* (such as text and digital images) and *continuous media* (such as audio and video). This book discusses recent technological developments that have made large scale deployment of such integrated networks closer to realization. Many issues in the design and implementation of integrated networks remain to be resolved and are the subject of current debate in the research community. This book discusses these unresolved issues as well.

The Challenges of Multimedia Networking

With networking developments advancing at almost a miraculous pace in the last two decades, one is tempted to think that constructing networks capable of transporting multimedia services is a possibility today. At least it may be a trivial extension of today's capabilities that will occur naturally once bandwidth and processor technologies cross a certain threshold in the not so distant future.

In fact, constructing networks to efficiently and cost-effectively transport multimedia services in a ubiquitous fashion requires a rethinking of some of the long-held paradigms of networking. We discuss the challenges in multimedia networking below.

Multimedia Traffic Demands

Multimedia traffic can present significant demand on network resources. This is due, in part, to the fact that much of this traffic is generated through the digitization (into bits) of originally analog signals such as video and audio. Unprocessed, the output of such digitization process can represent a very heavy load on networks.

As an example consider a video signal which requires 30 video frames per second and where each frame is represented by 512×512 pixels. If each pixel is represented by 16 bits, transmitting this video stream would require approximately 125 Mbps, a data rate in excess of the capacity of 12 Ethernets!, and this is for just one source. Table 1 shows some approximate numbers for some common types of audio and video signals. Clearly

Type	Digitized (unprocessed) Data Rate
Telephone Quality Audio	64 Kbps
CD Quality Audio	1.5 Mbps
Broadcast Quality NTSC Video	120 Mbps
Studio Quality NTSC Video	200 Mbps
HDTV, Scientific Visualization	1-2 Gbps

Table 0.1: Unprocessed Data Rates for Some Media Types

wide spread deployment of services that generate the data rates shown in this table is not feasible for existing or even future networks. It is, therefore, generally accepted that multimedia applications will need to process the data they generate in order to reduce the required data rate demands (i.e., *compress* the data) before presenting this data to the network. This compression does not in general come without costs. It can represent delays at the sources and receiver. It also adds a certain level of unpredictability to the volume of generated data.

Compression techniques and issues are the subject of the discussion of Chapter 1 of this book.

Intra- and Inter-media Synchronization Requirements

An important aspect of the delivery of multimedia data across a network is the timing of such delivery and particularly its relationship to the timing of delivery of other related data. Existing data networks never had to deal with such requirements. There are actually two types of timing (or *synchronization*) requirements.

- **Intra-media synchronization** where the timing of the delivery of data associated with a particular (audio or video) stream needs to correlate with the timing of other data associated with the *same* stream. For example if video frames are generated at a source at 30 frames per second, it would make the receiver's job easier if the data packets representing consecutive frames are received 1/30th of a second apart (with some tolerance factor).

- **Inter-media synchronization** where the timing of the delivery of data associated with a particular (audio or video) stream needs to correlate with the timing of data associated with *another* stream. This other stream could be originating at the same source or at some other source within the network. For example the video and audio of a particular movie may originate at different locations. They still need to be delivered and played back in synchrony with each other. The human senses can tolerate some asynchrony (a few 100 milliseconds) in the synchronization of audio and video and of audio and subtitles.

Chapter 2 of this book deals with multimedia synchronization problems and solutions.

Network Performance Guarantees

In the past, communication services were provided by different types of networks: telephone networks for voice, cable TV networks for video, and data networks for computer communications. In contrast, multimedia networks integrate communication services for all media types in a single network. It is widely accepted that multimedia networks should be based on a packet-switching (or cell-switching) paradigm[1]. Due to their ability to perform statistical multiplexing at network switches, packet-switching networks achieve a high utilization of networking resources. However, statistical multiplexing introduces (sometimes unpredictable) queueing delays at network switches, resulting in variable network delays of transmitted packet. This unpredictability can be a major hurdle in meeting the synchronization requirements discussed above.

Network applications that involve discrete media, such as file transfer or electronic mail, are only moderately sensitive to amount and variations of the network delays. However, most continuous media applications are delay-sensitive and require performance guarantees for the transmitted data. In addition,because many multimedia applications are interactive, they require an assurance on the maximum end-to-end delay to maintain their interactive nature.

The set of performance guarantees that a multimedia network can provide to its applications is generally referred to as *Quality-of-Service* (QoS). QoS may include guarantees on the throughput, the network delays, the delay jitter, and the error rate. Providing QoS in packet-switching multimedia networks without completely sacrificing the gain of statistical multiplexing has been the subject of much recent research activity. This work is discussed in Chapter 3 of this book.

Multipoint Communication Requirements

Many multimedia applications involve multiple (more than two) participants simultaneously. This is in contrast to many of the currently popular networked applications (like file transfer and electronic mail) which typically involve only two participants (a sender and a receiver or a client and a server). Multiple participant multimedia sessions have necessitated extending the underlying network support to allow for efficient point-to-multipoint and multipoint-to-multipoint data flows. Many well-understood network procedures like routing and flow and congestion control have to be revised considerably in order to provide efficient network support for multipoint communication. We examine developments in this area in Chapter 4 of this book.

[1] In the context of ATM networks, *cells* refer to small fixed-size packets.

Application Development and Deployment

The impetus for the development of networks capable of transporting multimedia traffic stems from the desire by many to see widespread use of multimedia applications. Even after the network support is in place development of such applications will present many challenges. We present some example multimdeia applications in Chapter 5.