Chapter 2

Current and Emerging Applications

In this chapter, we describe some application scenarios and experimental prototypes based on multimedia database management systems technology. This gives us an idea of typical application environments for multimedia database management systems and the different requirements that characterise these environments. We first give some examples of database support for applications with much user interaction. Secondly, we look briefly at experimental systems addressing content based retrieval of multimedia data. Finally, we summarise some other application domains that may use multimedia database systems. Obviously, there exist many more experimental systems which we cannot cover here. We only want to give the reader an idea of the wide spectrum of applications and their potential with respect to multimedia technology and database systems.

Multimedia databases are often needed to support interactive processes. First, we give an example of a publication scenario. Both the information providers and the customers make use of the same overall system to keep track of the information. The second example shows the support of design and evaluation of complex technical systems. In this example, information technology has been applied to assist a group of designers with the specification and documentation of the design task. Finally, we demonstrate the application of a multimedia database system to support teleservices.

Furthermore, multimedia database systems may be used to provide access to data that could previously only be accessed manually. First, we describe a system for the storage and retrieval of images. We will see how the application of modern techniques helps us to deal with huge amounts of images. Next, we look at a system for document archiving. We conclude the chapter with a sample prototype of a digital library, that is automatically filled with objects.

Finally, some other application domains with characteristics different from the previous examples are summarised.

2.1 A Multimedia Publication Environment

At GMD-IPSI a prototypical implementation of an electronic multimedia magazine, called `MultiMedia Forum` [S+94] has been developed. It consists of an *editing environment*, where information providers can produce and combine multimedia information, and of a *reading environment* where information consumers can access the information. Therefore, the `MultiMedia Forum` serves as an example for the entire *multimedia publication process* which goes far beyond the traditional desktop publishing paradigm and current solutions and approaches reflected by World Wide Web-based publication models. The publishing processes supported by the `MultiMedia Forum` make use of an underlying database system. The general impact of using advanced database concepts in such an environment is discussed in [ABH94]. Figure 2.1 shows the multimedia publication process. The functionality supported by this prototype can be grouped into three main functions as follows and it will be described in some more detail below.

- *Information import:* this process covers the creation and the acquisition of information from the authors or editors, and the transformation of import formats into formats used internally for further processing or information exchange with other systems.
- *Information processing:* this process consists of storing, indexing, retrieving, layout preprocessing, and manipulating documents, which are modeled by means of a rich semantic document model (e.g. a hypermedia document model).
- *Information export:* this process deals with the export of information to other environments and the distribution of multimedia documents to users. This includes important features like administrative functions such as access control mechanisms, workflow-management, and accounting.

Information Import: Various types of input information are imaginable. Paper or electronic documents, analog or digital media such as graphic, video clips or audio, numerical data and hypertext documents are accessible via various information bases. Such sources can be, for example, paper archives, the user's mailbox, private archives, heterogeneous globally accessible or online databases [YGM94], video/audio file-servers or public archives. The first step in a publication process is the *digitalisation* of analog data input. All digital data need then to be prepared

(i) to be converted into an appropriate document format (for example SGML - Standard Generalised Markup Language [ISO86a]) by proper tools (for example, DREAM [GF92]),

(ii) to be compressed according to some standard compression techniques (for example JPEG [ISO92c], MPEG [Gal91]), or

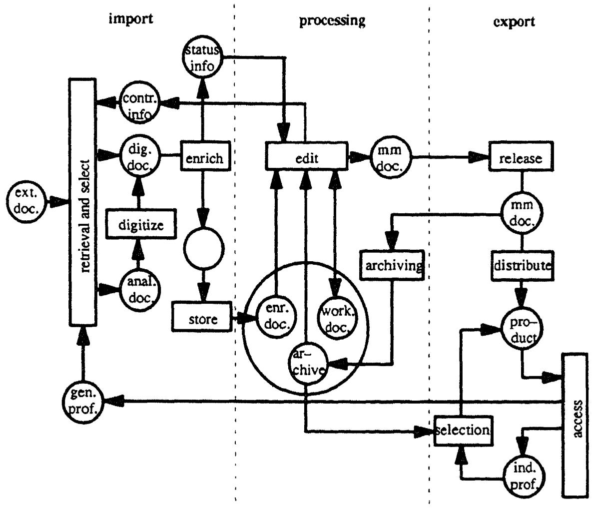

Figure 2.1: The Multimedia Publication Process Model at GMD-IPSI

(iii) to be preprocessed, for example, identifying and extracting the abstract of a multimedia document, categorisation of documents, the creation of metainformation (document content describing information), and the actual insertion into the multimedia document pool.

The input to the document pool of the `MultiMedia Forum` is currently hypermedia documents structured by means of SGML and containing linked multimedia content data that will be extended to HyTime [NKN91] in the future [BA94].

Information Processing: Documents are stored in a multimedia document pool based on an object-oriented database management system with multimedia extensions. Tools used by the authors and editors retrieve and load documents stored in the database by using an integrated retrieval interface providing access to the entire document pool.

Information Export Selection and retrieval facilities are provided by a specific export component, called the readers environment, as well as by the application programming interface of the `MultiMedia Forum`.

Retrieval facilities are based on formal query languages and/or navigational interfaces. The query language can be used in an ad hoc interactive fashion or embedded in a host programming language (for example, C^{++}). Application dependent interfaces, such as a particular document pool navigator, which are based on their own query engine, can better support user needs in a specific application context. The disadvantage of application specific interfaces is a lack of reusability and flexibility.

Another important function available is *releasing information* which is applied before distributing the final approved electronic product. Releasing information triggers amongst other things the (active) distribution of information, for example, a new issue of the magazine. From a technical point of view the distribution is mainly opening and transmitting data to a specific communication channel. These could be networks, an external storage device (for example, disk or CDROM), external or online databases [YGM94], or an electronic archive. Typical new distribution channels are new services such as electronic bookstores or digital libraries.

2.2 Multimedia and Database System Support for Systems Engineering

The design process of technical systems has to be supported by various different software tools. In traditional systems engineering environments the integration of these tools resolved the heterogeneity of the underlying operating and file systems. Integration was achieved by sharing common data (for example documents) by means of files and standard data exchange formats. Under these circumstances it was difficult to manage the dependencies between different documents and to ensure the consistency of documents in a multiuser environment.

The *MuSE*[1] project [DGJ+94] aims to provide integrated system support for the systems engineering process such that dependencies between different specification and documentation documents can be managed automatically and global consistency can be ensured. *MuSE* covers various phases of the systems engineering process including design, verification, animation, and simulation, and follows the concurrent engineering paradigm.

In the *MuSE* environment all information resulting from the design process, verification of specifications, system simulation and animation, and testing is stored in its underlying object-oriented database management system. This includes alphanumeric data as well as graphics, images, audio and video annotations which may originate from simulation and animation results. *MuSE* allows for the storage, retrieval and manipulation of highly-structured information like 3D-data, part structures, and multimedia and hypertext documents in a multiuser environment which needs to be supported by the underlying database management system.

The *MuSE* prototype uses hypermedia concepts to organise the documents of the system development process. The complete system model is represented as a hypernetwork containing the different specifications. The hyperstructure is visualised via a hypermedia authoring environment which provides the desktop for the whole *MuSE* environment. Figure 2.2 shows the architecture of the system. Components to describe technical systems using several problem

[1] *MuSE* is the acronym for a project entitled *Multimedia Systems Engineering*. The project is a joint effort of groups at the Technical University of Darmstadt and GMD-IPSI, Darmstadt. It is sponsored by the *Deutsche Forschungsgesellschaft* DFG, grant numbers He 1170/5-1 and He 1170/5-2.

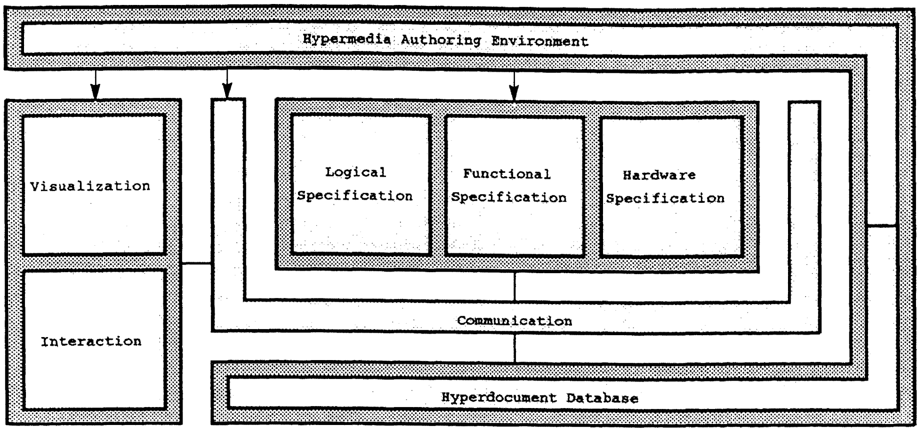

Figure 2.2: Architecture of the *MuSE* Environment

specific specification languages and a visualisation and interaction component are embedded into the hypermedia desktop and database environment.

Storage and manipulation of hyperstructures and the documents contained in the hypernodes are modeled at the database level[KAN93]. In addition to the management of conventional data, the underlying database management system VODAK[GMD95, RL94, AK94a, KNS90, BA94] offers mechanisms for the storage, manipulation and presentation of multimedia data. Therefore, the system allows an extended notion of documents including text, images, audio, video and data produced from the tools of the *MuSE* environment. The object-oriented model of the documents contains methods providing access to all the tools that are necessary to display and edit their content. This means that even complete simulations may be executed calling a method of an object that contains the system model.

2.3 A Multimedia Calendar of Event Teleservice

In the field of teleservices there is a high potential for multimedia applications. Progress in integrating network technology and multimedia information systems allows specific multimedia teleservices to be provided. One specific example from this application domain is a multimedia database system supported archiving teleservice which incorporates electronic mail [TR94]. This teleservice supports both interchange mechanisms of the multimedia mail system used, which is based on the principles of the CCITT Recommendation X.400 Message Handling System. It allows for *store* and *forward* operations to interchange complete multimedia documents and for *access by reference* in order to refer to large multimedia documents available in a global store instead of copying the documents for each e-mail message. The distinguished features of the teleservice are the integration of multimedia mail with simple stand-alone

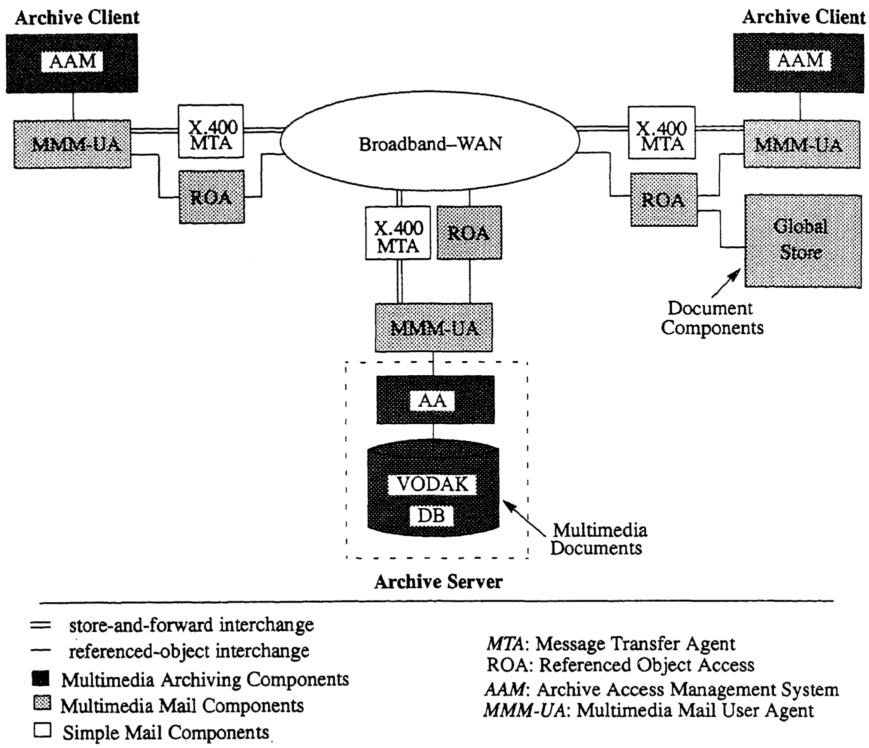

Figure 2.3: Multimedia Archiving Teleservice Architecture

archive clients for heterogeneous platforms and an archive server which is based on a multimedia database system. The archive contains structured multimedia documents which can be retrieved via e-mail requests. Answers to a request are returned to the user via multimedia electronic mail by means of composite multimedia documents.

Figure 2.3 shows the overall client-server based architecture of the teleservice prototyped in the context of the GAMMA project at GMD-IPSI. The multimedia document archive is based on the open object-oriented database system VODAK [GMD95, RL94, AK94a, KNS90, BA94]. The archive and the clients are connected to the Multimedia Mail User Agent (MMM-UA) providing the X.400 Message Handling System functionality. The archive contains complete structured multimedia documents and keeps track of received and sent multimedia mails, (parts of) documents, and access requests. Archived data is accessed by the VODAK Data Manipulation Language [GMD95] using application specific interactive tools at the clients. The global store contains only parts of documents for some limited period of time determined by the user access.

In this sample teleservice the database system used has to manage multimedia documents, but does not need to deliver multimedia data according to some temporal quality of service parameters. This is because of the asynchronous access mechanism which calls for the decoupling of the presentation environment from the database management system. Multimedia data is delivered via multimedia e-mail and subsequently processed by presentation tools at the client site. Obviously, this limits the applicability of this approach because huge amounts of multimedia data are copied within the network. However, a broad range of public multimedia applications, for example, product catalogs for teleshopping, subscription services for multimedia products (titles, publications, etc.), virtual travel agency, cooperative authoring of multimedia documents, etc. can cope with the e-mail delay and, hence, may be based on this approach to a teleservice.

2.4 The QBIC System

Multimedia database systems should provide means for content-based retrieval of the multimedia data. One of the first systems demonstrating such capabilities for the retrieval of images is the QBIC (*Query By Image Content*) system [NBE$^+$93]. This technology will be beneficial for several applications. In a medical environment, one might want to search for images that contain tumors with a certain texture. A travel agency might want to find attractive pictures for several destinations to use in their glossy magazines. This functionality would really enhance a publication system like the one described in section 2.1. Other applications may be found in journalism, art, fashion, retailing, and industry.

Previous systems for image retrieval used manually assigned keywords describing the objects in the database. Queries were expressed using languages like SQL. An obvious problem with this approach is that it is not very likely that people will be able to describe images with keywords in a standardised manner. One person will describe a picture as 'dark' while another person describes the same picture as 'somber'. The use of keywords to search the database will not retrieve the right pictures if the user has a different vocabulary from the system.

Searching for images in QBIC is quite different from querying 'normal' databases, but it illustrates the search process in multimedia databases well. Instead of being exact, searches are approximate and driven by similarity. If the user provides the system with an initial image, the database system retrieves similar images. The user may then select some images that approximate the kind of image he had in mind, and may ask the system for more images similar to this selection. This iterative procedure, using a combination of browsing and searching, is called 'Query by Example'. Basically, the database system does not retrieve images based on explicitly defined characteristics, but assists the user in narrowing the search space.

Developing a QBIC application consists of three logical steps: database population, feature calculation and image query. The first step loads the images into the system. Usually, a 'thumbnail' version of the image is prepared, to be shown to the user during the iterative search process (a thumbnail is a smaller version of the original image). These thumbnails can be stored on hard disk, while the huge (terabyte!) collection of original images may be stored remotely using a tape robot. The user can judge the relevance of the pictures using the thumbnails, instead of having to wait for a slow retrieval of the original picture every iteration.

Of course, other attributes can be added to the images. The user can provide textual attributes for an image and optionally identify the outline of important objects. Unfortunately, detection of the outlines of objects in an image cannot be automated completely, but image analysis techniques can assist the user so that he only has to give an inexact outline.

The second step in a QBIC application is feature calculation. This step can be performed completely automatically and calculates measures indicating the colour, texture and shape of images. An example feature that is easy to understand is the colour histogram of the image. Other features that can be used are contrast, directionality, and area. The purpose is to select features that have approximately the same values for images that a human would call similar. Also, the values should be far from each other if the user judges the images to be dissimilar. More information on suitable features and similarity functions can be found in chapter 8. The features can be indexed with the techniques from chapter 10.

The final step in QBIC is the querying process. As described before, the user can supply an example image to the system. If the outlines of objects have been added during database population, the user can also sketch some dominant lines or edges. The system then calculates a distance between the feature values for the example image and the images in the database. In principle, the user does not have to see the values for underlying features. However, some features may be intuitive enough to be used for direct input from the user. An example of such a feature could be the amount of red in an image. The returned images are displayed as thumbnails. In each step, the user can choose to refine the search using one of the returned images as a new query, display the full version of the image or keep some of the thumbnails in a temporary holding area for further processing.

2.5 Multimedia Document Archives

Within the BERKOM II initiative[2] and the POLIKOM research programme[3] concepts and prototypes addressing the general problem of document archiving are being developed. One sample application for a multimedia document archive was also developed in the **GAMMA** project (see also section 2.3). The documents are event descriptions composed of text, images, graphics, auditions, and video clips. Examples of such document parts are a snapshot of a theatre performance, a video clip about the actors, digitised newspaper critiques, and

[2]see DeTeBerkom GmbH, Berlin.

[3]POLIKOM is a national research and development program for providing telecooperation and telepresence for the German government distributed between Bonn and Berlin. See [HBS92]

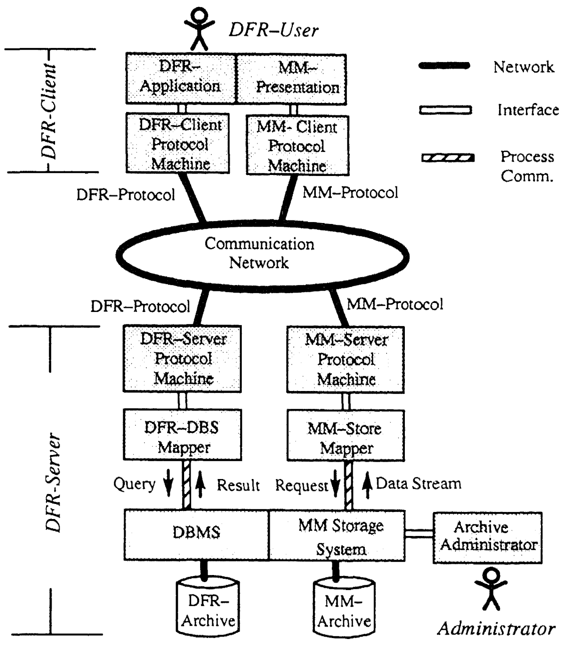

Figure 2.4: Multimedia Document Archiving System

an audio sample presenting a prominent song being performed. The functionality of the archive includes *storage* and *presentation* of multimedia documents, *navigational access* within a hyperlinked structure, *document retrieval*, and *concurrent access* of several users to the same multimedia documents. In contrast to the Multimedia Calendar of Event Teleservice, this application is based on synchronous access to the multimedia document archive: i.e., the archive becomes responsible for the timely delivery of multimedia data according to quality of service parameters relevant for the presentation to the end-user. The archive structure is based on extensions of the ISO/IEC standard 'Document Filing and Retrieval' (DFR) regarding multimedia data handling [R^+94].

Figure 2.4 shows the architecture of the multimedia archive system. The prototype is based on an extended VODAK DBMS [GMD95] which allows for the handling of multimedia documents. The extensions include a multimedia storage component, DBMS interface components including appropriate multimedia transport protocols, and appropriate components at the client site which receive multimedia data via the network and provide for their presentation to the user.

2.6 The Informedia Project

A large project that focuses on content-based retrieval of multimedia documents is the *Informedia Project* at Carnegie Mellon University [HS95a]. Informedia is building a *digital library* that will contain over a thousand hours of digital video, audio, images and text materials. The problem is that simply storing all this data does not provide enough functionality to make such a library useful. Therefore, new technology has been developed that adds search capabilities to this large collection of video data.

As you can see in figure 2.5, the Informedia project is divided into two phases. The first phase is library creation and the second phase is library exploration. In this section, we focus on the first phase.

The Informedia approach to library creation is to apply a combination of speech recognition and image analysis technology to transcribe and segment the video data. The project uses the Sphinx-II speech recognition system [HRT$^+$94] to transcribe the audio. The transcribed data is then indexed to accomplish content-based retrieval. Initially, a highly accurate, speaker independent speech recogniser transcribes the video soundtracks. This transcription is then stored in a full-text information retrieval system.

Because speech recognition is not an error-free process and formulating a query that captures the user's information need is very hard, not all retrieved videos will satisfy the user. However, if the user has to watch the whole video document before he can judge the value of the retrieved information, the system cannot be used efficiently.

When dealing with traditional text documents, we use tables of contents, indices and skimming to quickly find the pieces of information that we need. The time needed to scan a video cannot be dramatically shorter than the real time of the video, so the notion of 'video skimming' is introduced in Informedia. Using image analysis techniques like automatic shot detection (section 8.5.1) in conjunction with analysis of the speech recogniser's output, the essence of the video content can be expressed in a small number of frames. This small sequence of frames is called a 'film strip'. Using the film strip, fast browsing of the video material is possible.

In the News-on-Demand prototype system [HWC95], a library with television news is created completely automatically using the Informedia technology. Of course, an automatic data collection system based on speech recognition is prone to errors. Errors found in experiments with the system included the wrong identification of the beginning and the end of news stories, and incorrect words in the transcripts.

Despite the recognition errors, the prototype system shows big changes in the way we will deal with television data in the future. The system allows us to navigate the information space of news stories interactively based on our interest. Compare this with waiting passively until the next news broadcast, following a path through the same information space that has been planned by somebody else.

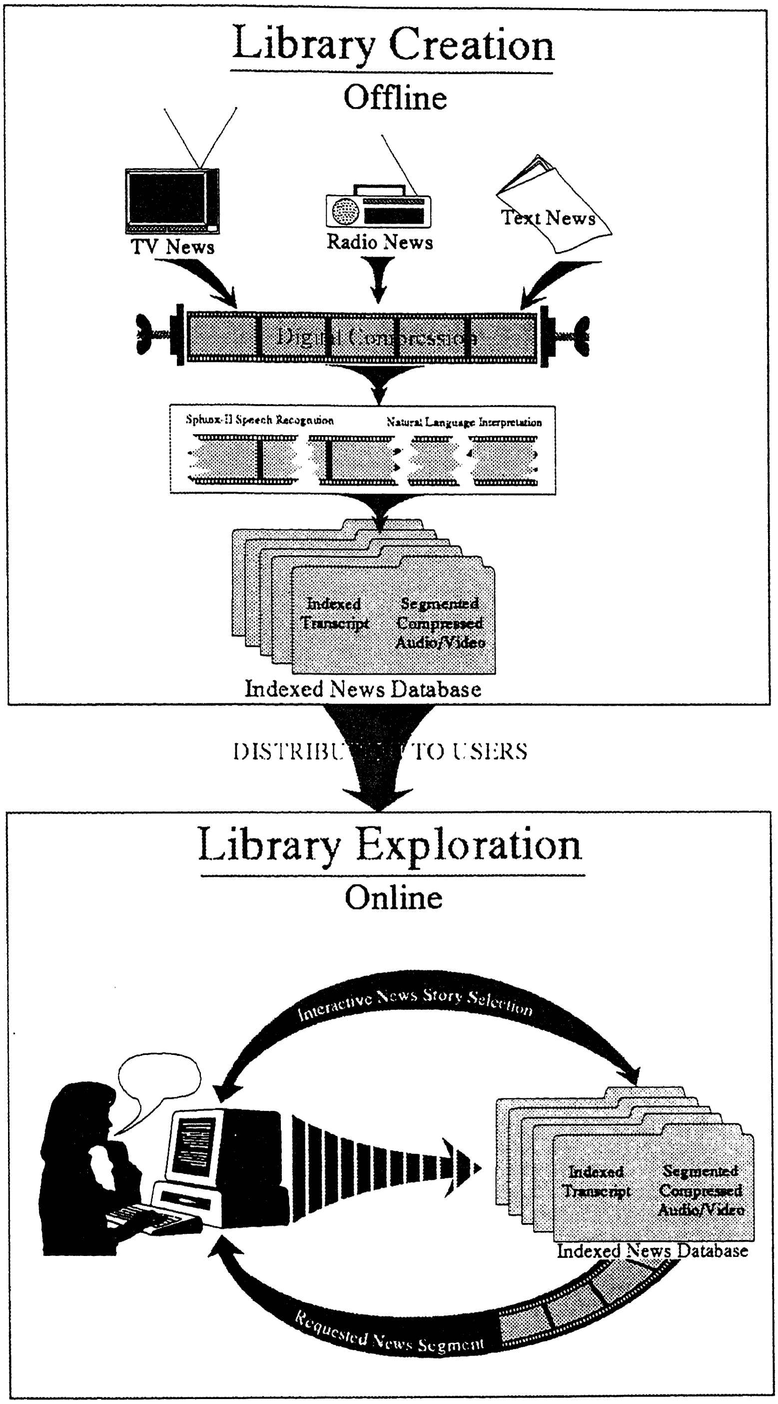

Figure 2.5: The Informedia Architecture

2.7 Some Virtual World Applications

In this subsection we consider some applications that should serve as specific examples for the wide range of virtual environments. The subsection is divided into three parts discussing applications developed at the Visualisation and Media Systems Design (VMSD) department at GMD-IMK: the Responsive Workbench, TELEPORT and a virtual studio system (3DK).

The Responsive Workbench

The Responsive Workbench [KF94, KBF+95] is not a virtual world application, but a virtual environment. It integrates a computer display into a real workbench allowing designers, architects or surgeons to work in a virtual world which simulates their normal working situation. The workbench is a projection-based system in which stereo images of virtual scene and control panel are rear-projected onto an enlarged horizontal display. Projection parameters are set to make objects appear above the workbench surface. The maximum height of objects is about half a meter, limited by the size of the display system.

Users stand in front of the workbench wearing glasses to perceive images in 3D: they can wear a data glove to grasp, move and turn objects and employ a stylus to control the parameters of the display interface. Data glove and glasses have a Polhemus sensor attached to identify the location of the hand and the user's viewpoint for rendering[4] the correct perspective. Since the data glove does not allow for haptic feedback, object grasping is indicated by colourend userisation.

The applications developed so far for the workbench originate from three domains: medicine, architectural design and automotive applications. The medical applications are intended for medical education and aspects of body process representation (for example heart and blood flow). Figure 2.6 shows a model of a human body which is rendered with transparent skin to allow a clear view of the human skeleton. Architectural design applications assist in the visualisation of designs in architecture and environmental design and allow repositioning of objects within a model. Automotive engineers can use the workbench for the visualisation of crash tests (air bag simulations), engine mixing process dynamics and aerodynamical studies.

TELEPORT

TELEPORT [BGA96] is a teleconferencing system allowing users at different sites to meet as if face-to-face in a semi-real and semi-virtual meeting space. The system consists of a real room of which one wall is entirely covered by a display surface. Onto that surface a virtual extension of the real room is projected. A viewer in the room is tracked by an optical tracking system in order to allow for correct perspective rendering. Since a viewer is able to

[4] The process that maps a model to an image is referred to as *rendering*, the software or hardware implementation of this process is referred to as *rendering pipeline*.

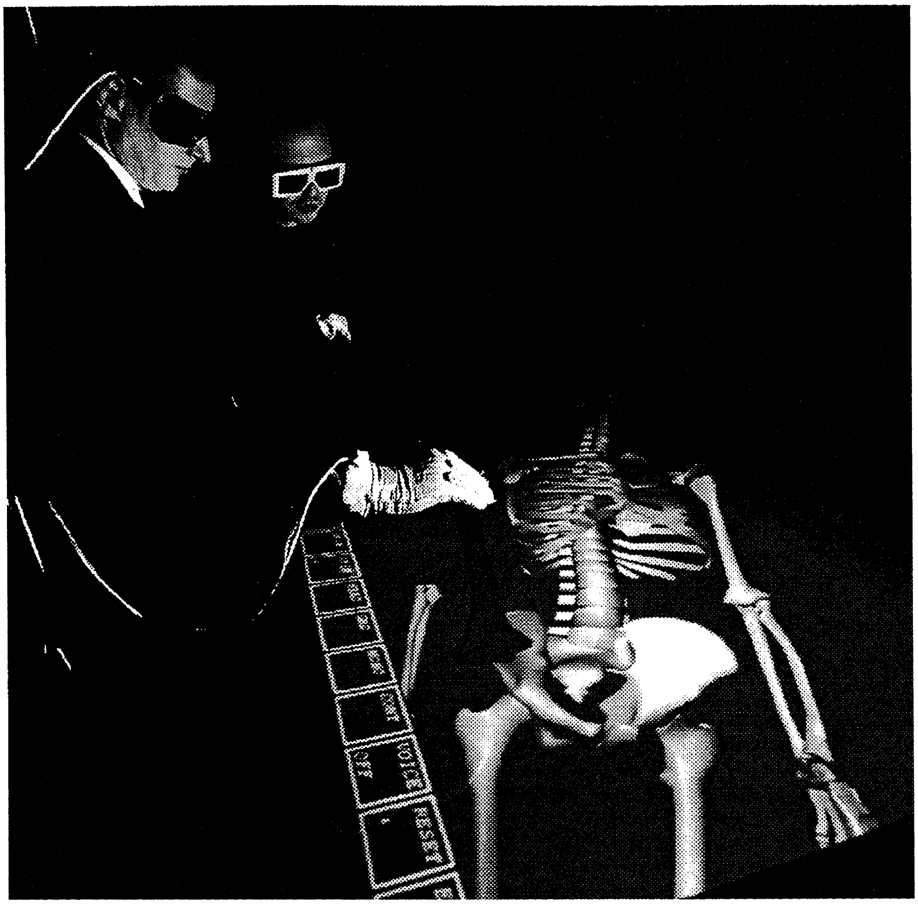

Figure 2.6: The Responsive Workbench

perceive the real room and projected extension at the same time TELEPORT can be called an augmented reality system.

The virtual scene and the real room are carefully furnished and designed to match as closely as possible in colour, lighting and general appearance. Room borders, walls and ceiling are aligned to the real walls and ceiling surfaces. The system is best viewed by wearing stereo glasses (active and passive glasses are possible). The video image of a remote participant is positioned into the virtual extension, giving the illusion of speaking to a person sitting around the same table. In order to produce this illusion chroma-keying[5] of the par-

[5] Chroma-keying refers to situations where objects or more frequently persons are shot in front of a uniformly lit blue background. A chroma-keyer produces a mask signal that is black where the original video frames are blue and white otherwise. The mask signal is used to mix the camera signal (foreground) with any other video signal (background). Wherever

ticipant's video image is required. Since keying normally requires the use of a blue background which is not easily available in an office, a real-time keying technique has been developed that allows a person to be keyed out of any background provided the video camera remains in a fixed position. The keying method, called delta-keying, produces a mask that is used to project the video image on a transparent grid within the model of the virtual extension. As a consequence, the image of the participant is part of the model and follows the same geometric transformations as the virtual extension.

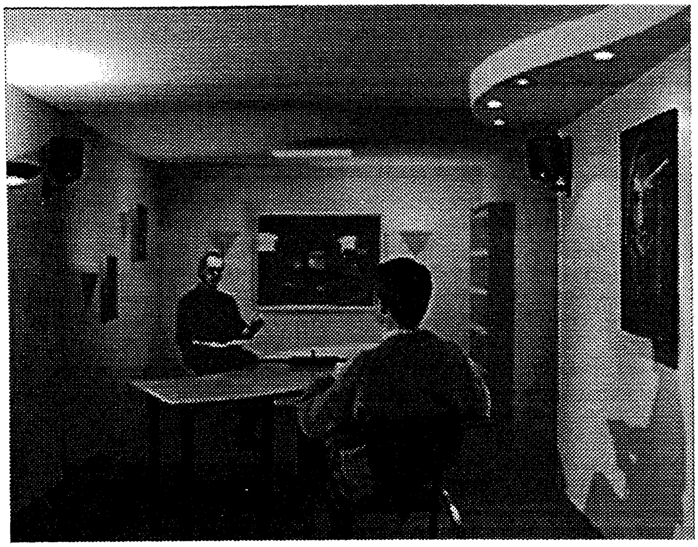

Figure 2.7: TELEPORT

The real room is further equipped with a video camera, loudspeakers and microphones. Video image and audio of a meeting participant are sent over the network. For this purpose codecs have to be used to compress the video. Since the video processing and video encoding introduce delay, audio has to be delayed to synchronise with the corresponding video information. Trials have been made between GMD headquarters at Sankt Augustin (near Bonn) and the University of Geneva. For these trials 24 Mbps ATM lines were used.

A Virtual Studio System

The term 'virtual studio' refers to a new production technique exploited by broadcasters and video studios in order to reduce production costs. Instead of using real sets and props that have to be built, removed, stored and probably reused, productions take place in a blue room and sets and props are generated

the mask signal is white the foreground signal is used for the final mix, wherever the mask is black the background appears.

in real-time by a graphics supercomputer. The technique is based on traditional blue room technology that has existed for more than three decades, but does not allow for camera movement. The use of camera operations such as pan, tilt and zoom, would make it apparent that there is no coupling between the foreground camera signal (for example a moderator) and the computer generated set in the background. To overcome this disadvantage, virtual studios use sensors attached to the cameras used to measure all the changes of camera parameters. These changes are transferred to the graphics computer to adjust the view point in the geometric model from which the background images are rendered. Foreground and background video images correspond to each other even if the camera is moving and in doing so, give the impression of a moderator actually being within the computer generated set.

Although the underlying ideas of virtual studios date back to the late eighties when the Japanese Broadcast Corporation NHK developed *Synthevision*, it was only three years ago that hardware was developed that met broadcasters' requirements concerning image quality and system usability in general. At the moment there are about half a dozen players in the virtual studio market, selling products to video studios and broadcasters and offering virtual studio services. [GAB$^+$96] gives a state-of-the-art report.

Broadcasters' requirements for highest image quality have a series of consequences: first, rendered images must not have any jumps or jitter, and have to be produced at full frame rate (25 frames per second for PAL video). To be more precise, images are rendered at a rate of 50 frames per second, since rendered images have to be scan converted from graphics to video.[6] Second, high quality mixing of foreground and background requires the use of digital mixers, so all virtual studio systems operate with digital video. Third, scene modelling has to compromise between rich and pleasant visual appearance and the complexity of the scene. If models include too many polygons, rendering may not keep up with fast camera movements and the illusion of a person actually being in the virtual set may be disturbed or even lost.

2.8 Other Emerging Applications

There are many other application scenarios which involve the storage, processing, and retrieval of multimedia data. We will briefly review a few additional application domains to give a more complete picture of the potential use of multimedia database management system technology.

- *Multimedia document management and processing* is a quite natural and general application domain for multimedia databases. Multimedia document management will be needed in various, more specific application domains like CAM, technical documentation of product maintenance,

[6]The necessity to render 50 fps results from the different nature of RGB and PAL/NTSC video signals. PAL/NTSC video frames consist of two interlaced fields. One rendered frame is used to produce one of these fields.

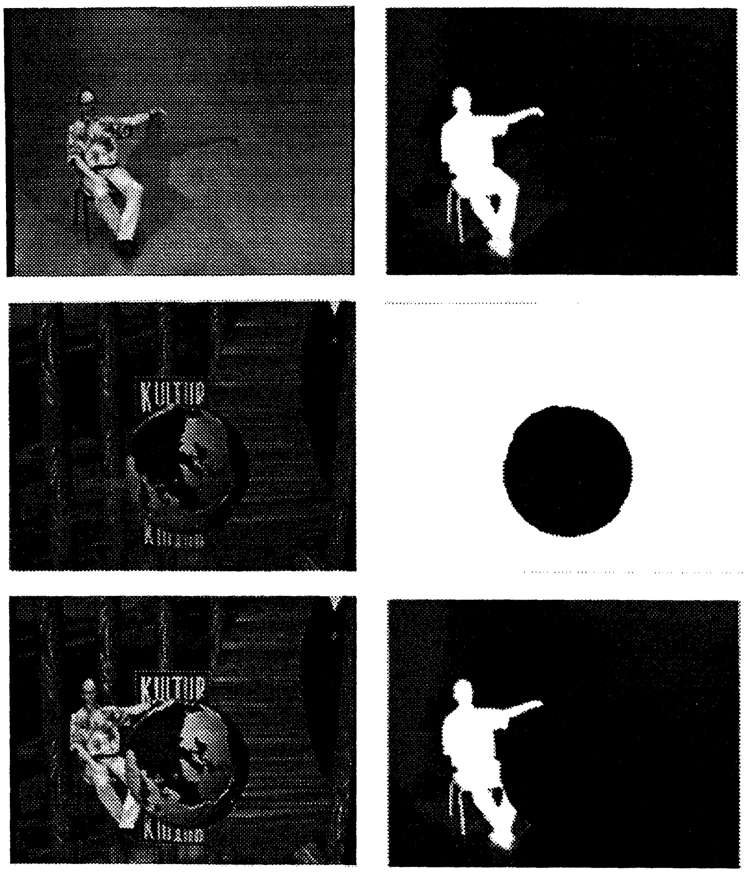

Figure 2.8: Virtual Studio Signals: Foreground Signal (Top Left), Foreground Mask (Top Right), Background Signal (Middle Left), Background Mask (Middle Right), Mixed Signal (Bottom Left) and Mixed Mask (Bottom Right). Since all Images are Presented as Halftones it Has to be Noted that only Mask Signals are Actually Halftones.

education, geographic information systems, teleservices, etc. All these emerging applications share the need for management of multimedia documents. [BA94] discusses the technical issues involved in multimedia document management based on HyTime. [TR94] presents an approach to multimedia document management offered as part of a teleservice.

- *Multimedia Mailing Systems* are an advanced form of electronic mailing systems which integrate various applications like multimedia editing and voice mail. Examples of related projects are the voice-mailing system Etherphone developed at Xerox PARC [Vin91], the MICE (Modular Integrated Communications Environment) project of Bellcore, and the multimedia mail and archive projects of BERKOM (for example, [TRR94]). This type of communication environment may significantly benefit from the availability of a multimedia storage system which serves as a repository for multimedia messages in the network.

- *Teleconferencing* involves multimedia data in several ways. The participants in a teleconference can communicate to each other via audio and video channels at the same time. In practice this was achieved by using dedicated equipment and lines, and specially designed conference rooms. The techniques of communicating the material used by the participants within a teleconference session may be integrated with the technology providing for the communication between the participants. That is, one can imagine that teleconferencing may be based on multimedia workstations which provide an integrated workbench for computer-based teleconferencing. [SGHH94] discusses a prototypical implementation of such a teleconferencing scenario for integrated meeting support across electronic whiteboards, local and remote workstations. In [HBS92] technical challenges for computer-based telecooperation and telepresence in a highly distributed government administration are described.

- *Kiosk information systems* are an example of a highly distributed multimedia information system which consists of servers managing and maintaining huge amounts of multimedia data which are made available to customers at clients' sites. The document archive containing a Calendar of Events [R$^+$94], as described in section 2.5, is an example of such a multimedia information system which allows customers to enquire about events like concerts, festivals, and exhibitions in particular metropolitan regions, and to order or buy tickets for such events. A special type of kiosk information system could be set up for the purpose of home shopping. Customers would dial the retailer of their choice and connect to the retailer's server in order to conduct a sale, or at least to prepare a final sale. In the future, we may also be able to visit trade fairs and exhibitions by means of an online virtual trade fair and virtual exhibition which may offer teleconferencing features to establish contacts between customers and product suppliers.

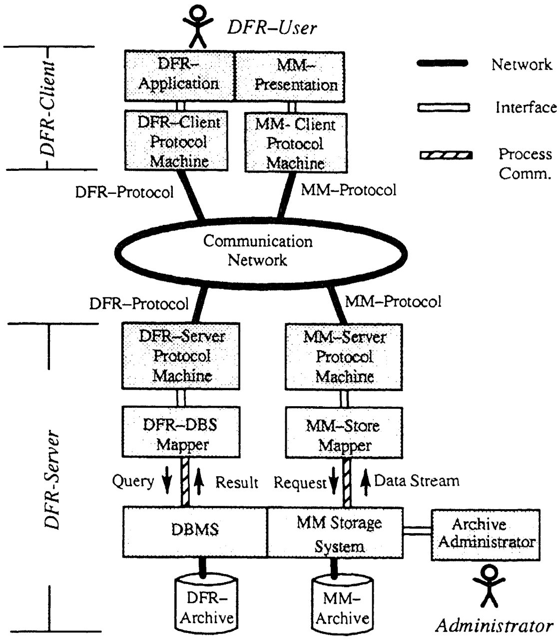

Figure 2.4: Multimedia Document Archiving System

an audio sample presenting a prominent song being performed. The functionality of the archive includes *storage* and *presentation* of multimedia documents, *navigational access* within a hyperlinked structure, *document retrieval*, and *concurrent access* of several users to the same multimedia documents. In contrast to the Multimedia Calendar of Event Teleservice, this application is based on synchronous access to the multimedia document archive: i.e., the archive becomes responsible for the timely delivery of multimedia data according to quality of service parameters relevant for the presentation to the end-user. The archive structure is based on extensions of the ISO/IEC standard 'Document Filing and Retrieval' (DFR) regarding multimedia data handling [R^+94].

Figure 2.4 shows the architecture of the multimedia archive system. The prototype is based on an extended VODAK DBMS [GMD95] which allows for the handling of multimedia documents. The extensions include a multimedia storage component, DBMS interface components including appropriate multimedia transport protocols, and appropriate components at the client site which receive multimedia data via the network and provide for their presentation to the user.