Client-Server Metadata Management for the Delivery of Movies in a Video-On-Demand System

T.D.C. Little and D. Venkatesh
Boston University*

Abstract

In this paper we describe a metadata mechanism for location, identification, and delivery of continuous media in the form of digital motion pictures and in the context of a distributed system. The metadata mechanism and supporting protocols are based on the client-server model. The mechanism can be applied to numerous application domains including multimedia-based home entertainment, catalog shopping, distance learning, and distributed-interactive classrooms. We demonstrate the practicality of our the mechanisms through a prototype application based on home entertainment.

Keywords: Internet-based services, metadata management, client-server architectures, video-on-demand.

1 Introduction

The Internet is expanding at a very rapid pace, with almost a million new users being added every month [2, 9]. Many services are being added to satisfy the varying needs of the different groups of customers who use the network. Some of these services make use of continuous media like video and audio that typically consume enormous bandwidth and strain the resources of the network as well as local data stores.

However, technology has kept pace with this demand and communication services capable of supporting high bandwidth applications have been developed. This communication ability, coupled with the decrease in memory costs, and increase in storage device capacities will create an excellent opportunity for information providers to develop applications that will use distributed resources. Interactive multimedia applications are fast becoming the most popular applications being developed to exploit the new infrastructure. These applications including services such as video-on-demand (VOD), distance learning, and catalog browsing will soon become the prime consumers of network resources [2, 9].

As these bandwidth and memory intensive applications evolve, it becomes important to develop schemes which efficiently utilize the system resources while simultaneously providing the best service to their users [1, 3, 10, 11]. To be successful, these applications will need to make enormous volumes of data easily accessible to their users. To make the handling of data manageable, the information will need to be distributed across different sites. One view, specific to the distribution of home entertainment, is shown in Fig. 1 [5].

This model consists of a local database connected via a high speed backbone network to archives where multimedia information are stored. The local databases cache the information locally (on a very large scale), and make them available to their users on-demand [14].

The choice of data distribution is based on the desire to support many users. The support of a single VOD session[1] requires substantial bandwidth [11]. To support many users, significant system communication and storage I/O are required. Local storage on the end user's computer is not feasible due to the high cost per user. The distribution of data and use of multiple access points increases the system's overall I/O capacity to support interactive sessions. Furthermore, a distributed system can have more access points which the users can use to connect to the system. Experimental testbeds and deployments to date are either simplistic in functionality, or do not address issues of scale [4, 8, 12, 13]. Future systems of this type must support thousands of users.

*The authors can be contacted at the Multimedia Communications Laboratory, Department of Electrical, Computer and Systems Engineering, Boston University, Boston, Massachusetts 02215, USA, (617) 353-9877, *tdcl@bu.edu*. This work is supported in part by the National Science Foundation under Grant No. IRI-9211165.

[1] A session is described as a single, truly interactive connection from the system to the user

0-8186-5835-5/94 $03.00 © 1994 IEEE

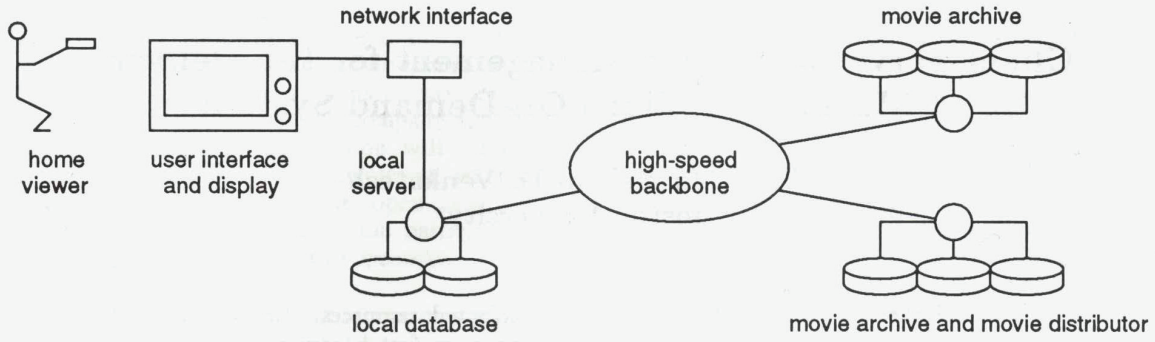

Figure 1: A Distributed Database Scenario for VOD

Using this system model we seek to support information access and delivery services for interactive multimedia applications. To this end, we investigate service mechanisms for building such applications which require location, identification, and delivery of multimedia data. More specifically, we examine services for retrieval of distributed multimedia data in a client-server architecture in which a server coordinates different client requests. We also present an experimental prototype we have developed in our laboratory which demonstrates a subset of the aforementioned concepts.

The remainder of this paper is organized as follows. In Section 2, we develop our query architecture based on the concept of resource and metadata servers. Section 3 describes the application of these concepts in the context of a client-server VOD architecture. In Section 4, we describe the prototype testbed we have built in our lab. Finally, Section 5 concludes the paper.

2 Services for Supporting Interactive Multimedia Applications

A critical issue in the design of multimedia information systems is the data distribution architecture. Due to the enormous volumes of data involved, it is unwieldy to completely store and manage all information distribution from a single site. Furthermore, the interactive nature of the queries to the database requires the database manager to provide both database management and continuous media delivery functionality.

Services required to support interactive multimedia applications include data distribution, metadata management, and resource management. These are described in the following subsections.

2.1 Metadata Management

To simplify the query mechanism, and increase interactivity, most database systems employ a "metadata" mechanism. Metadata are "data about data." In our context, this means that they contain concise information about the location and characteristics of the data to be retrieved (e.g., movie titles, where they are stored) [6]. By using such a scheme, the data delivery process can be decoupled from the database management functions.

The advantage of this metadata approach is that it allows the user examine the contents of a database without necessitating the retrieval of the actual data objects. The user can quickly go over an information summary and select the information to be delivered. This concept is particularly appealing with respect to stored audio and video which can be costly to deliver and can consume significant portions of a data server's I/O bandwidth.

There are several requirements for the metadata management scheme necessary to support interactive multimedia applications. First, we require basic services to support object identification, location, naming, and distribution [7]. Second, application-specific models are required to support application functionality (e.g., searches on movie scene content). In the former, we can also consider the limitations of available I/O bandwidths for individual storage devices at data servers. For example, this limitation can be managed by a central site in the form of a resource state-table for sessions in progress (e.g., movies played out off a specific server). The attributes managed by the central site include: *location, availability, current state*