SEPTEMBER 1991, Volume 14, No. 3

the quarterly bulletin of the
IEEE Computer Society
technical committee on

Data Engineering

CONTENTS

1951-1991
40 YEARS OF SERVICE
IEEE COMPUTER SOCIETY

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC
IEEE

SPECIAL ISSUE ON MULTIMEDIA INFORMATION SYSTEMS

Data Engineering Bulletin is a quarterly publication of the IEEE Computer Society Technical Committee on Data Engineering. Its scope of interest includes: data structures and models, access strategies, access control techniques, database architecture, database machines, intelligent front ends, mass storage for very large databases, distributed database systems and techniques, database software design and implementation, database utilities, database security and related areas.

Contribution to the Bulletin is hereby solicited. News items, letters, technical papers, book reviews, meeting previews, summaries, case studies, etc., should be sent to the Editor. All letters to the Editor will be considered for publication unless accompanied by a request to the contrary. Technical papers are unrefereed.

Opinions expressed in contributions are those of the individual author rather than the official position of the TC on Data Engineering, the IEEE Computer Society, or organizations with which the author may be affiliated.

Membership in the Data Engineering Technical Committee is open to individuals who demonstrate willingness to actively participate in the various activities of the TC. A member of the IEEE Computer Society may join the TC as a full member. A non-member of the Computer Society may join as a participating member, with approval from at least one officer of the TC. Both full members and participating members of the TC are entitled to receive the quarterly bulletin of the TC free of charge, until further notice.

A MESSAGE TO TCDE MEMBERS

As you all know, the Technical Committee on Data Engineering (TCDE) is concerned with the role of data in the design, development, management, and utilization of information systems. The main activities of TCDE in the past have been the publishing of this Technical Bulletin and sponsoring the Data Engineering Conference.

Unfortunately, since the resignation of the last TCDE Chair in June 1991, TCDE has been without an organizational structure. The Computer Society budget for 1992 is to be approved by the Board of Governors by the end of October 1991, and there is no budget request from TCDE. A possible outcome of this situation could be that this issue of the Technical Bulletin may be the last one.

To get us out of this situation, the Computer Computer Society Technical Activities Board (TAB) appointed a reorganization committee consisting of the undersigned. The primary goal of the reorganization committee is to see how TCDE may be reactivated. In particular, we need to formulate an operating plan for technical activities for 1992, develop an interim budget, and work with TAB to get necessary financing. We also need to develop necessary infrastructures (organizational and financial) so that TCDE remains on sound footing in future. We may also use this opportunity to define sharply the role and objectives of TCDE. We will try our best to devise means for TCDE to continue publishing the Bulletin and sponsoring the Data Engineering Conference.

We are writing this to bring you up to date on the current state of affairs and to solicit advise from you on how you would like TCDE to operate in future. Please provide us with input on directional, organizational, and financial issues and what you think should be the role of TCDE.

Rakesh Agrawal, ragrawal@ibm.com
Nick Cercone, nick@cs.sfu.ca
David Choy, dmc@ibm.com

Letter from the Editor

Multimedia is becoming a central theme in computing environments of the future. We have been experiencing tremendous increase in computing power in recent years. This increase in computing power enables us to handle advanced data types, which we call *multimedia*, even in personal computers. Multimedia typically include unformatted data types such as images, video, text, audio, and graphics. We expect in the near future that they will become essential elements in multimedia documents, electronic publishing, facsimile, videotex, videophone, videoconference, and many broadband service (such as ISDN) applications.

This issue of the Data Engineering Bulletin is devoted to helping the readers understand how this new technology affects the research in the database and information systems area. The database research community started looking at these problems only recently. I hope the papers presented in this issue provide an insight into research problems on multimedia potentially related to the database and information systems research.

The papers are listed in the order of presenting the low-level systems issues first, progressing to high-level applications. The paper by Christodoulakis, Ailamaki, Fragonikolakis, Kapetanakis, and Koveos presents an object-oriented architecture for multimedia object servers. In particular, it describes features of Object Virtual Machine (OVM) being developed by the authors and presents the techniques for multimedia handling and memory management in OVM.

The paper by Putz and Neuhold describes a prototype multimedia information system called is-News (Individualized Science News). It addresses many aspects of a future multimedia information system, such as electronic publishing, remote information access, knowledge-based information retrieval, hypertext, and video handling, with a particular emphasis on the integration of various services.

The paper by Little, Ghafoor, Chen, Chang, and Berra surveys interesting issues related to multimedia presentation, in particular, synchronization. Synchronization of multimedia objects from various sources, (e.g., from over the network), has been recognized as an important requirement of a multimedia information system for seamless presentation to the user. It discusses the issues at the physical, service, and human-interface levels of integration. It also addresses the problems of storing and delivering multimedia objects in synchrony to/from the database over the network.

The paper by Wiederhold, Beech, and Minoura reports the findings of a group of database experts on the database technology in Japan. The study group was organized by the Japan Technology Education Center (JTEC) at Loyola College and sponsored by NSF. The paper identifies two key areas of database research in Japan: multimedia and object-oriented databases. It surveys prototypical efforts from both industry and academia in these two areas providing a good overview of on-going database research activities in Japan.

Finally, the paper by Huang introduces an interesting application of the multimedia technology: the multimedia classroom of the future. The paper notes that people's skill is the prime resource in the nation of Singapore, and effective education is of prime importance for the development of people's skill. Recent advances in the multimedia technology help achieve that objective. The paper discusses how multimedia can enhance classroom productivity, presents a prototype multimedia educational system, and suggest research issues.

I would like to thank all the authors for their contributions and cooperation. In particular, I appreciate the contribution from Gio Wiederhold, who spared no time to get the paper out on such a short notice, despite his busy schedule before leaving for DARPA on a two-year research leave from Stanford. I regret that a paper by Shoshana Hardt-Kornacki at BellCore was withdrawn at the last minute due to strict time constraint and the author's extremely busy schedule. I believe it would have been an excellent contribution if included. I appreciate Fragonikolakis' effort for changing the UNIX troff format of his paper to the Postscript format to deliver the revised version on line from Greece in time. Finally, I want to acknowledge the help from Gary Sockut at IBM Santa Teresa Laboratory in making the last-minute revision of the paper by Huang, which was in IBM's SCRIPT format and delivered on line.

Kyu-Young Whang

Associate Professor, Computer Science Department
Director, Database and Knowledge-Engineering Laboratory
Center for Artificial Intelligence Research
Korea Advanced Institute of Science and Technology (KAIST)
email: kywhang@csd.kaist.ac.kr, kywhang@eclipse.stanford.edu

An Object Oriented Architecture For Multimedia Information Systems

Christodoulakis Stavros[1,2]
Ailamaki Natassa[1]
Fragonikolakis Manolis[1]
Kapetanakis Yorgos[1]
Koveos Leonidas[2]

ABSTRACT

We present aspects of the design and implementation of a Multimedia Object Server. We first present the design and implementation of the Object Virtual Machine (OVM), a multimedia object manager. OVM considers environments with large main memories that can contain a large number of objects. However, multimedia objects have to reside in secondary and tertiary storage due to their size. Multimedia objects are treated as first class objects, and methods and access structures for them are supported by the system. The secondary storage information is modeled in three layers: logical files, physical files and devices; thus allowing flexibility, clustering and parallelism in the multimedia data base design. Next, we describe experimental and analytical performance studies for a server environment for multimedia data. We outline results on optimal data placement on optical disks, buffering for delay-sensitive multimedia data environments and scheduling aspects in a server based on secondary and tertiary optical storage (jukeboxes).

1. INTRODUCTION

OVM (Object Virtual Machine) is a general purpose object-oriented DBMS kernel. The OVM kernel directly supports the object-oriented paradigm and provides persistence and sharability for objects. All well known features from object oriented architectures [Kim et al.90], [Deux et al.90], [Fishman et al.87] like complex objects, object identification, encapsulation, typing, multiple inheritance, methods, overriding and extensibility are supported [Christodoulakis et al. 91a]. It also supports an extensive set of database features useful for data intensive applications, including queries, rules, multimedia information management [Christodoulakis et al.86] and versions [Katz 88].

We have already modeled a multimedia information server on top of the OVM kernel. Similar attempts have been carried on in the past [Woelk et al.86]. Meanwhile, we have been prototyping and experimenting multimedia applications on an operational distributed testbed. The results of these experiments are discussed in the sixth section of this paper. Our final objective is to use the knowledge obtained in order to enhance OVM capabilities towards the

1. Technical University of Crete
P.O. Box 133, 731 00 Chania, Greece
Electronic and Computer Engineering Dept.
2. University of Waterloo, Ontario, Canada N2L 3G1
Computer Science Dept.

development of an efficient multimedia information server.

The rest of this paper is organized as follows: Section 2 gives in brief all the fundamental concepts of OVM, including objects, classes, object constructors, several object properties, typing and object relationships. Section 3 outlines the way we support multimedia objects on top of OVM. Sections 4 and 5 refer to the main memory and secondary storage managers of the OVM kernel. Finally, the last section is devoted to analysis and experimental studies on the implementation of a high performance multimedia server on top of the OVM kernel.

2. FUNDAMENTAL CONCEPTS IN OVM

Objects

Objects in OVM correspond to entities in the world. The designer charged to develop an application on top of OVM kernel is expected to set up a natural correspondence between the entities of the application and objects in the OVM kernel model.

Every object has a unique identifier (OID). This identifier (usually referred to, as object id) is used to identify the object through its lifetime in the system. Because objects in the system refer to other objects by their object ids, an object id never changes. However, this object id is for internal use into the system, thus the application layer does not need to know anything about it. References to all objects are done by using these objects' properties or their relationships to other objects. In general, every object may contain object ids to describe its relationships to other objects, and also any non-id data desired.

Objects are grouped into classes. The basic role of classes is to act as "typing" mechanisms for objects: a class description constraints attributes applicable to its instances and their values. Hence, there are class objects in which we maintain information about the attributes such as names and types. From that point of view, class objects provide a natural repository for summary information and meta-data. Class objects can be considered as instances of "metaclasses" and have their own object ids.

Classes are organized into subclass hierarchies. When there are many classes, the resulting hierarchy organizes the class objects and ensures that subclass instances belong to superclasses. Also, the notion of inheritance eliminates repetition of definition, since the application layer need

Insofar, we have considered two basic object categories. There are class objects and instance objects that are generated from the former through the process of instantiation. The OVM kernel also supports objects that do not obey this tight coupling. To be more specific, the notion of "object constructors" is also present. By this, we imply the existence of objects that are generated by constructors and they don't belong to any class. Object constructors maintain a specified internal structure, and objects generated by them inherit this structure along with a set of object properties. There are object constructors for several abstract types like sets, lists and multi-dimensional arrays.

Special attention in the design has been given to first class objects like multimedia objects and rules (for specifying integrity constraints). First class objects are considered objects that have special meaning with respect to the applications that the system wish to effectively support. These objects are like the ordinary ones but because of their importance, the system gives special treatment to them.

The application layer has also the ability to specify several object properties. Depending on whether these properties have been set or not, decisions can be taken and special semantics can be imposed on the operations performed on these objects.

Object Attributes

The smallest unit in which information is organized within an object is the attribute. The names of the attributes along with the domains of the values that an attribute can hold are defined by the application layer at the class definition time. Values are given at instance creation time to the attributes according to their definition and the constraints imposed. The attributes can be either "typed" or "uninterpreted", in a sense that the domain of the values that an attribute can hold is explicitly defined or not.

The primitive types that exist in OVM kernel are: integers, reals, booleans, strings, bitstreams and objects ids. We support both constant and variable length strings. By using constant string types instead of variable length, whenever this is possible, the performance and storage utilization of the system are improved. The primitive type bitstream is mainly used for long data items which are composing multimedia objects.

Any "typed" attribute in OVM can be single-typed or multi-typed. In the former case the type of the attribute can be one of the primitive types provided. In the latter, the attribute type can be either any combination of the primitive types, or a set of specific values of some primitive types, or any combination of the primitive types and multi-types created by the application layer. In addition, typed attributes can be single-valued, multi-valued and tuple-valued with respect to the structure of the data stored. Single-valued are these attributes which hold a single value (as the name implies) belonging to some primitive or multi-type. Multi-valued attributes are allowed to hold more than one value. The structure of these values can be either a set (set-valued) or a list (list-valued) where all properties are preserved and all operations are provided that are known for these abstract data types. On the other hand, tuple-valued attributes are these attributes that hold tuples as values. By tuples we mean a list of pairs (label, value) of predetermined length. Moreover, OVM supports any nesting or combination of the aforementioned structures.

A class definition may contain derived attributes. Derived attributes, like ordinary attributes, have a name and a data type. In addition, each derived attribute has an expression, possibly parameterized, used to calculate the value of the attribute. Such expressions can be rules, functions, programs etc.

In addition to standard attributes we can also have exceptional attributes (defined at instance level), which are not defined in the class objects. Exceptional attributes satisfy the need that instance objects may have a slight deviation from the standard frame given at their class at definition time.

Methods

In class objects definition, in addition to attributes, we also define a number of methods. A method is a piece of code attached to a specific class and can be applied to (instance) objects of this class. Methods are used for the manipulation of objects. We can also have exceptional methods defined at instance level. In order to define a method, the application layer needs to specify the class to which it is attached, the signature (name, type and parameters) and the body (source code) of the method.

Object Relationships

There are several relationships among objects. These relationships form lattices (hierarchies) with objects as nodes.

One type of object relationship forms the PART-OF hierarchy which holds only among instance objects. The PART-OF relationship represents the notion that an object is a part of another object. Links associating instance objects in a PART-OF hierarchy can be either essential or non-essential. Essential links interconnect objects in a PART-OF hierarchy that constitute the main part of the complex object (strongly related). All other PART-OF links are non-essential.

In addition we consider another type of relationship (and therefore another type of link), called reference. It holds between instance objects only, and implies that an object simply refers to another object.

The IS-A hierarchy can be composed of class objects and instance objects. Based on the IS-A hierarchy, we can inherit both values and methods from some objects (superobjects), to their descendents (subobjects) with respect to the hierarchy. A class object can inherit the type definitions and the default values from another class object, an instance object inherits the structure and the default values from a class object (through the process of instantiation), and finally an instance object can inherit methods, attributes (both standard and exceptional) and data from another instance object. We have to emphasize the fact that in any of the above cases an object can inherit information from more than one other objects (multiple inheritance).

3. MULTIMEDIA OBJECTS

We pay special attention to multimedia object management in our design. The first aspect of multimedia objects is that they are long objects. In addition there is a need to support several useful but non-trivial operations on them.

The OVM kernel supports a portion of the class hierarchy dedicated to long multimedia objects. A parent long object class is defined on the top of this portion of the class hierarchy, and all types of long multimedia objects are connected to the former via ISA links. Consequently, we can assume that there is the parent long object class and underneath several subclasses, one for each multimedia object class such as text, voice, image, video and graphics.

Each one of these multimedia subclasses contain a number of attributes. First, there is an attribute that contains descriptive information concerning the value of the specific object. For example, this could be a small but representative portion of the value of a text object, or a small piece of text describing the contents of a picture, or even a small but also representative window of an image. The values of this attribute can be either static or derived. Namely, this small portion of the long data value can be prespecified, precomputed and stored or dynamically computed using a function. In general, this descriptive information is used to inform the application layer for the contents of the long data item that is the actual value of the multimedia object. This way we can avoid fetching the whole long data item from the archive as far as the application layer judge that, this value does not satisfy its needs. (More on this can be found in the subsection of Main Memory Management discussing Long Data Faults.)

In addition to that, a field holding the size of the long data item is also present. Finally, the actual physical location of the value (long data item) is maintained in another attribute. This physical location may denote a location on a magnetic or an optical disk. In this attribute, there also may be a list of addresses, if the long data item has been split up in more than one segments, each located in a separate disk (either magnetic or, more frequently, optical). This can be done for achieving parallelism.

Along with the general long object properties we already described, there are several operations of the multimedia objects that have to be dealt with. All these operations, are supported as methods on these classes. Below we discuss two examples.

The first common aspect present in every multimedia object class, is that of interactive queries. Interactive are these queries, for which we first provide the user with the first part of the long data item as soon as possible, and while the user is busy examining this first part, the system transparently fetches the whole data item from the secondary or tertiary storage. It is obvious that this capability must be present for all multimedia objects because of performance reasons.

In large images operators that specify windows within the image must be allowed. This requires the transfer of a part of the image which is enclosed by a rectangle in main memory. It is obviously a performance oriented feature. A similar requirement exists for digitized sound, where a window of consecutive sound samples has to be transferred to main memory. Similar requirements exist for graphics images.

4. MAIN MEMORY MANAGEMENT

Object Representation in Main Memory

As objects move from secondary storage to main memory and backwards, their representation is likely to change. Both instances and classes hold several attribute values. Values that have a fixed length (e.g. integer-type values) are held inside the object space, while other values (list or set-type values) the size of which is likely to change are stored in a "tail" at the end of object's primitive cluster. When a value in this tail has to be changed, the tail is spread in main memory. When the object is to be stored in secondary storage, the tail is compacted again.

There is another type of object conversion with respect to their object ids as they move from secondary storage to main memory and backwards. Objects often refer to other objects using their object ids. While these references in secondary storage consist of plain object ids, in main memory they are transformed to pointers on demand. This means that when an object's reference to another is activated in main memory, the referenced object id is transformed to a direct main memory pointer to its memory location. Therefore, later accesses to the referenced object is much faster from this point on (there is no need to access the primary indexing methods).

Primary Indexing for Objects

In a similar way to secondary storage object format, we maintain contiguous storage allocation for main memory primitive clusters. Primitive clusters are small groups of objects with strong relationships among them. They are further analyzed in section 5 of this paper. When an object is requested to be fetched from secondary storage to main memory, OVM fetches the entire primitive cluster it belongs. Of course, speaking for class objects, there is no notion of primitive cluster, so we fetch the class object itself as needed. The reasons why we fetch the primitive cluster instead of a single instance object are the following:

- Users that define themselves their primitive clusters do so because they want them to be fetched as units in main memory.
- Even when primitive clusters are created using the default rules provided by OVM, it is more likely that the user needs or will need in the near future the set of an object's strong related relevants in PART-OF hierarchy.

Objects are accessed in main memory using a hash function with an argument which is a part of the object identification that constitutes its primitive cluster identification. This way, we have faster access on strongly related object groups, while we drastically reduce the number of entries in the hash table.

Class Objects In Main Memory

To efficiently represent class IS-A hierarchy in main memory we use the technique of constructing bitvector codes, introduced in [Ait-Kaci 89]. As classes are being created and deleted, a bitvector table in main memory containing a bit code for each class grows and shrinks. The table represents the class IS-A hierarchy and is very efficient in performing lattice operations.

Long Data Faults

We assumed that it is profitable to keep all related data in main memory at all times. This is not true for long objects such as for multimedia because a single multimedia object (as for example an image) may occupy a very large proportion of main memory (several megabytes) and therefore cause extensive swapping in main memory. We will only bring multimedia objects in main memory selectively based on the user's explicit demand. In that respect, requests for multimedia objects to secondary storage are viewed as intentions of the application layer. Thus, instead of fetching all the data available in the object, we fetch all but the actual value (that is the long data item). In addition to other information, such as the size of the long data item and the physical location of it in secondary storage, a description of the value is also fetched. This description can be a piece of text describing the nature of the information contained in the object (e.g. describing the picture contained, or whose voice has been recorded). Alternatively, this can be a small but very descriptive part of the data (e.g. a small window of the image or a few seconds of voice). This small piece of information serves as providing a user with a general idea of what the long object is about, preventing the user demanding the whole object just to see what it is.

In addition to that, we also support partial fetching of multimedia objects. For example, one can fetch a specified window of an image. This provides our system with an additional feature, which offers great flexibility. As a matter of fact, we reduce space overhead in main memory by fetching only the part of the long object we actually need. However, another characteristic of the multimedia objects is that some of them are needed at once or not at all. Thus, the strategy that has to be used is "use it and toss it".

5. SECONDARY STORAGE MANAGEMENT

All objects in OVM enjoy the property of persistency. The Secondary Storage Manager groups objects in logical and physical units according to the needs of the application layer. Below, we describe the basic units of information clustering as well as the way in which the Secondary Storage Manager maps these units to physical files and devices.

Logical Files

The application layer organizes and accesses information, on the top of the OVM kernel, based on the notion of logical files. A powerful set of operations and a script language are the tools used to support this notion.

A logical file, which can be considered as a form of information clustering, may, for example, contain all information related to project X, or all recent communication, or even all objects of type A. Logical files can be organized hierarchically. This is in accordance to ISO standards for office system architectures. A logical file may contain objects that are frequently accessed together. For example, someone may want to work with a given project X at a certain point in time so he invokes the logical file X in main memory.

Access methods can be associated with a given logical file. The user can navigate through the objects of the logical file, using the access methods provided for the logical file, and ignoring any other data or access methods in the database. This is different than what is provided in traditional databases. There, an index is assigned to a particular relation (class). Here, we can assign an index across diverse classes that exist in the same logical file.

Clustering

Every PART-OF hierarchy is logically divide into strongly connected parts, that usually need to be fetched together in main memory.

Thus for achieving better performance and also for modeling reasons, we introduce clusters. Clusters are these logical but also physical data units, consisting of a number of relevant objects with respect to their relationships in a PART-OF hierarchy. There are several ways for defining clusters in OVM. The application layer may explicitly specify the clustering policy. If no indication is given, the system generates a default clustering policy which is mainly dependent on the notion of essential and non-essential links.

Mapping Logical Files to Physical Files

A logical file may be mapped into more than one physical files for performance reasons. The mapping, partitions the set of clusters of the logical file into subsets residing on physical files of the given logical file. The rules of the mapping are specified with a script language. For example the contents of a single logical file may be partitioned into two physical files each located in a different devise in order to achieve parallelism in the data transfer to main memory. This is important when a single device does not have the bandwidth to transfer data in main memory, in the rate that this data is needed (for example video data have very large bandwidth requirements).

Mapping Physical Files to Devices

A physical file may be placed on one or more devices (possibly removable). The part of a physical file that resides on a single device is called an extend.

This mapping operates on a three level store basis: main memory, magnetic disk and tertiary storage (i.e. optical disks). The distribution of the physical units among the three levels depends on the semantic information that is kept on every physical unit and the distribution criteria that the application layer defines. The motivation for such a storage organization is based on the fact that future object managers will be called on to manage very large object bases in which time critical objects reside permanently in main memory, other objects are disk resident, and the remainder occupy tertiary storage. Migration mechanisms (synchronous or asynchronous) are investigated for transferring the physical units to the appropriate level.

Finally, the ability of placing physical units to different devices provides parallelism, achieving this way much better performance. These diverse devices may be classified in different levels of the memory hierarchy, with respect to access speed and storing capacity. This ability is very important, mostly for the storage management of large multimedia objects.

6. ANALYSIS AND EXPERIMENTAL STUDIES

Our long term objective is the implementation of a high performance multimedia server in a distributed workstation environment. OVM (which is already operational) will be the data model used in this environment. In parallel, we are pursuing prototypical, experimental, and analytical studies for the design and implementation of high performance features for multimedia data management in the distributed workstation environment. A second version of a distributed testbed for multimedia applications operates in this environment and has been used extensively for prototyping, experimenting, and understanding the functionality and performance requirements of this environment. Solutions that prove successful will be integrated in the new system. In the following we present a summary of our theoretical and experimental results so far.

Secondary and Tertiary Storage Devices

Multimedia data requires very large storage capacities. Optical disks of various kinds are appropriate media for the archival of multimedia data. They provide inexpensive storage, large storage capacities, and reasonably fast random access. There are currently analogue and digital optical disk storage devices in the market. Digital optical storage of multimedia data presents advantages in the long run. Within the category of digital optical storage devices one can find devices better suited to publishing (CDROM's), to archival of information (WORM's), or to changing information (rewritable) as well as variations in between. With present technology often it is not possible for any single device to store the large volume of multimedia data produced and processed by organizations daily. For this reason, optical disks are organized in larger arrangements (jukeboxes) from where disks are fetched on demand via a (slow) control mechanism [Byte 90].

Optical disks have different performance characteristics than magnetic disks. Therefore their retrieval performance must be understood clearly and modeled to allow for maximum utilization of their capabilities [Christodoulakis 87]. Based on the performance models developed, better solutions to data base performance problems may be provided.

Such a problem is that of optimal data placement on CLV optical disks. CLV (Constant Linear Velocity) optical disks, unlike CAV (Constant Angular Velocity) optical disks, have non-uniform storage capacity distribution across a radius of the disk. In addition the angular velocity with which the disk rotates is not constant, but it depends on the location of the disk head. When the disk head is near the outside tracks, the disk is rotated slower than when the disk head is near the inside tracks. In this environment the problem of optimaly placing the data on top of the CVL disk is very different than the problem of optimal data placement on CAV disks (most magnetic disks for example). The distribution of sectors to tracks in an optimal allocation satisfies the following two properties: First, in an optimal arrangement there cannot exist two different positions on the disk, x and y, such that both a $P1$ and a $P2$ mass elements are assigned to position y (consecutivity of the allocation probability masses). Second, the optimal arrangement of the probability masses is unimodal. (For the details, see [Ford and Christodoulakis 91]).

We have derived optimal and heuristic solutions to the problem of optimal placement, and we have shown that it is sensitive to the ratio of seek time and rotational delay, to the rate of increase of the disk storage capacity, and to the skewness of the access probability distribution. As the rotational delay increases (with respect to the seek time), the location of the mode of the probability distribution of tracks moves towards the inside track of the disk. As the skewness of the storage capacity increases, the location of the mode will be moved towards the inside tracks or the outside tracks, depending on the ratio of seek time to rotational delay. Finally, as the skewness of the access probability increases, the mode moves towards the outside tracks [Ford and Christodoulakis 91].

Delay-Sensitive Multimedia Data Retrieval and Buffering Requirements

Multimedia data often poses real-time transfer and synchronization requirements. The real-time transfer requirements are soft (delay-sensitive). It is desirable for the data to reach the receiver end (within the workstation itself, or across the network) always in time so that there are no interruptions in the presentation. For example if voice or video is retrieved from a secondary storage device, the next chunk to be presented should reach the receiver end before the currently played chunk finishes, else interruptions in the presentation will be heard or seen (for voice presentation and video presentation respectively).

We have investigated the impact of the retrieval requirements for delay-sensitive data on the buffering requirements of the system, as well as on the database design and ability of the system to multiplex users [Gemmel and Christodoulakis 91]. A formal model was developed and analytic formulae derived that give solutions and answers to the above problems. The model takes into account data compression with variable compression rates. Conditions were derived to guarantee that no interruptions will occur in the playing sequence. Experimentation verified our results. These results also apply to multiple output channels.

Scheduling in the server

Scheduling may take place at several places in a computer system. Scheduling in the server aims to the maximization of the system performance. There are very important reasons for scheduling in an environment of multimedia data. First, because data may be delay-sensitive and therefore on-time retrieval must be guaranteed. Second, because multimedia data is long and therefore queuing delays may be frequent. Third, because retrieved data may not be absorbed fast by interactive users and therefore occupy large system resources. Fourth, because tertiary storage devices (jukeboxes of optimal disks) may be involved, and therefore the retrieval may be very slow.

The problem of optimal scheduling is difficult even in relatively simple cases. Consider for example a simple type of jukebox with a simple read write head. In this environment optimal scheduling algorithms that minimize the system time or the response times are possible under certain conditions [Christodoulakis et al. 91b]. The difficulty of the problem however increases rapidly, and heuristic solutions are required even in this simple case.

In our studies we have developed and tested in a prototype environment of a distributed multimedia information system the performance of intermediate and long-term scheduling algorithms. The scheduling algorithms calculate priorities for the disks to be fetched next, based on several factors like the priorities and the number of jobs that access information from the disk, the time that a job has already in the system, the amount of information that has been delivered to the job so far, the cashed information for the job, etc. Experimental results show also strong dependence of the response times on the size of the slice retrieved per job, as well as on the type of the data to be retrieved [Sam Chee 91]. It has also been shown that the ability to interrupt the execution of queries is important in this environment, where users may not be interested in more information after they find what they want [Mok 91]. Protocols for information staging from the server to distributed cash, to the client stations that improve response times and system throughput have also been developed.

7. SUMMARY

We have described aspects from the design of a multimedia object server. We described the design and implementation of OVM, a general purpose object management system, which is part of a multimedia server under implementation. We also described results of performance studies related to the design and implementation of such a server. Several issues related to performance remain to be investigated (like parallelism, data placement, data migration, server architecture, cache management etc).

We pursue a strongly prototypical and experimental approach since the interelationships of the design parameters make the analytical solutions hard and unreliable. The results of the performance studies pursued will show what are the important parameters that must be integrated in the final system implementation. We implement our system following an object oriented framework. In parallel, we investigate and implement particular multimedia information system applications (like geographic information systems, document management systems etc) to verify our design and find new requirements.

8. REFERENCES

[ACM 91]
"State of the Art report on Standards for Multimedia Systems", (E. Fox editor), CACM, April 1991.

[Ait-Kaci 89]
"Efficient Implementation of Lattice Operations", ACM Trans. on Programming Languages and Systems, January 1989, 115-146.

[Byte 90]
"State of the Art Report on Magnetic versus Optical Disk Technologies", BYTE, November 1990.

[Christodoulakis et al. 86]
S. Christodoulakis, F. Ho, M. Theodoridou, M. Papa, A. Pathria: "Multimedia Document Presentation, Information Extraction and Document Formation in MINOS: A Model and a System", ACM TOOIS, October 86, pp.345-383.

[Christodoulakis 87]
S. Christodoulakis: "Analysis of Retrieval Performance for Records and Objects Using Optical Disk Technology", ACM TODS, 1987, pp.137-169.

[Christodoulakis et al. 91a]
S. Christodoulakis, L. Koveos, M. Fragonikolakis, N. Ailamaki, G. Kapetanakis: "Revised Design Documentation of the OVM", Technical Report TUC/DECE 91-008, Department of Electronics and Computer Enginnering, Technical University of Crete.

[Christodoulakis et al. 91b]
S. Christodoulakis, D. Anestopoulos, S. Argyropoulos: "Scheduling Algorithms in Jukebox Architectures", Technical Report TUC/DECE 91-009, Department of Electronics and Computer Enginnering, Technical University of Crete.

[Deux et al.]
Deux O. et al., "The Story of O2", IEEE Trans. on Knowledge and Data Engineering, April 1990, 91-109.

[Fishman et al.87]
"IRIS: An Object-Oriented Database Management System", in Readings In Object-Oriented Database Systems, edited by S. Zdonik and D. Mayer, pp. 216-226.

[Ford and Christodoulakis 90]
D. Ford and S. Christodoulakis: "Optimal Data Placement on CLV Optical Disks", ACM TOIS, September 1990.

[Gemmell and Christodoulakis 91]
J. Gemmell and S. Christodoulakis: "Principles of Storage and Retrieval for Delay Sensitive Data", ACM TOIS, 1991

(to appear).

[Katz 88]
Randy H. Katz, "Towards a Unified Framework for Version Modeling", Report No. UCB/CSD 88/484, December 1988.

[Kim et al. 90]
W. Kim, F. Garza, N. Ballou, D. Woelk: "Architecture of the ORION Next- Generation Database System", IEEE Transactions on Knowledge and Data Engineering, April 1990, pp. 109-124.

[Mok 91]
Steven C.W. Mok, "Protocols for Data Archival and Retrieval in the MINOS Distributed Testbed", MMATH Thesis, Department of Computer Science, University of Waterloo, 1991.

[Sam Chee 91]
Michael A.L. Sam Chee, "Scheduling in the Server of a Distributed Multimedia Information System", MMATH Thesis, Department of Computer Science, University of Waterloo, 1991.

[Woelk et al.86]
D. Woelk, W. Kim, W. Luther: "An Object-Oriented Approach to Multimedia Databases", Proceedings of the ACM SIGMOD Washington D.C. May 1986, pp. 311-325.

is-News: a Multimedia Information System

*Wolfgang Putz, Erich J. Neuhold**

Abstract

Many research groups are currently working on prototypes, which try to utilize the emerging technologies in the communications and personal computing area. Common to most of these research activities is the attempt to build multimedia aspects into a single system. Within both the systems themselves as well as the user interface many different metaphors and modelling paradigms are explored. Especially for the areas of electronic publishing and teaching/training multimedia aspects are to be supported homogeneously in a distributed electronic communication environment.

is–News – Individualized Science News – has been developed by the Integrated Publication and Information Systems Institute (IPSI) of the GMD, the German Research Centre of Computer Science. It is a multimedia application in the electronic publishing domain and emphasizes multimedia and integration aspects. Access to heterogeneous multimedia databases and remote online database as well as a hypertext writing environment and an editor's workbench were integrated in a prototypical system that will explore different human computer interface metaphors. is–News is used as an experimental environment of ongoing research activities and will be enhanced continuously by our own efforts or those of our partners. We describe the first prototype of is–News emphasizing its multimedia aspects followed by a short outline of future work.

Introduction

The potential of multimedia systems increases with the availability of high-performance workstations, high-resolution displays and the emerging communication technology. Parallel to these hardware enhancements, the emergence of object-oriented database systems, the implementation of image processing and animation software and – as a combination of hardware and software developments – the development of compression and decompression algorithms for images encourage the construction of multimedia applications. We consider multimedia systems as systems that handle time-dependent information like video and audio and time-independent information like text, graphics and images and that provide computer-based manipulation, management, presentation and communication of information. Many international research institutes are working on prototypes of multimedia systems emphasizing different special topics. For example, the Electronic Magazine of Bellcore Laboratories concentrates on the support of public switched networks and possible future network based services with emphasis on filtering techniques to show possible ways of dealing with an ever-increasing information flow [9]; the Newspace application of the MIT Media Lab is a system that offers a broadsheet sized electronic news presentation using "paper quality" displays [1]; VANGOGH, a system for heterogeneous multimedia messaging, developed at CWI Amsterdam, is concentrating on document composition, transfer and delivery of text, image and voice data [3].

*GMD–IPSI Integrated Publication and Information Systems Institute
Dolivostr. 15 D–6100 Darmstadt, FRG
e–mail: {putz,neuhold}@darmstadt.gmd.de

This paper describes a prototypical system and its multimedia aspects, with particular emphasis on the integration of the various steps of designing, producing and using an electronic newspaper. Many services like access to heterogeneous local multimedia databases and remote online databases, a hypertext writing environment and an editor's workbench have been integrated in the system.

The is–News prototype

The central idea of is–News (Individualized Science News) is to develop an experimental environment, which demonstrates our common vision of a future multimedia information and publication system. New techniques in the hardware and the software area and new approaches in multimedia information storage, retrieval, processing and presentation are combined in a comprehensive experiment, which comprises all research activities of our institute and which includes results of our research partners.

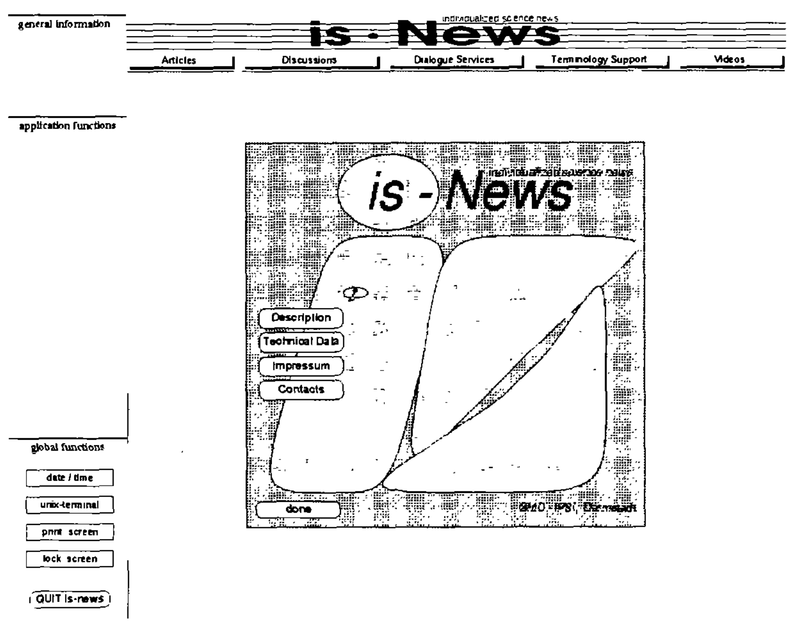

Figure 1: is–News logo

The very first prototype of is–News (see Figure 1) was completed at the end of 1990. The emphasis was on the reader's point of view. The content we selected were articles featuring the electronic publishing area, reports on recent events, news, background information like dictionaries and information of interest to the members of our institute. Important points were that the very first version would be a multimedia system and that it would integrate new services, not available in existing newspapers. The resulting system would allow for access to an article pool, to an international bulletin board and to remote databases. The system supports nonlinear documents, a knowledge-based information system and provides access to an electronic publishing dictionary. In addition to text and image information we support access to video sequences on optical discs.

It should be mentioned that two different aspects of multimedia systems are considered in is–News. The first is the aspect of different information types. As mentioned, is–News supports time-independent information of text and image and time-dependent information of video and is therefore a true multimedia system. The other aspect is multimedia support for the user interface. The visualization of information on the screen and the functionality available to the user for handling the system include many possibilities to integrate different metaphors like hypermedia structures or newspaper layout.

ing an article, the topics and "keywords" found can be used as a starting point to search in the databases for more detailed publications or for information on future conferences.

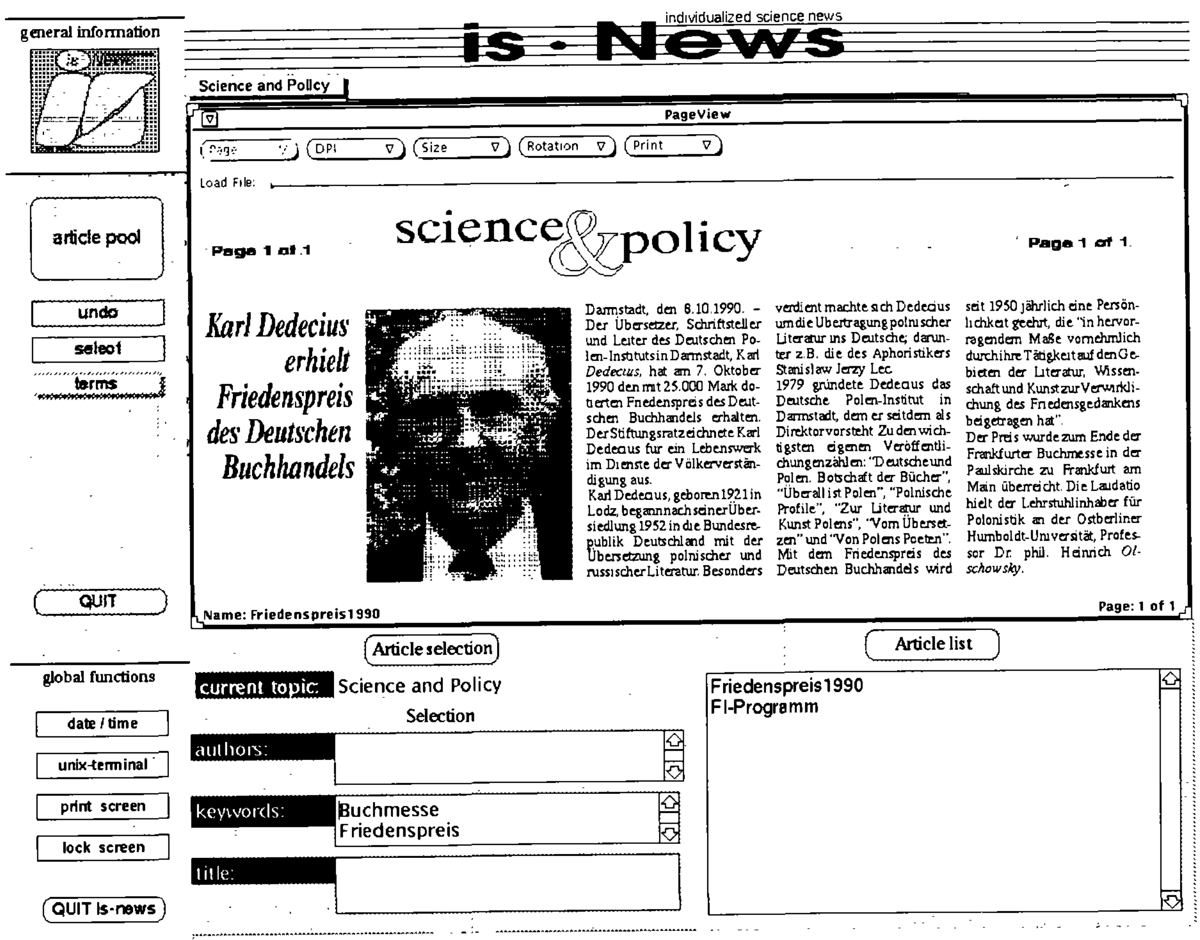

Figure 3: Screen layout of an is–News newspaper

The crucial point using external information sources is that the information should uniformly be available to the user independent of the actual information source. Both types of implemented services deliver only unstructured text information. Articles of the USENET bulletin board and the search results of the remote databases are pure ASCII information without marked structures. Within the electronic newspaper however the information must be handled like all other information in a structured form. It must be presented in the user defined context and must be available for user-specific operations like copying or inclusion in a personal database. The solution implemented in is–News is to transform the selected information to SGML-structured documents. This transformation is done automatically. Every selected document is parsed and the content structure is recognized [6]. The marked structured document can then be considered as an object and managed using the object-oriented database system.

At the moment, only time-independent information is available from public information services. In the first version of is–News we were interested in demonstrating that access to external information services is a valid function for electronic publishing products. With the enhanced communication technology, however, remote dynamic information like videos, animation, sound and voice will be available and will be treated in the same manner as we currently handle local information of these types.

Knowledge-based information system

Very often information in a newspaper is the starting point for decisions – just think of the stock market information found in your daily newspaper. The electronic version of the newspaper can be used as an environment to analyze problems and to find solutions. In addition to the service on conference information is–News is therefore used as a framework to support decision making – the reader should be able to prepare his attendance of these conferences.

The hypermedia service SIC ("System for Information about Conferences") utilizes a knowledge-base containing general and specific information about conferences. The user chooses a "perspective" offering different aspects of conferences like topics of the conference, workshops, tutorials, persons like referees, attendants, or lecturers. The knowledge pertaining to the selected perspective is presented by graphical objects, showing the information at various levels of detail [15]. As the user can navigate in the generated object space only according to principles of topical coherency, a dialogue is achieved which realizes to some extent a graphical conversation with the conference information-base.

The user interface to the conference information service combines graphical query constructs with exploration oriented graphical components. Working with structured information-bases and with object-oriented databases requires new approaches to user interfaces. Due to the intensive interrelationships between the information objects, the stored structures often show such a high degree of complexity that traditional query languages are not flexible enough to handle these structures. Especially, imprecise and incorrectly phrased requests have to be accepted and have to be rectified during query processing.

The knowledge necessary to solve these problems is represented in a concept-oriented form. The concept classes workshops, tutorials, conference topics, persons and institutions are covered. Selecting one of these concepts determines the perspective, from which the subtopic of the information objects can be derived.

Using a form sheet, the user can restrict the chosen perspective to reduce the amount of relevant objects by stating attribute restrictions. The system checks whether the attribute values are valid and, if they are not, the knowledge-base is used to present a list of semantically similar, acceptable values. The results of these restrictions are those objects of the information-base which comply with the user's demand. Since these objects are linked to other objects in the information- and knowledge-bases, we can progress to other information as controlled by the user's dialogue.

The results are transformed into a hypermedium format, i.e. the visualizations are constructed as cards within hypermedia stacks (see Figure 4). The user can navigate through these stacks by following a thematic pattern which is determined by a transition network called a state graph [13]. This graph describes the potential sequence of dialogue states the user follows when gathering information about his subject of interest.

Terminology support

When reading textual information in an electronic newspaper, two problems can be encountered. The reader finds a term, for which he would like to get a definition or an explanation. He may also be confronted with a foreign language document. In both cases the electronic product will give direct support. In the first case the user needs a dictionary to locate the term in question and in the second case a machine translation system will help.

In the first version of is–News a dictionary on electronic publishing is supported. It can be used as a separate service, but in most cases other services are the starting point for using the dictionary. Every entry of the dictionary contains a term, its German or English translation, the definition of the term and further information like synonyms, homonyms, narrower or broader terms, and so on. Like all other documents handled in the system every entry of the dictionary and the dictionary itself are structured documents.

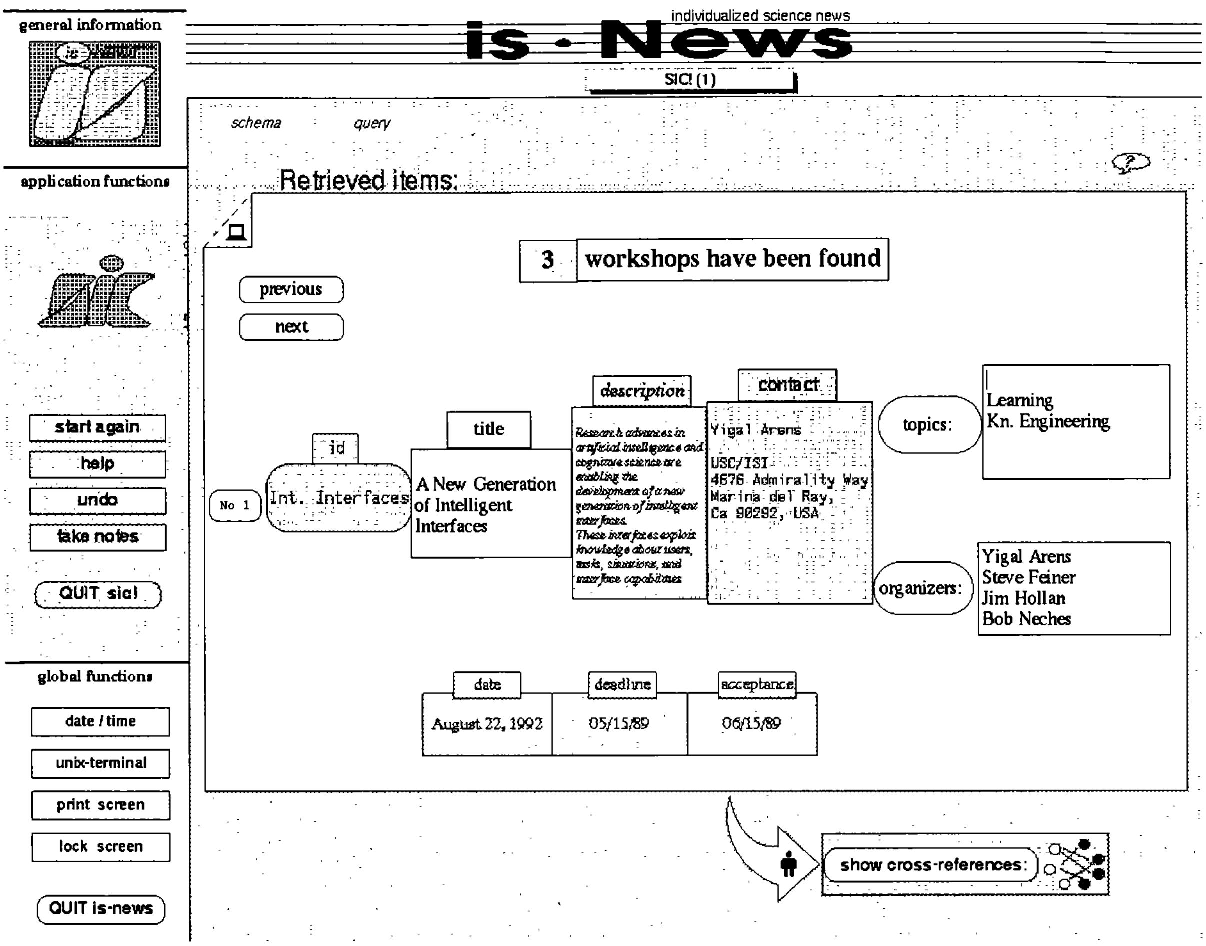

Figure 4: Knowledge-based information system

At the moment only the reader's point of view of the dictionary is supported. But different readers can have different views of the dictionary depending on their information needs. Furthermore authors and editors have yet another view even when only "reading" the dictionary. For example, an editor may choose an appropriate synonym for a term depending on the potential readers. A document written in colloquial language normally allows a broader interpretation of a term than will be expected in a scientific document. Therefore, the view of the dictionary depends on the actual user and his role in the publication environment.

Thus, the electronic version of a dictionary opens up handling possibilities at two different levels. Firstly, the information is not restricted to text and image information but can include dynamic information and, secondly, multiple views of the same information can be supported. Both are technically founded on the structured document approach and the object-oriented database.

The translation service uses an existing machine translation program which allows per sentence translation of textual information. Our goal is to test how such a service even with the quality limitations of current machine translation programs can be integrated in a newspaper environment.

Hypertext

A multimedia system cannot be thought of containing linear documents only. The concept of hypertext, or more exactly hyperdocument, offers new possibilities for writing and using all types of information [4,7]. A hyperdocument consists of a network of nodes, each node containing a part of the document produced by the author. The author's/ editor's additional task is it to define semantic links between the nodes, thus opening new browsing and navigating possibilities. Nodes can contain all types of information, text, image, video and audio, even executable programs.

Although the reading and browsing point of view is emphasized in the other services of is–News, the hyperdocument service is oriented toward concepts for writing argumentative texts. Argumentative texts emphasize the reasoning structure of information and therefore need a specific model that supports the comprehension of arguments and their relationship. The approach used is based on the schemata proposed by Toulmin [16] and Kopperschmidt [11]. Whereas the Toulmin schema provides an analysis of arguments at the micro level, the Kopperschmidt schema, expanded by our own abstraction hierarchy is used at the macro level. The presentation and browsing styles for argumentative texts were produced with our authoring tool and are included in the first version of is–News. We used an existing argumentative text, the "chinese room" dialogue on artificial intelligence. Different levels of argumentation and the micro structure of the argumentation are presented graphically. Starting from an introductory text, the reader can browse the graphical or textual representation at different levels of abstraction. When switching the abstraction level, the graphical presentation will be adapted to the new environment. With this approach we believe that we can avoid the problem of being "lost in hyperspace." The graphical presentation is used to follow links between different nodes. In the textual presentation links are represented by highlighted terms. There is a special node type "executable program", where a separate process can be started.

For this kind of hyperdocument powerful authoring tools are needed to support this extensive structuring process. Otherwise, authors will reject creating such documents or will produce badly structured ones. One of the central ideas in designing our authoring system is to externalize as many internal (mental) states of the writing process as possible. This reduces the intellectual work on the part of the author and supports activities, for which he normally needs notes and scribbles. From the various cognitive aspects of the writing task, different "activity spaces" are provided [14] in a publication environment component called SEPIA (Structured Elicitation and Processing of Ideas for Authoring).

Video

The only time-dependent information type currently supported by is-News is video. Video sequences are stored as analogue coded information on tapes or optical discs and can be distributed locally by an analogue video/audio network. The workstation converts the information in real time to digital information to allow viewing and manipulating a video sequence or individual frames.

The main problem in integrating this type of information is the synchronisation of analogue and digital information. To allow retrieval of and access to the analogue information, the video sequences are described in a database. The approach chosen is to store each video sequence as an object. Attributes describe the author, the abstract and the classification of the content. Additional information about the storage on tape or optical disc is needed. For example, the duration in seconds or the starting and final frame number on the optical disc or the position on the tape. This information is dependent on the storage medium used. Operations on a video sequence like starting, stopping, advancing frame by frame are then methods of this object type and are transparent to the user.

For an interface metaphor we are exploring the VCR controls displayed as icons on the screen. But, in addition, the user will also be able to capture his own copies of individual frames and include them into his personal information space.

Future work

The version of is–News described is a very first prototype of a multimedia information system. As has been pointed out before the present functionality is mainly restricted to the reader's point of view. The emphasis was on the integration of new media and new services in a single electronic product. Future work continues to follows a general model of the publication process, which can be described as preparation, production and communication of documents and which involves authors, editors, designers, reprographers and users. These groups must be provided with homogeneous access to their tools to develop a publication environment for all steps in the publishing process.

Scientifically our future work will concentrate on the completion of the environment and its enhancement with more "intelligent" components. It will be used as a testbed to implement new and more automated services and to test their acceptance. One such new service within is-News will be a multimedia messaging service. With the new communication technology, information distributed over the network is not restricted to text and image. Prototypical systems which include audio are in existence. Video sequences will be integrated with improved network capacities. As a consequence is-News will not be restricted to an inhouse newspaper. We plan to make it available to external users of the local science community and, at a later stage, to commercially interested people.

Another direction regarding the functional enhancement of is-News will be the author's and editor's workbench. The authoring tool was already mentioned. It will be completed to become a cooperative authoring tool including video/audio support. Editorial support of integrated image, video and audio information will be realized step by step. Using existing tools like image processing systems and digital video editing systems our important point is that the electronic multimedia newspaper has to provide new functionality in order to be commercially successful. A few of them have been discussed here, others will have to be explored in the future.

One such feature will be an automatic text-generation component for visualization of information. Answering a user's request the selected fragments of the information-base and their semantic relations will be transformed into automatically generated sentences. Another area of interest is to introduce newer multimodal capabilities like touch screens, tablets, and audio.

The main goal when selecting a system architecture is to get a high degree of flexibility for adding new components. At the moment we use a server-client architecture for the system. The different service clients use the user interface server, the data management server etc.. Another approach would be the assistant-agent approach [12], where assistants directly support users and agents provide services to these assistants. We will test these approaches by developing alternate architectures for the multimedia applications in our is–News environment.

Acknowledgement

The experimental environment is-News is a joint venture of all research departments of the Institute. Special thanks are given to all members of the departments who contributed actively to the system. There are too many of them to be mentioned by name, but a few are identified in the list of references.

References

[1] W. Bender, H. Lie, J. Orwant, L. Teodosio, N. Abramson. Newspace: Mass Media and Personal Computing. *to appear in USENIX*, 1991.

[2] E.A. Bier, A. Goodisman.Documents as User Interfaces. In R. Furuta (Ed.) *EP 90 Proceedings of the International Conference on Electronic Publishing, Document Manipulation & Typographie,* pp. 249–262, The Cambridge Series on Electronic Publishing, Cambridge, 1990.

[3] D.C.A. Bulterman. VANGOGH: A System for Heterogeneous Multi-Media Messaging. *Report of Centre for Mathematics and Computer Science CWI Amsterdam,* Amsterdam, 1990.

[4] J. Conklin. Hypertext: An Introduction and Survey. *IEEE Computer Magazine*,20(9):17–41,1987.
[5] W.D. Dodd. Convergent Publication, or the Hybrid Journal: Paper plus Telecommunications. *Electronic Publishing*,3(1):47–59,1990.
[6] T. Goettke. DREAM V1.0 – User Manual. *to appear as Working Paper of the GMD (in german)*, GMD, 1991.
[7] F.G. Halasz. Reflections on Notecards: Seven Issues for the Next Generation of Hypertext Systems. *Communications of the ACM*,31(7):836–852,1988.
[8] ISO. Information Processing – Text and Office Systems – Standard Generalized Markup Language (SGML) (ISO 8879). *International Standards Organization*,1986.
[9] T.H. Judd,G.C. Cruz. Customized Electronic Magazines: Electronic Publishing for Information Grazing.*Advanced Printing of Paper Summaries, Electronic Imaging '89*,pp. 504–509, 1989.
[10] W. Klas, E.J. Neuhold. Designing Intelligent Hypertext Systems using an Object-Oriented Database Model. *Working Paper of the GMD No. 489*, GMD, 1990.
[11] J. Kopperschmidt. An Analysis of Argumentation. In T.A. van Dijk (Ed.),*Handbook of Discourse analysis: Vol 2. Dimensions of Discourse*, pp. 159–168,London : Academic Press, 1985.
[12] P. Muth, T.C. Rakow, W. Klas, E.J. Neuhold. A Transaction Model for an Open Publication Environment. *IEEE Data Engineering, Special Issue on Unconventional Transaction Models for Advanced Applications*, 14(1):34–38, 1991.
[13] S. Sitter,A. Stein. Modeling the Illocutionary Aspects of Information-seeking Dialogues. *to appear in Information Processing & Management*, 1991.
[14] N.A. Streitz, J. Hannemann, M. Thüring. From Ideas and Arguments to Hyperdocuments: Travelling through Activity Spaces. In *Proceedings of HYPERTEXT '89* , pp. 343–364, Pittsburg, 1989.
[15] M. Travers. A Visual Representation for Knowledge Structures. In *Proceedings of HYPERTEXT '89* , pp. 343–364, Pittsburgh,1989.
[16] S. Toulmin. *The Uses of Argument*. Cambridge: Cambridge University Press, 1958.

Multimedia Synchronization

T.D.C. Little, A. Ghafoor[1], C.Y.R. Chen, C.S. Chang, and P.B. Berra

Department of Electrical and Computer Engineering
Syracuse University
Syracuse, New York 13244 USA

[1]School of Electrical Engineering
Purdue University
West Lafayette, Indiana 47907 USA

Abstract - Multimedia synchronization has been recognized by many researchers as a significant requirement for applications using time-dependent media. The orchestration of static data elements such as images and text, and the "lip-sync" of audio and video are examples of such synchronization. In this paper, many aspects of multimedia synchronization are reviewed at the physical, service, and human interface levels of integration. Applicable areas include temporal modeling with intervals and abstractions, conceptual and physical models for databases, and systems support for synchronization including real-time scheduling and communications protocols.

1 Introduction

Multimedia refers to the integration of text, images, audio, and video in a variety of application environments. These data can be heavily time-dependent, such as audio and video in a movie, and can require time-ordered presentation during use. The task of coordinating such sequences is called multimedia synchronization. Synchronization can be applied to the playout of concurrent or sequential streams of data, and also to the external events generated by a human user. Fig. 1 shows an example of time-ordered multimedia data for a sequence of text and images.

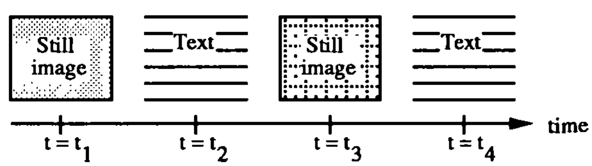

Fig. 1. Time-Dependent Presentation

Temporal relationships between the media may be implied, as in the simultaneous acquisition of voice and video, or may be explicitly formulated, as in the case of a multimedia document which possesses voice annotated text. In either situation, the characteristics of each medium, and the relationships among them must be established in order to provide synchronization in the presence vastly different presentation requirements. Consider a multimedia slide presentation in which a series of verbal annotations coincides with a series of images. The presentation of the annotations and the slides is sequential. Points of synchronization correspond to the change of an image and the end of a verbal annotation, representing a coarse-grain synchronization between objects. A multimedia system must preserve the timing relationships among the elements of the object presentation at these points of synchronization by the process of temporal integration.

In addition to simple linear playout of time-dependent data sequences, other modes of data presentation are also viable, and should be supported by a multimedia database management system (MDBMS). These include *reverse, fast-forward, fast-backward,* and *random access.* Although these operations are quite ordinary in existing technologies (e.g., VCRs), when non-sequential storage, data compression, data distribution, and random communication delays are introduced, the provision of these capabilities can be very difficult.

In this paper, we describe the multimedia synchronization problem for time-dependent MDBMSs with respect to three levels. These are the *physical level,* the *service level,* and the *human interface level* [1]. At the physical level, data from different media are multiplexed over single physical connections or are arranged in physical storage. The service level is concerned with the interactions between the multimedia application and the various media, and among the elements of the application. This level deals primarily with intermedia synchronization necessary for *presentation* or *playout*. The human interface level describes the random user interaction to a multimedia information system such as viewing a succession of database items, also called *browsing*. We also overview important temporal models necessary to describe time-dependent media, and survey various approaches for their specification. Furthermore, we describe the implications of time-dependent data retrieval when data can be distributed.

The remainder of this paper is organized as follows. In Section 2, we define appropriate terminology for synchronization of time-dependent data. Section 3 surveys conceptual models for describing multimedia synchronization including temporal intervals, process models, user interaction, and temporal abstractions. Section 4 provides a examination of the database issues related to multimedia synchronization, including logical and physical models of temporal data, and systems support. Section 5 concludes the paper.

2 Classification of Time-Dependent Data

Time-dependent data are unique in that both their values and times of delivery are important. The time dependency of multimedia data is difficult to characterize since data can be both static and time-dependent as required by the application. For example, a set of medical cross-sectional images can represent a three-dimensional mapping of a body part, yet the spatial coordinates can be mapped to a time axis to provide an animation allowing the images to be described with or without time dependencies. We must therefore develop a characterization of multimedia data based on the time dependency both at data capture, and at the time of presentation.

Time dependencies present at the time of data capture are called *natural* or *implied*. Audio and video recorded simultaneously have natural time dependencies. Data can also be captured as a sequence of units which possesses a natural ordering (e.g., see the medical example above). On the other hand, data can be captured with no specific ordering e.g., a set of photographs. Without a time dependency, these data are called *static*. At the time of playout, data can retain their natural temporal dependencies, or can be coerced into *synthetic* temporal relationships. A synthetic relation possesses a time-dependency fabricated as necessary for the application. For example, a motion picture consists of a sequence of recorded scenes, recorded naturally, but arranged synthetically. Similarly, an animation is a synthetic ordering of static data items. A *live* data source is one that occurs dynamically and in real-time, as contrasted with a *stored-data* source. Since no reordering, or look-ahead to future values is possible for live sources, synthetic relations are only valid for stored-data.

Data objects can also be classified in terms of their presentation and application lifetimes. A *persistent* object is one that can exist for the duration of the application. A *non-persistent* object is created dynamically and discarded when obsolete. For presentation, a *transient* object is defined as an object that is presented for a short duration without manipulation. The display of a series of audio or video frames represents transient presentation of objects, whether captured live or retrieved from a database. Henceforth, we use the terms *static* and *transient* to describe presentation lifetimes of objects while *persistence* expresses their storage life in a database.

In the literature, media are often described as belonging to one of two classes; *continuous* or *discrete* [2-4]. This distinction is somewhat vague since time ordering can be assigned to discrete media, and continuous media are time-ordered sequences of discrete ones after digitization. We use a definition attributable to Herrtwich [5]; continuous media are sequences of discrete data elements that are played out contiguously in time. However, the term *continuous* is most often used to describe the fine-grain synchronization required for audio or video.

3 Conceptual Models for Describing Multimedia Synchronization

The problem of synchronizing data presentation, user interaction, and physical devices reduces to satisfying temporal precedence relationships under real timing constraints. In this section, we introduce conceptual models that describe temporal information necessary to represent multimedia synchronization. We also describe language and graph-based approaches to specification and survey existing methodologies applying these approaches.

3.1 Modeling Time

In information processing applications, temporal information is seldom applied towards synchronization of time-dependent media, rather, it is used for maintenance of historical information or query languages [6-7].

However, conceptual models of time developed for these applications also apply to the multimedia synchronization problem. Two representations are indicated. These are based on *instants* and *intervals* [8], described as follows.

A time instant is a zero-length moment in time, such as "4:00 PM." By contrast, a time interval is defined by two time instants and therefore, their duration, e.g., "100 ms" or "9 to 5." Intervals are formally defined as follows: let $[S,\leq]$ be a partially ordered set, and let a, b be any two elements of S such that $a \leq b$. The set $\{x \mid a \leq x \leq b\}$ is called an *interval* of S denoted by $[a,b]$. Time intervals can be decoupled from isolated instants by specifying durations, leading us to temporal *relations*.

There are thirteen ways in which two intervals can relate in time [9], whether they overlap, abut, precede, etc. These relations are indicated graphically by a timeline representation shown in Fig. 2 [10]. Only seven of the thirteen relations are shown since the remainder are inverse relations.

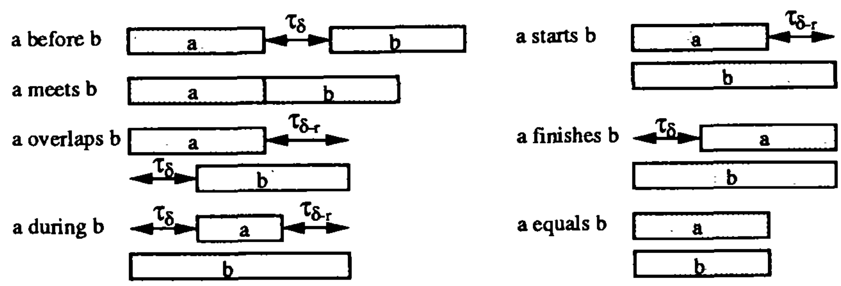

Fig. 2. Temporal Relations

Temporal intervals can be used to model multimedia presentation by letting each interval represent the presentation time of some multimedia data element, such as a still image or an audio segment. These intervals represent the time component required for multimedia playout, and their relative positioning represents their time dependencies. Fig. 3 shows audio and images synchronized to each other using the *meets* and *equals* temporal relations. For continuous media such as audio and video, an appropriate temporal representation is a sequence of intervals described by the *meets* relation. In this case, intervals abut in time, and are non-overlapping, by definition of a continuous medium.

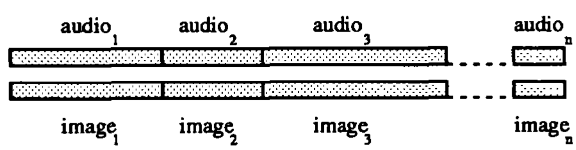

Fig. 3. Synchronization of Audio and Images

With a temporal-interval-based (TIB) modeling scheme, complex timeline representations of multimedia object presentation can be delineated. The notion of temporal intervals can also support reverse and partial playout activities. For example, a recorded stream of audio or video can be presented in reversed order. For this purpose, *reverse* temporal relations can be defined. These relations, derived from the forward relations, define the ordering and scheduling required for reverse playout. Furthermore, *partial interval* playout is defined as the playout of a subset of a TIB sequence

Under some conditions, it may be desirable to introduce incomplete timing specifications using the TIB model, as can often arise when a time-dependent stream is to be played out in parallel with a static one. For example, if an audio segment is presented in synchrony with a single still picture, the time duration for image presentation would nominally be greater than or equal to that of the audio segment. Incomplete specification can allow the static medium to assume the playout duration of the continuous medium. It is always possible to incompletely specify the timing for the parallel *equals* relation when only one medium is not static. For other types of relations, more information is required to describe the desired temporal result.

3.2 Process Synchronization

Temporal intervals and instants provide a means for indicating exact temporal specification. However, the character of multimedia data presentation is unique since catastrophic effects do not occur when data are not available for playout, i.e., deadlines are *soft* in contrast to specification techniques which are designed for real-time systems with *hard* deadlines [11]. When such specification approaches are used, the presentation of a non-decomposable multimedia element is assigned to an executable process and the processes are indicated as synchronized via inherent language constructs. To be applicable to multimedia synchronization, these methods must allow synchronization on component precedence and on real-time constraints, and provide the capability for indicating laxity in meeting deadlines. The primary requirements for such a specification methodology include the representation of real-time semantics and concurrency, and a hierarchical modeling ability. The nature of multimedia data presentation also implies further requirements including the ability to reverse presentation, to allow random access (at a start point), to incompletely specify timing, to allow sharing of synchronized components among applications, and to provide data storage of control information. Therefore, a specification methodology must also be well suited for unusual temporal semantics as well as be amenable to the development of a database for timing information.

Although some language-based models satisfy these requirements, graphical models have the additional advantage of pictorially illustrating synchronization semantics, and are suitable for visual orchestration of multimedia presentations. A graph-based model satisfying many of the requirements listed above is the Petri Net [12] which is both a graphical and mathematical modeling tool capable of representing concurrency.

3.3 Interaction and Synchronization

When a human user interacts with a multimedia system there is a requirement to synchronize the application with the user or external world. This can take the form of starting or stopping the presentation of an object, posing queries against the database, browsing through objects, or other inherently unpredictable user or sensor-initiated activities. For continuous-media systems, user interaction also implies random access to a sequential form of information. Consider a database of video stills representing scenes from an automobile, shot while looking out at a city's streets [13]. If the scenes are recorded at regular intervals, then a virtual "drive" down the street is possible through animation. When the database contains images from all possible orientations, e.g., all streets of a city, "driving" may include "turns" and corresponding jumps out of the sequential nature of the sequence of images corresponding to a street. Synchronization in this case requires coordination of the multimedia presentation with random external events created by the user. This application has been implemented [14] for interactive movies by using the hypertext paradigm [15]. The hypertext paradigm is a mode of information access and manipulation suitable for facilitating user-level interaction and database browsing.

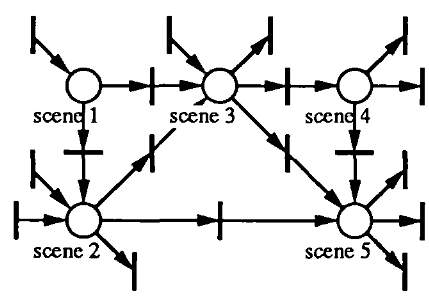

Fig. 4. PNBH Petri Net

The essence of hypertext is a nonlinear interconnection of information, unlike the sequential access of conventional text. Information is linked via cross-referencing between keywords or subjects to other fragments of information. One hypertext representation also uses Petri nets [16]. Such a Petri-Net-Based-Hypertext (PNBH) expresses information units as net places and links as net arcs. Transitions in PNBH indicate the traversal of links, or the browsing of information fragments. For example, in Fig. 4 we show a PNBH network consisting of segments of the aforementioned interactive movie. These segments can be played-out in a random order, as selected by the user and restricted by the semantics of the net.

3.4 Temporal Abstractions

Some of the requirements for multimedia presentation are not well described by either the graph or language-based specifications. For example, to reduce (slow motion) or increase (fast-forward) the speed of a multimedia presentation, the temporal models are deficient. These requirements can be addressed by *temporal abstractions*, which are means to manipulate or control the presentation of a temporal specification via time reference modification.

Various *virtual time* abstractions have been described in the literature [5,17-18]. These describe the maintenance of a time reference that can be scaled to real-time and adjusted to appropriate playout speeds. If real-time is defined as nominal clock time as we perceive it, then virtual time is any other time reference system suitable for translation to real-time (see Fig. 5). For example, a unitless reference can be converted, or projected [18] to real-time system by any scaling or offsetting operations. In this manner, the output rate and direction for a sequence of data elements can be changed by simply modifying this translation, i.e., an entire temporal specification, language or graph-based, can track a specific time reference or translation process.

We now describe existing methodologies for representation intermedia synchronization of time-dependent data.

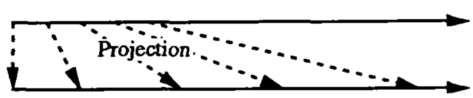

Fig. 5. Projection of Virtual Time to Real Time

3.5 Existing Temporal Specification Methodologies

An instant-based temporal reference scheme has been extensively applied in the motion picture industry, as standardized by the Society of Motion Picture and Television Engineers (SMPTE). This scheme associates a virtually unique sequential code to each frame in a motion picture [19]. By assigning these codes to both an audio track and a motion picture track, intermedia synchronization between streams is achieved. This absolute, instant-based scheme presents two difficulties when applied to a computer-based multimedia application. First, since unique, absolute time references are assumed, when segments are edited or produced in duplicate, the relative timing between the edited segments becomes lost in terms of playout. Furthermore, if one medium, while synchronized to another, becomes decoupled from the other, then the timing information of the dependent medium becomes lost. This scenario occurs when audio and image sequences are synchronized to a video sequence with time codes. If the video sequence is removed, the remaining sequences do not have sufficient timing information to provide inter-media synchronization. Instant-based schemes have also been applied using MIDI (Musical Instrument Digital Interface) time instant specification [20] as well as via coupling each time code to a common time reference [21].

The other existing approaches to timing specification for multimedia either rely on simple time precedence relationships or are based on temporal intervals. Of the ones based on intervals, most only provide support for the simple parallel and sequential relationships. Synchronization can be accomplished using a purely TIB representation, with explicit capture of each of the thirteen temporal relations [22], or with additional operations to facilitate incomplete timing specification [2].

For language-based schemes, an extension for the language CSP has been proposed to support multimedia process synchronization, including a resolution of the synchronization blocking problem for continuous media [3]. Various other language-based approaches have also been proposed, e.g., specification using LOTOS (Language Of Temporal Ordering Specification) [23], and process-oriented synchronization in CCWS [24].

The OCPN (Object Composition Petri Net) [25] represents the only graph-based specification scheme that we are aware of. The particularly interesting features of this model are the ability to explicitly capture all of the temporal relations, and to provide simulation in both the forward and reverse directions. Each place in this Petri net derivative represents the playout of a multimedia object while transitions represent synchronization points. For example, the audio and image sequence of Fig. 3 can be represented by an OCPN and is shown in Fig. 6.

Unlike the OCPN, which is a form of *marked graph* [12], net places in PNBH can have multiple outgoing arcs, and therefore can represent nondeterministic and cyclic browsing. Instead, the OCPN specifies exact presentation-time playout semantics, useful in real-time presentation scheduling. Clearly these two models complement each other for specifying both user interaction and presentation orchestration.

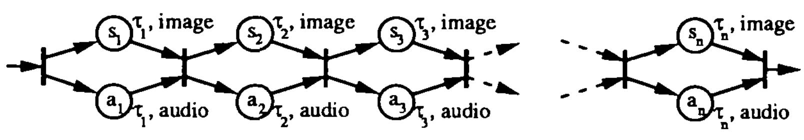

Fig. 6. Image and Audio Synchronized with the OCPN

Standardization activities have resulted in several approaches to synchronization for electronic documents, including hypermedia [15]. For electronic document representation and interchange, the Office Document Architecture (ODA) has been standardized [26]. This standard describes parallel, sequential, and independent temporal control [27-28] but does not support synchronization for continuous types. However, work is underway to extend the ODA model for this purpose [2]. HyTime (Hypermedia/Time-based Structuring Language) [18] and the Hytime application SMDL (Standard Music Description Language) [29] are language-based approaches to synchronization based on SGML (Standard Generalized Markup Language, ISO 8879). The HyTime specification provides a scripted form of language specification.

4 Database Aspects of Multimedia Synchronization

Once time-dependent data are effectively modeled, a MDBMS must have the capability for storing and accessing these data. This problem is distinct from historical databases, temporal query languages [6-7], or time-critical query evaluation [30]. Unlike historical data, time-dependent multimedia objects require special considerations for *presentation* due to their real-time playout characteristics. Data need to be delivered from storage based on a prespecified schedule, and presentation of a single object can occur over an extended duration, e.g., a movie. In this section we describe database aspects of synchronization including conceptual and physical storage schemes, data compression, operating system support, and synchronization anomalies.

4.1 Conceptual Storage Models

A conceptual data model for time-dependent data must support forward, reverse, and random access in addition to conventional DBMS queries. Temporal intervals can be described by a timeline representation in an unstructured format, or in a structured format such as the OCPN. Using the OCPN, temporal hierarchy can be imparted to the conceptual schema as sets of intervals bound to a single temporal relation can be identified and grouped. For example, this process is applied to the OCPN of Fig. 6, resulting in the conceptual schema of Fig. 7.

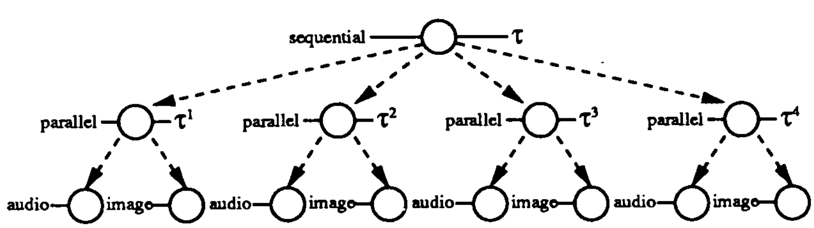

Fig. 7. Conceptual Schema Based on an OCPN

With this approach, the conceptual schema forms a temporal hierarchy representing the semantics of the OCPN. Subsets or subtrees of this hierarchy represent subnets of the OCPN, illustrating the capability of composing complex multimedia presentations. Terminal elements in this model indicate base multimedia objects (audio, image, text, etc.), and additional attributes can be assigned to nodes in the hierarchy for conventional DBMS access. Timing information is also captured with node attributes, allowing the assembly of component elements during playout.

Temporal information can also be encapsulated in the description of the multimedia data using the object-oriented paradigm [5]. Temporal information including a time reference, playout time units, temporal relationships, and required time offsets are maintained for specific multimedia objects. For stream types, this approach can define the time dependencies for an entire sequence by defining the period or frequency of playout (e.g., 30 frames/s for video) analogous to a set of intervals bound to a single temporal relation.

4.2 Physical Storage Models

Given the conceptual synchronization requirements for a multimedia application, the physical system must meet these requirements. Problems arise due to the strict timing requirements for playout of time-dependent data. In this section, approaches to physical storage of time-dependent data are discussed.

The multimedia types of audio and video require very large amounts of storage space and will exist, when not live, in secondary storage. In order to meet the presentation requirements for these data, there are some obvious storage organizations to facilitate data transfer from secondary storage to display. For example, data can be stored in contiguous blocks on disk in the same order as playout. If disk transfer rates are not attainable for a certain data type, then disk interleaving can be used to produce the necessary data rates, as has been successfully implemented for a monochrome video-in-windows (VW) display [31]. When multiple streams originate from the same storage device, interleaving of data is necessary both to maintain data rates suitable for the quality of the stream, as well as to prevent conflict between the interacting streams. An approach to the placement of two audio data streams on a disk is described by Yu et al. [32-33].

These data placement schemes rely on maintaining a fixed transfer rate between disk and display. In the event that the system becomes busy with some other task, it is possible to corrupt the playout sequence, causing a perceptible shortage of data resulting in a blank screen or silence in video and audio output. An alternate approach proposed resolves this contention problem by providing a variable quality of data transfer if there is contention in the system [34]. The key to this scheme is the storage of data frames in such a manner as to provide high or low-resolution data retrieval of the same sequence.

4.3 Data Compression

Since multimedia data types have enormous storage and communications requirements, data compression is desirable, if not essential to enable multimedia applications. In this section, some of the implications of data compression on continuous-mode data are discussed.

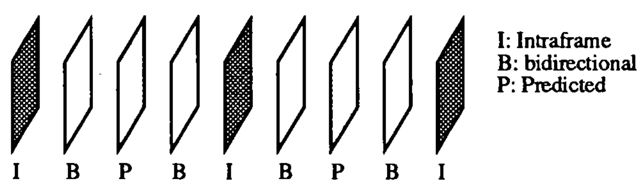

Fig. 8. MPEG Coding Scheme

Many data encoding formats exist for stream-type data. For most uncompressed formats we can delineate fixed data units, for example, video frames at 30 frames/s or audio samples at 3000 samples/s. Since these data are produced at regular intervals, they are called constant bit rate (CBR). Compression schemes for stream-type data can result in CBR or variable bit rates (VBR). The advantage of compressed data is clearly the savings in storage space and communication bandwidth. However, it becomes more difficult to identify points of synchronization between streams requiring synchronization when VBR compression is applied. The reason is as follows. Compression schemes use both intra and inter-frame coding. Intra-frame coding applies compression schemes within a single time-dependent frame. Therefore, a timing specification can apply to the self-contained frame before, during, and after compression. For inter-frame coding, compression schemes apply across a sequence of frames. For the proposed MPEG (Moving Picture Experts Group) coding scheme [35], differential values are generated for sequences of frames between inter-frame coded ones. This approach presents several problems for synchronization. First, it is desirable to have the ability to start at an arbitrary point in the continuous stream. With inter-frame coding this is not possible without first regenerating intermediate frames. Second, in order to provide reverse presentation, differential values must be available in both directions. The ability to begin presentation at an arbitrary point in a stream or to choose direction is part of the larger problem of providing random access or random insertion points into a stream-type object. These problems are approached in the MPEG scheme. To provide random insertion points, intra-frame coding is used at intervals as often as required for access as specified by the application. Reverse playout is accommodated by bidirectional differential frames (see Fig. 8).

4.4 System Support for Synchronization

For supporting time-dependent media, a MDBMS must deal with storage device latencies including ones due to data distribution across a network. MDBMS support for time-dependent data requires an operating system that is tailored to the specific requirements of real-time multimedia data. Unlike *hard* real-time systems, the inability to meet a deadline for multimedia data is unlikely to cause a catastrophic result. However, the design of such a system must account for latencies in each system component in the delivery of data from storage to the user. Specific scheduling is required for storage devices, the CPU, and communications resources. Recent work in the design of systems support for multimedia data includes [36-38], the details of which are beyond the scope of this paper.

Similarly, providing a transport mechanism for time-dependent data requires managing the resources of a computer network. For delay-sensitive media, these resources are communication bandwidth and end-to-end delay. The problem of synchronization across a network is most acute when providing intermedia synchronization for multiple independent stored-data sources. In this case, to achieve intermedia synchronization, random network delays on each connection must be overcome, in spite of variations in clock rates at each remote data source. Typically, the delay variations on each channel are estimated during connection set-up, and an end-to-end delay, called a *control time* is introduced, representing an interval over which buffering is applied. The result of this buffering is a reshaping of the channel delay distribution to reduce variance. This process is shown schematically in Fig. 9, where *p(t)* is the delay density function and *w(t)* is the reconstructed, playout density.

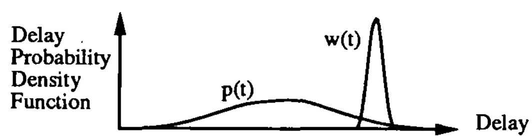

Fig. 9. Result of Buffering

Recent proposals to provide continuous media transport in a network rely on a connection establishment phase that reserves sufficient resources to provide a multimedia connection [36-38]. These reservation protocols evaluate system resources including available bandwidth and channel delay properties. In the event that sufficient resources are not available, the connection is refused. Once a connection is established, buffering is configured to provide the requested level of service in terms of delay and late-arriving data elements. For stored-data applications, the system has flexibility in scheduling the retrieval of data since the playout schedule is known *a priori*. For live multimedia sources, existing stream communication protocols can be applied. Since these protocols assume that some data can arrive late, the system must be able to accommodate shortages of data rather than introduce anomalous playout behavior.

4.5 Synchronization Anomalies

When data are delayed and are not available for playout, a synchronization anomaly occurs. At the output device, this can result in a gap in the sequence of presented elements, or a shortage of data to present per unit time. Policies for handling late-arriving data include discarding them or changing the playout rate to maintain a constant number of buffered elements. When data are lost or discarded, reconstruction can also be used. Steinmetz [3] proposes performing some alternate activity when a data element is not available, such as extending the playout time of the previous element. Generally, when gaps in a data sequence are ignored with respect to the playout rate, the loss of data elements when subsequent data are available advances the sequence in time, and can be corrected by slowing the playout rate until the schedule is correct. Approaches to synchronization of received packets include varying the playout rate and the utilization of received data [39]. The *expansion* method lets each packet be played out even if late. The result is the delay of all successive packets and an accumulation of skew with time. Another approach is to *ignore* some data since much redundant information is contained in the data streams, thereby preserving the duration of the overall sequence. This is analogous to a reduction in packet utilization [40]. One further gap-compensating technique for continuous media reconstructs the missing data elements. This approach is to substitute alternate data for the missing data in the stream. The data are chosen as null or non-null values (zero amplitude and waveform stuffing), or are interpolated from previous values [41].

5 Conclusion

In this paper, an overview of the many aspects of the multimedia synchronization problem are reviewed at the physical, service, and human interface levels of integration. Synchronization of this type has been recognized as an important requirement for enabling multimedia applications. Significant issues remain for providing time-dependent delivery of multimedia data in a general multimedia information system. The primary issues are specification and storage of temporal information describing the time dependencies of multimedia data, provision of an enforcement mechanism for temporal specifications, and accommodation of the laxity in the retrieval of time-dependent data when they are not available for playout.

Acknowledgement

This research was supported in part by the New York State Center for Advanced Technology in Computer Applications and Software Engineering (CASE) at Syracuse University, and the NYNEX Corporation.

6 References

[1] Sventek, J.S., "An Architecture for Supporting Multi-Media Integration," *Proc. IEEE Comp. Soc. Office Automation Symposium*, April, 1987 pp. 46-56.

[2] Herrtwich, R.G., Delgrossi, L., "ODA-Based Data Modeling in Multimedia Systems," *International Computer Science Institute Tech. Rept.* TR-90-043, August 1990.

[3] Steinmetz, R., "Synchronization Properties in Multimedia Systems, "*IEEE J. Selected Areas in Comm.*, Vol. 8, No. 3, April 1990, pp. 401-412.

[4] Nicolaou, C. "An Architecture for Real-Time Multimedia Communication Systems,"*IEEE J. Selected Areas in Comm.*, Vol. 8, No. 3, April 1990, pp. 391-400.

[5] Herrtwich, R.G., "Time Capsules: An Abstraction for Access to Continuous-Media Data," *Proc. 11th Real-Time Systems Symp.*, Lake Buena Vista, FL, December 1990, pp. 11-20.

[6] Snodgrass, R., "The Temporal Query Language TQuel," *ACM Trans. on Database Systems*, Vol. 12, No. 2, June 1987, pp. 247-298.

[7] Tansel, A.U., Arkun, M.E., Ozsoyoglu, G., "Time-by-Example Query Language for Historical Databases," *IEEE Trans. on Software Engineering*, Vol. 15, No. 4, April 1989, pp. 464-478.

[8]. Anderson, T.L., "Modeling Time at the Conceptual Level," *Improving Database Usability and Responsiveness*, P. Scheuermann, Ed., Academic Press, New York, 1982, pp. 273-297.

[9] Hamblin, C.L., "Instants and Intervals," *Proc. of the 1st Conf. of the Intl. Society for the Study of Time*, J.T. Fraser, et al., Ed, Springer-Verlag, New York, 1972, pp. 324-331.

[10] Allen, J.F., "Maintaining Knowledge about Temporal Intervals," *Comm. of the ACM*, November 1983, Vol. 26, No. 11, pp. 832-843.

[11] Faulk, S.R., Parnas, D.L., "On Synchronization in Hard-Real-Time Systems," *Comm. of the ACM*, Vol. 31, No. 3, March 1988, pp. 274-287.

[12] Peterson, J.L., "Petri Nets," *Computing Surveys*, Vol. 9, No. 3, September 1977, pp. 225-252.

[13] Lippman, A., "Movie-Maps: An Application of the Optical Videodisc to Computer Graphics," *Computer Graphics*, July 1980, pp. 32-42.

[14] Sasnett, R.M., "Reconfigurable Video," *M.S. Thesis*, MIT, 1986.

[15] Conklin, J., "Hypertext: An Introduction and Survey," *IEEE Computer*, September 1987, pp. 17-41.

[16] Stotts, P.D., Furuta, R., "Temporal Hyperprogramming," *J. of Visual Languages and Computing*, Vol. 1, 1990, pp. 237-253.

[17] Anderson, D.P., Tzou, S.Y., Wahbe, R., Govindan, R., Andrews, M., "Support for Continuous Media in the Dash System," *Proc. 10th Intl. Conf. on Distributed Computing Syst.*, Paris, France, May 1990, pp. 54-61.

[18] Committee Draft International Standard, "Information Technology - Hypermedia/Time-based Structuring

Language (HyTime)," ISO/IEC CD 10743, April 1, 1991.
[19] SMPTE Recommended Practice: RP 136-1986, "Time and Control Codes for 24, 25, or 30 Frame-Per-Second Motion picture Systems," April 1986, *SMPTE J.*, August 1986, pp. 862-864.
[20] Moore, D.J., "Multimedia Presentation Development Using the Audio Visual Connection," *IBM Systems J.*, Vol. 29, No. 4, 1990, pp. 494-508.
[21] Hodges, M.E., Sasnett, R.M., Ackerman, M.S., "A Construction Set for Multimedia Applications," *IEEE Software*, January 1989, pp. 37-43.
[22] Schramm, C., Goldberg, M., "Multimedia Radiological Reports: Creation and Playback," *Proc. SPIE Medical Imaging III: Image Capture and Display*, SPIE Vol. 1091, Newport Beach CA, January 1989, pp. 191-201.
[23] Weiss, K.H., "Formal Specification and Continuous Media," *Proc. 1st Intl. Workshop on Network and Operating Support for Digital Audio and Video*, Berkeley, CA, November 1990.
[24] Poggio, A., Garcia-Luna-Aceves, J.J., Craighill, E.J., Moran, D., Aguilar, L., Worthington, D., Hight, J., "CCWS: A Computer-Based Multimedia Information System," *IEEE Computer*, October 1985, pp. 92-103.
[25] Little, T.D.C., Ghafoor, A., "Synchronization and Storage Models for Objects," *IEEE J. on Selected Areas in Comm.*, April 1990, pp. 413-427.
[26] International Organization for Standardization, ISO Document No. 8613, ISO, Geneva, March 1988.
[27] Garcia-Luna-Aceves, J.J., Poggio, A., "Multimedia Message Content Protocols for Computer Mail," *Proc. IFIP WG 6.5 Working Conf. on Computer-Based Message Services*, Nottingham England, May 1984, Hugh Smith, Ed., North-Holland, New York, pp. 87-98.
[28] Postel, J., "A Structured Format for Multimedia Documents," Report RFC767, DDN Network Information Center, SRI International, Menlo Park CA, 94025, USA, August 1980.
[29] Committee Draft International Standard, "Information Technology - Standard Music Description Language (SMDL)," ISO/IEC CD 10743, April 1, 1991.
[30] Hou, W.C., Ozsoyoglu, G., Taneja, B.K., "Processing Aggregate Relational Queries with Hard Time Constraints," *Proc. ACM SIGMOD Intl. Conf. on the Management of Data*, Portland, OR, June 1989.
[31] Katseff, H.P., Gaglianello, R.D., London, T.B., Robinson, B.S., Swicker, D.B., "An Overview of the Liaison Network Multimedia Workstation," *Proc. 1st Intl. Workshop on Network and Operating Support for Digital Audio and Video*, Berkeley, CA, November 1990.
[32] Yu, C., Sun, W., Bitton, D., Yang, Q., Bruno, R., Tullis, J., "Efficient Placement of Audio Data on Optical Disks for Real-Time Applications," *Comm. of the ACM*, Vol. 32, No. 7, July 1989, pp. 862-871.
[33] Wells, J., Yang, Q., Yu, C., "Placement of Audio Data on Optical Disks," *Proc. 1st Intl. Conf. on Multimedia Information Systems '91*, Singapore, January 1991, pp. 123-134.
[34] Park, A., English, P., "A Variable Rate Strategy for Retrieving Audio Data from Secondary Storage," *Proc. 1st Intl. Conf. on Multimedia Information Systems '91*, Singapore, January 1991, pp. 135-146.
[35] Le Gall, D., "MPEG: A Video Compression Standard for Multimedia Applications," *Comm. of the ACM*, Vol. 34, No. 4, April 1991, pp. 46-58.
[36] *Proc. 1st Intl. Workshop on Network and Operating Support for Digital Audio and Video*, Berkeley, CA, November 1990.
[37] Little, T.D.C., Ghafoor, A., "Multimedia Synchronization Protocols for Broadband Integrated Services," to be published in *IEEE J. on Selected Areas in Comm.*, 1991.
[38] Anderson, D.P., Herttwich, R.G., Schaefer, C., "SRP: A Resource Reservation Protocol for Guaranteed-Performance Communication in the Internet," *International Computer Science Institute Tech. Rept.* TR-90-006, February 1990.
[39] Naylor, W.E., Kleinrock, L., "Stream Traffic Communication in Packet Switched Networks: Destination Buffering Considerations," *IEEE Trans. on Comm.*, Vol. COM-30, No. 12, December 1982, pp. 2527-2524.
[40] Little, T.D.C., Ghafoor, A., "Network Considerations for Distributed Multimedia Object Composition and Communication," *IEEE Network*, November 1990, pp. 32-49.
[41] Suzuki, J., Taka, M., "Missing Packet Recovery Techniques for Low-Bit-Rate Coded Speech," *IEEE J. on Selected Areas in Comm.*, Vol. 7, No. 5, June 1989, pp. 707-717.

Multimedia Database Development in Japan

Gio Wiederhold *, David Beech †, and Toshi Minoura ‡

Contributions by Charles Bourne, Nick Farmer, Sushil Jajodia,
Nick Farmer, Duane Shelton, and Diane Smith

1. Summary

This report presents some of the findings of a group of database experts, based on an intensive study trip to Japan in March 1991. The study was organized be the Japan Technology Evaluation Center (JTEC) at Loyola College, and was sponsored by the National Science Foundation (NSF). Academic, industrial, and governmental sites were visited. We found much activity in two key areas, multimedia and object-oriented databases. We can expect to see future export of Japanese database products based on this technology, typically integrated into larger systems.

Technology transfer in Japan is enhanced through governmental support. Although direct research funding is surprisingly small, the governmental support is intended to encourage activities of various committees in which academic and industrial researchers and developers together participate. Frequent and in-depth communication takes place among committee participants.

The full report will be available later this year from the National Technical Information Service (NTIS), 5285 Port Royal, Virginia 22161. It contains sections on the role of the Japanese government in database activities, on the various database technologies being pursued in software and hardware, and on the market drivers for import and export. Short summaries of activities at 31 sites are also included, as is a bibliography of relevant papers and reports. This paper has been prepared from the draft of the full report. Its content is thus preliminary and may be corrected and augmented in the final report.

We start by presenting some relevant technology topics and finish with general observation on support of these activities in Japan.

2. Object Databases

There is considerable activity in Japan in the field of object databases. Indeed, there is a worldwide sense that objects are likely to be at the heart of the next generation of database technology. This may well be a self-fulfilling prophecy, since anything can be described as an object, and the technical sense of the term is still negotiable in the object database world. Even in Japan, the land of consensus, there is no commonly agreed definition of what an object database is — consensus takes time. In all cases there should be significant extensions over current, value-oriented or record-oriented databases.

There are three major trends in the development of object database systems — namely to concentrate (a) on multimedia objects, (b) on complex objects, or (c) on programming language objects. There is considerable conceptual overlap among these approaches, but the differences in emphasis can produce very different systems. In this extract we limit ourselves to multimedia objects.

The simplest use of object concepts occurs in the storage and retrieval of multimedia information. The binary representation of a voice message, a color image, or a video recording is intuitively thought of as an *object* due to its individuality and its synthetic nature — it is as though it springs into life when created, and it is highly unlikely that it is the same as other objects unless there has been explicit copying. In everyday

* Department of Computer Science, Stanford University, Stanford, CA 94305; gio@cs.stanford.edu.
† Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA 94064; dbeech@us.oracle.com.
‡ Department of Computer Science, Oregon State University, Corvallis, OR 97331-3902; minoura@mist.cs.orst.edu.

life, except to a handful of mathematicians, it is too unwieldy to consider all such multimedia phenomena to be instances of a preexisting domain of immutable *values*.

Since there is marketing value in claiming object support, one is liable to encounter such claims on behalf of relational or other systems capable of handling such "binary large objects" (BLOBs).

The positive side of this use of the term "object" is that it is the first step in the direction of a more general conceptual model that would fully justify the terminology. If a relational system is extended, not just with a type `LONG`, but with types `VOICE`, `IMAGE`, and `VIDEO`, say, and these types have appropriate functions defined on them, and can be given more specialized subtypes, then such a system begins to share the rich model of other object systems, although perhaps offering only limited extensibility.

When this perspective on multimedia objects is allied to the market potential for their support by databases in the office — and the home — of the future, it should be clear that this can be reputable use of the terminology of object databases.

2.1 Industrial Projects in Object Database Systems

Although there are not yet any Japanese object database products, whereas in the United States there are several, we found a number of substantial prototype systems in Japan, and a great deal of research activity which generally showed a good blend of theoretical and practical considerations. Japanese researchers are well informed in this area and appear to find it congenial territory, which helps contribute to rising standards in both content and presentation of their database literature.

Three industrial object database prototype systems were described to us: MANDRILL (Hitachi), JASMINE (Fujitsu), and ODIN (NEC).

Mandrill

The MANDRILL system at Hitachi's Systems Development Laboratory at Kawasaki aims to improve database productivity, applicability, and performance. The range of applications is seen as very wide, including not only Computer-Aided Design, Manufacturing, and Software Engineering (CAD/CAM/CASE), but also Office Information Systems (OIS), Management Information Systems (MIS) and Geographic Information Systems (GIS). Many of the points of the Object-oriented Database System Manifesto (presented by Atkinson et al. in Kyoto, 1989) are being addressed. A preliminary version of MANDRILL (version 1) was built in 1988/89, with the version 2 prototype scheduled for completion in 1991, and application studies in C planned for 1992.

The approach at Hitachi has been to build an object layer over an existing DataBase Management System (DBMS) and its file system interfaces, with the intention of supporting multimedia information as well as conventional data. The central part, MANDRILL/CORE, first provides a logical media layer to establish system-independence from the underlying database and file systems, and then builds its object layer over this. Applications access the object layer via statements of the MANDRILL/QUEST database language embedded in programming languages.

The QUEST language is somewhat like SQL, although no attempt has been made at a close integration or compatible extension. QUEST commands are treated as system-defined methods (functions) that can be applied to objects. Users can also define their own methods if they wish, and the object model has the usual class hierarchy. In addition, attributes of a class may be tagged as inverses of other attributes or as referring to parts of an aggregate object. A complex object defined by means of attributes with `has_parts` tags may be processed as a whole, for example, being retrievable by a single `complexselect` statement. Database integrity is supported by offering user-defined triggers, and authorization controls are applied at the granularity of the object instance.

Jasmine

The JASMINE system is being developed by Fujitsu Laboratories Ltd., a Fujitsu subsidiary in Kawasaki. It is an object database system that supports set-oriented queries containing path expressions. Constraints and demons (triggers) may be defined, and the system may be extended by new nodes, attributes, and links defined procedurally in JASMINE/C. Both textual and image data are supported in the window interface, with hypermedia applications in mind.

At the conceptual level, papers have been published during the past three years on schema design, object views, complex objects, behavioral consistency, natural joins, versions ("alternative objects"), and ILOG (a language for generating and manipulating object identifiers)

Extensions to relational systems have included a BLOB approach to storing and retrieving PostScript objects in the database, and SQL-Navigator which dynamically attaches methods to objects to support a hypertext style of navigation.

Recent work has been on TextLink which supports hypertext on an object database system. We observe the trend to move from a good general understanding of the object database field to specific applications in the hypermedia world.

Library and Information Science

The Library and Information Science University in Tsukuba specializes in information management. Preliminary work over the past six years has led to a project to develop a system named OMEGA (Object-oriented Multimedia database Environment for General Application). Current emphasis is on the part-of relationship and referencing and synchronization issues in dealing with complex objects.

Hokkaido

Current activity at Hokkaido University is largely in the multimedia area and was preceded by more general object database investigations. In particular, their object data model ODM brings together object ideas from SMALLTALK-80 and set-oriented ideas from the relational database model. The concept of a relation as a set of tuples is generalized to that of the "u-set" (uniform set) of objects. The uniformity of a u-set resides in the fact that all its elements must be conformable to the base class of the u-set, i.e. must satisfy its interface. The u-set itself is defined as an object class with appropriate methods, both for the usual set operations, and for relational operations such as projection, Cartesian product, and join on the instance variables of objects regarded as tuples.

National Ethnology Museum

The National Ethnology Museum operates a system to store high quality 3-dimensional color images of museum artifacts, along with the associated descriptive information. A file of over 50 000 such images has already been assembled. With 7 megabytes/image, this represents a total storage volume of 350 gigabytes. The in-house retrieval capability includes search and display of specific color images based on their indexed attributes. A dialup capability was planned for researchers who were remote from this facility.

3.2 Obase Consortium

A novel development is the formation of the Obase consortium, intended to run from September 1990 through August 1993. Its organizer is the Senri International Information Institute (Osaka) and participants include universities and industries. The objectives are to investigate current ODBMS and applications, to develop a new ODBMS (OBASE) together with object database design tools, and to carry out related research.

The project began by investigating the advantages and disadvantages of a number of conventional Object-oriented DBMSs and then moved on to the development of some specific applications. The emphasis in the current, preliminary applications study appears to be hypertext, video, and all forms of document management. The Obase researchers are using the VERSANT system to implement these applications, including a fuzzy retrieval system for a movie database, a LATEX document database, etc.

4. Interoperable Databases

We expect that in the future most database activity will involve multiple computers. When computing activities are distributed, then issues of joint operation must be addressed. In object and multimedia databases new standards are needed to achieve effective interoperation. The Japanese understand the importance of interoperability among computers of different sizes and from different vendors. Interoperation is essential in the "information-based society" that they envision.

4.1 Interoperability and Standards

The most notable move toward the vision of an information-based society is the National Research and Development Program on Interoperable Database System (Interoperability Program), managed through INTAP (Interoperability Technology Association for Information Processing). INTAP is an industry consortium created to perform Research and Development focused on open network systems architecture. It was started by MITI in 1986 with a total projected budget of $100M. The official goal of the 5-year program is "To establish fundamental technology indispensable for realizing multimedia reliable distributed database systems on interoperable computer network systems."

The projects supported by the program are divided into four areas: 1) database systems architecture, 2) multimedia technology, 3) distributed systems technology, and 4) open network systems architecture. Most initial focus has been on the fourth area. Important is an emphasis on the RDA standard for remote data access.

Few results have been reported that are directly related to multimedia technology. It is obvious that first the infrastructure has to be established before higher-level standards-proposing and -setting activities proceed.

4.2 Adaptation to the Japanese Market

To reach the Japanese market foreign products have to be adapted at least to handle Japanese character sets. In earlier foreign products little consideration was given to handling non-ASCII or non-EBCDIC character representations. The absence of a single standard for Japanese characters hinders foreign developers as well. The simpler conventions do not provide for complications in practical japanese text as

- integration of roman numbers and terms
- vertical and horizontal orientation
- the font changes concomitant with changes in orientation
- the complication induced by the combination of the first two points
- the layout problems induced by mixed orientations.

Some Japanese companies have invested greatly in making these adaptations. The experience means that they had to address representation problems within the DBMS. This experience is very helpful when moving to support large, variable sized data elements (BLOBS).

Having BLOBS is essential for multimedia applications. It will depend on the effectiveness of these extensions, whether established DBMSs will be used for the multimedia services of the future. Otherwise, it will be necessary for the developers of multimedia systems to develop their own DBMSs. The availability of standards, as SQL and RDA, makes entry of new DBMSs that satisfy these standards feasible. Even if they are less mature, having multimedia capability can be a decisive factor in the market.

Intermediary solutions do exist. Conventional DBMSs can reference images and other large objects in distinct files, and these can be accessed indirectly. However, such solutions are more complex to manage and are likely to be transient. Furthermore, if associative access to image and voice data becomes a reality, than the indirect approach will no longer be feasible.

5. Professional Support

Two professional societies promote database research: Information Processing Society of Japan (IPSJ) established in 1970 and the Institute of Electronics, Information, and Communication Engineers (IEICE) established in 1917. Two special interest groups within them focus on databases: the Special Interest Group on Database Systems (SIGDBS) associated with IPSJ, and the Special Interest Group on Data Engineering (SIGDE) within IEICE. The IEICE is roughly equivalent to the IEEE in the U.S. Each society has its own journals and annual conferences and holds workshops on various aspects of databases.

There is now a Japanese-printed English-language publication: the monthly *IEICE Transactions*. It publishes papers in four categories:

1. Fundamentals of Electronics, Communication and Computer Sciences,
2. Communications,
3. Electronics,
4. Information and Systems.

This Japan-originated English publication is not well known now. The *IEICE Transactions* also contains English abstracts of Japanese-language papers published in the two journals *IEICE Computer* and *IEICE Information Processing*, which cover the IEICE areas of Information and Systems. The *IEICE Transactions* also list the English titles of the unrefereed workshop reports from the 47 technical groups of the IEICE.

The industry providing databases for information services is represented by the Japan Database Industry Association (DINA). Established in 1979, DINA has over 100 members and is quite active.

6. Government and Database Research and Development

Most of the work in multimedia databases is sponsored by industry. Government participation is financially small, but effective. There are of course substantial government-funded research laboratories for which governmental support is critical.

The Ministry of Industry and Technology (MITI) carries out planning and promotion of technology directions. Noteworthy is the existence of a Database Promotion Center (DPC), established in 1984. DPC has as its goals the promotion of database construction, research and development of the basic technologies related to databases, establishment of efficient clearing services, education propagation and training associated with databases, and international information exchange informing other countries about Japanese databases.

Japanese online databases in science and technology were first developed by MITI's information service (JICST), and the Japan Patent Office, in the mid 1970s, after business databases had become operational. MITI publishes an annual Database Directory, which provides a comprehensive list of databases accessible in Japan, both Japanese and foreign. The 1989 Directory, in Japanese, is about 3000 pages long and lists over 3000 databases. An English summary version is about 100 pages.

DPC keeps track of various database statistics such as the following: "MITI's survey reports that the Japanese database service industry's sales reached Y106.3 billion in 1988 ... In their *Outlook of the Information Industry for the year 2000*, submitted in 1987, estimated sales were placed at Y144.5 trillion for the entire information industry in the 2000's, and at Y3.4 trillion for information provision services including database services. MITI also runs the Patent Office which is highly automated with a huge database operation.

Furthermore, MITI sponsors certain National Projects. ICOT is one (see below) with an Object-Oriented Database as a subproject. Another is the Interoperable Database System Project (INTAP), which endeavors to establish an open system architecture in Japan, although MITI is also supporting the International Multimedia Association, which is an industry association and is expected to initiate some international multimedia project. One of the most important functions of ICOT is the ability of young database researchers to come together outside of their parent companies. Several working groups have been formed and meet periodically.

The Science and Technology Agency (STA) of the prime minister's office supports, a,o., a new project titled: *Development of Self Organizing Information Base to Aid Researchers in Creative Research.* It has the goal of creating a more intelligent large hypertext database system. Project members include STA's information service (JICST), Ryukoku, Tsukuba, Kobe, and Hokkaido Universities, ETL, NTT, FujiXerox, and other industrial partners.

The Ministry of Education also provides support for information dissemination. The NACSIS center at the University of Tokyo provides information retrieval. It includes about three million entries from Japanese university library catalogs. This service is available in the U.S. through NSF.

The communication part of this is the Science Information Network (SIN) which is a privately operated packet switching network employing packet multiplexers installed throughout Japan at major research areas. This network will serve as the main communication facility for the Inter-University Computer Network of Japan (an N-1 network). Plans are to permit full texts, graphics, voice and image communications.

Government and Progress

The government, overall, seems to have less influence than perceived from the outside. The funding supports some laboratories and projects, but academics are able to use the support for their own directions. The amounts for industrial laboratories are relatively small and do not influence market-driven priorities. However, these projects do cause regular meetings of academic, government, and industrial researchers, increasing mutual awareness and understanding.

7. Conclusion

An important driving mechanism in multimedia database development is the Japanese capability in electronic products. High quality image acquisition, transmission, storage, and display are emphasized. Voice data are also of concern. Database management systems are being expanded to provide such "multimedia" support. While fundamental database management systems are not being advanced in Japan, the incorporation of multimedia support will change their character greatly. We can expect that purchasers of systems with multimedia requirements — a field likely to grow rapidly — will, with the Japanese image-processing hardware, acquire Japanese database software. Certain subareas, as computer-aided-design and computer-aided-education, that are critically dependent on graphics, will be the initial users of this technology. multimedia technology is likely to provide an important path for the introduction of Japanese database software into European and U.S. markets.

Acknowledgment

The study was supported by NSF and DoC, and indirectly by Oracle Corp. The authors wish to thank Kyu-Young Whang and Sang K. Cha for their help in producing the final version of this paper.

Bibliography

For an overview of the topic, readers are referred to the following article:

[1] Shojiro Nishio, "Status Update on Database Technology in Japan," *IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*, Victoria, B.C., Canada, May 1991.

The full report will contain an extensive list of references. A selected subset is attached below.

[2] Aoshima, M., Izuonida, Y., Makinouchi, A., Suzuki, F., and Yamana, "The C-based Database Programming Language Jasmine/C," Fujitsu Labs; Kyushu Un., rcvd Mar.1991.

[3] Chang, T-S. and Tanaka, K., "Alternative Objects in Object-Oriented Databases," *Proc. Int. Symp. Database Sys. for Adv. Appl*, Seoul, Korea, 1989.

[4] Michael A. Cusumano, "Japan's Software Factories: A Challenge to W.S. Management," Oxford Un. Press, 1990.

[5] Database Promotion Center, "Database in Japan 1990," MITI, Tokyo, Japan, 1990, 81pp.

[6] Mohammed ElSharkawi, "Answering Queries in Temporal Object-Oriented Databases," *Proc. 2nd Int. Symp. Database Sys. for Adv. Appl*, Tokyo, Apr.1991.

[7] Yochihiro Hirosawa, "Overview of ODIN," presentation material, NEC Corp, C&C Systems Res. Labs, rcvd. Apr.1991.

[8] ICOT, "Outline of Fifth Generation Project," TR, 1990.

[9] Akikazu Ida, "An Overview of Full Color Image Database System – SR-ret," presentation material, NTT Data, Mar.1991.

[10] Iizawa, A. and Shirata, Y., "Optical Disk-Based Multi-media Database," TR, Ricoh, Software Res. Ctr, Tokyo, rcvd. Mar.1991.

[11] Hiroshi Ishikawa, "An Object-Oriented Knowledge Base Approach to a Next Generation of Hypermedia System," *IEEE Compcon*, 1990.

[12] David Kahaner, "Japanese Database Activities," rcvd. Apr.1991.

[13] Kato *et al.*, "TRADEMARK: Multimedia database with abstracted representation on knowledge base," *Proc. 2nd Intl. Symp. on Interoperable Info. Systems*, 1988.

[14] Kato, K., Fujisawa, H., Ohyama, M., Kawaguchi, H., Hatakeyama,A., Kaneoka,N., and Akizawa,M., "An Index-Free Full-Text Search Machine for Large Japanese Text Bases," *Data Engineering*, December 14, 1989

[15] Kato, T. and Mizutori, T. "Multimedia Data Model for Advanced Image Information System," *Adv. Database Sys. Symp*, Kyoto, 1989.

[16] Kato, T. and Kurita, T., "Visual Interaction with Electronic Art Gallery," *Database and Expert Systems Applications*, Tjoa and Wagner (Eds), S-V, 1990.

[17] Kato et al., "A Cognitive Approach to Visual Interaction," *Proc. Intl. Conf. on Multimedia Info. Systems*, Jan. 1991.

[18] ICOT Librarian, "Publication List," *ICOT report*, Nov.1990. (List of available TRs and memoranda in English)

[19] MITI, "Report of the Research Committee on the New Information Processing Technology," Industrial Electronics Div., Machinery and Information Industry Bureau, MITI, Tokyo, Mar.1990. (Basis for 6th generation project, calling for international cooperation)

[20] MITI, "Interoperable Database System," rcvd. Mar.1991. (INTAP overview)

[21] Maeda, D., Tanaka, S., Hirto, T., Intatsumata, T., and Shibayama, J., "A Multimedia Database System Featuring Similarity Retrieval " *Proc. INTAP*, 1988.

[22] Makinouchi, M. and Aritsugi, M., "The Object-Oriented Persistent Programming Language for Multi-media Databases," TR CSCE 91-C04, Dept CS and Comm.Eng. Kyushu Un., 1991.

[23] Masunaga, Y., "Object Identity, Equality and Relational Concept," *Deductive and Object-Oriented Databases*, Kim, Nicholas, and Nishio (eds), Elsevier N-H, 1990.

[24] Masunaga, Y., "Design Issues of OMEGA: An Object-Oriented Multimedia Database Management System," *J. of Inf. Proc.*, Vol.14, No. 1, 1991.

[25] Morita, Y, Haniuda, H., and Yokota, K., "Object Identity of Quixote," *DBS8-12 IPSJ*, Nov.1990

[26] Mozaffari, M. and Tanaka, Y., "OBM: an Object Oriented Data Model," *NGC*, Vol.7, 1989, pp.3–35.

[27] National Committee of FID of Japan, "The National Information Policy in Japan," Tokyo, 1990, 9pp.

[28] Sakai, T., and Aichi, Y., "Multi-media System IMES," 1988. (A proposal from SAKAI lab at Kyoto, basis for directions at Ryukoku Un.)

[29] Tanaka, K., "Storing and Manipulating Multimedia Database Objects by Postscript and Relational Databases," *Proc. 2nd Int. Conf.on Interoperable Information*, INTAP, Tokyo, 1988.

[30] Tanaka, K., Yoshikawa, M., and Ishihara, K., "Schema Visualization in Object-Oriented Databases," *IEEE DE* 4, Los Angeles, 1988.

[31] Tanaka, K., and Yoshihara, M., "Towards Abstracting Complex Database Objects: Generalization, Reduction, and Unification of Set-type Objects," An extended abstract in *Proc. 2nd Int. Conf. Database Theory*, Bruges, S-V Lecture Notes in CS, Vol.326, 1988.

[32] Tanaka, K., Yoshikawa, M., and Ishihara, K., "Schema Design, Views, and Incomplete Information in Object-Oriented Databases," *J. of Inf. Proc*, Vol.12 No.3, 1989.

[33] Tanaka, S., Shima, M., Shibayama, J., and Maeda, D., "Retrieval method for an Image Database based on Topological Structure," *SPIE* Vol. 1153, *App. of Dig. Im. Proc. XII*, 1989.

[34] Tanaka, Y. and Imataki, T., "Intelligent Pad: A Hypermedia System Allowing Functional Compositions of Active Media Objects through Direct Manipulations," *Information Processing 89, IFIP*, North Holland 1989.

[35] Tanaka, Y., Takahashi, K., and Mozaffari, M., "Transmedia Machine," *NATO OSI Series F24, Database Machines*, Sood and Qureski(eds) Springer, 1986.

[36] Tanaka, Yuzuru "A Tool Kit System for the Synthesis and the Management of Active Media Objects," *Deductive and Object-Oriented Databases*, Kim, Nicolas, and Nishio (eds), Elsevier N-H, 1990.

[37] Tanaka, K. and Chang, T-S., "On Natural Joins in Object-Oriented Databases," *Deductive and Object-Oriented Databases*, Kim, Nicholas, and Nishio (eds), Elsevier N-H, 1990.

[38] Akio Tojo, "National Research and Development Plan on Interoperable Database Systems," *Workshop on Interoperability in Multidatabase Systems (IMS'91)*, Kyoto, Japan, Apr.1991.

[39] Wakimoto, K., Shima, M., Tanaka, S., and Maeda, D. "An Intelligent User Interface to an Image Database Using a Figure Interpretation Method," *IEEE 10-th Conf. on Pattern Recognition*, Jun.1990.

[40] Wakimoto, K., Shima, M., Tanaka, S., Shibayama, J., and Maeda, D., "Knowledge-Based Multimedia Information Retrieval Applied to Plant Diagrams," *Workshop on Interoperability in Multidatabase Systems (IMS'91)*, Kyoto, Japan, Apr.1991.

[41] Yamamoto, Y., Namioka, M., Moki, K., and Sata, K., "An Experimental Multimedia Database System: MANDRILL – its Architecture and Language," *Int. Symp. on Database Sys. for Adv. App.*, Seoul, Korea, Apr.1989.

[42] Yamazaki, M., Yobe., M., and Fujiwara, Y., "The Image-based Database on Electronic Materials using CD-ROM," *CAMSE* 1, Tokyo, Aug.1990.

[43] Yokota,K., Kawamura, M., and Tanaka, H., "Overview of R and D activities in Databases and Knowledge bases in ICOT," ICOT, 3rd Res. Lab., rcvd Mar.1991.

Multimedia Classroom of The Future

Kuan-Tsae Huang

Institute of Systems Science
National University of Singapore
Kent Ridge, Singapore 0511
and
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

1. Introduction

Why are children, even adults so fascinated with video games? Nobody fully understands how. To reach the generation of children, parents and teachers who are raised on TV, multimedia can play very important role. The marriage of the interactive media of audio/visual equipments and the brain of personal computer is giving birth to dazzling offspring that could revolutionize the classroom environment and the art of instruction.

The multimedia classroom is an effort to integrate the quality of electronic technologies and the brain of the personal computer into a seamless environment to increase the classroom productivity. Such a system provides the authoring tools and control command software to help the instructor interactively, enhance consistency, stimulate the learning process and improve student retention. Innovative teachers will develop new uses of such systems in stimulating thoughts and learning in the classroom. This paper describes the efforts of building the interactive multimedia classrooms that could enhance learning by supplying the teacher with relatively inexpensive technology.

2. Today's Classroom

Until recently, most classroom presentations were made on chalk boards, overhead projectors, and 35mm slide projectors. With TV sets, video and audio equipments, and computers dominating our lives today, many of classrooms have installed such equipments in the classroom, available to the instructor to augment the lecture presentations. The use of conceptual visuals, diagrams and even live video model presentations appeared to attract students and make them more attentive to the accompanying explanations by the instructor. It is generally agreed that such visual aids would improve the effectiveness of the instruction and achieve better learning results.

However, the logistical difficulties for the instructor to use multimedia in the classroom resulted in non-productive time and were frequently disruptive to the continuity of the learning process. The objective of the multimedia classroom project is to investigate the use of new technology to enhance the instruction process and to improve the productivity and effectiveness of the learning.

3. *Advance of Multimedia Technology*

Multimedia, or media integration involves delivery of information in a computer-based presentation that integrates two or more kinds of media including text, graphics, motion video, still video, video recognition, animation and sound. Multimedia brings the television and radio technologies to the desktop computer. But doing so can be complicated because multimedia integration requires a variety of hardware and software. Data can be stored in a variety of devices. The most important software to integrate computers with peripherals such as CD-ROMs and laser videodiscs, is the authoring software. Multimedia authoring environments let users program their computers to control and synchronize numerous external devices by using English instead of a programming language. Authoring software is the glue that binds all the elements together. Authoring tools can be used to create portions of a presentation, like sound and animation sequences, and to integrate products that were created with dedicated animation, video or audio software. Some multimedia authoring environments also contain hypermedia capabilities, or the ability to link items together in a non-linear fashion. This gives viewers of a multimedia presentation control over both the sequence and contents of the information they receive. By using pointing device to select a word or graphics anywhere on the screen, they can access more information on a particular topic.

Multimedia pose interesting new problems in data storage and retrieval. In practical systems today, most of the data has been stored in file systems, which do not provide the security, integrity, or query features of database management systems. Traditionally, database management systems deal with formatted data and are oriented towards access through transaction-oriented programs. Information explosion is coming and inevitable in this high-tech century. We have to rely heavily on the computer for information management. Multimedia object management systems in school are used to replace school documents and teaching objects with an electronic facsimile or likeness, and then digitize, process, compress, store, retrieve, display, print, and distribute the electronic information object as we have done with the current paper office operation. With the advance of multimedia research, the ability to provide compact storage and quick access of graphic data and video will determine the feasibility of many multimedia applications.

There are a various reasons for an enterprise to use a multimedia object management system. The most common one is reduction of the cost of handling papers. The other is the speed of retrieval and transmission of the objects. Another one is integration of different media. The challenge of object management can be summed up in three Rs: repositories, retention, and reference. All the schools have libraries of objects that are retained for a period of time and then purged from active files. These objects are then archived for a retention period. Before the object's life ends, it may be used or referred to by students and teachers. It may have been a piece of homework that was received, such as a trip report or an art drawing. The object itself may then have been sent to a poster board for presentation, or to a competition for evaluation by judges. Looking at this object may have required reference to historical objects, photo images and video and audio recording in a central file or repository. When processing is completed, the object requested is filed in a folder. When the student inquires about the status of the object, a search server will get from the document object management system. The heart of this storage system is a document recognition system that can take the scanned image, separate the text area, and then recognize the characters in the text area, follow by a text summary and abstract to create the index for storage and retrieval.

One key problem that remains to be solved is integration of all types of data stored in different types of storage devices and data management systems. Users cannot be expected to remember multiple different access methods and access paradigms in order to use these separate databases. Nor can the user expect all these databases to be converted to a single common model with a single access method. One common object abstraction and indexing scheme is used for all kind of objects, including a relational table. The traditional free text retrieval technique is then applied to identify related objects. A research effort is to make the user creation of queries easy by exploiting the use of images as visual feedback for query formulation and confirmation. Our search strategy is to use the probabilistic retrieval model.

4. *Vision of The Multimedia Classroom*

To transform the traditional teaching environment into a state-of-art classroom, we need to combine the interaction between the student and the teacher with sound, instructional design, and productive use of technology. The "electronic blackboard" is the focal point of the system. What make it work is a built-in computer, an instruction workstation, which is as powerful as an engineering workstation.

The machine will allow an instructor to simultaneously monitor the electronic databases for course materials or other information of related interest, to watch for the student learning pace from student participations, to activate the student interaction with specific questions and answers, and to orchestrate graphics, text, videotape, videodisc, audiotape, and lighting through a single control device at his disposal. Such a classroom environment would allow the instructor to use instructional materials effectively, create a variety of learning atmospheres interactively and coach the learning process adaptively. The emphasis of classroom instruction will change from pure knowledge transfer to skill building.

In addition, the machine will also offer the students a means to actively participate in the learning process. There will have an electronic input device at students' desks to allow them to initiate and respond to questions. Using data generated by students, the system will tabulates and displays all responses, and compares those of the current class to those of previous classes. Such a classroom environment will stimulate the active processing of data, information, ideas, viewpoints and beliefs while the learning is taking place. The differences in answers and opinions as a group can be observed and discussed immediately. Such indicators of both individual and class learning progress will encourage positive effects of self-assessment and competition among students. The student can also participate in the group discussion.

5. *What Makes The Multimedia Classroom Smart?*

In essence, the multimedia classroom is an ambitious effort to give the classroom some real smarts. The architecture of such a new classroom environment has four important features:

(A) Powerful instructor's program

The instructor's program will guide the teacher through a series of interactions that help him to create an effective and efficient instruction. The program needs to support the development efforts by systematically moving the author down through a seven-level course structure (course, module, unit, lesson, component, activity and media event). Much of the instructional design work on a course occurs at the top six levels. All the work to implement the design in the classroom is done at the media event level. Media event records are created by identifying the type of media, such as the laser disk, or the keypad, and then providing detailed information about how the event will function.

(B) Student response system

A student response system, using a keypad, keyboard or touch screen to input, would enable each student to respond to questions during the learning process and to become a more active participant in the process. The systems will poll the student response units, tabulates the results and presents graphics simultaneously to the students and instructors. There are several categories and types of questions that can be used to elicit input from students during a course: (1) Yes/No or True/False (2) multiple choice (3) mean numeric entry (4) correct numeric entry (5) rating scale from 1 to N, etc. The student can also use the keyboard to type in descriptive answers without having to worry about being seen by others. The participation in group interaction is encouraged.

(C) Audio/Visual Database

It is important to provide a seamless integration of all audio/video equipments. A simple software that can control the switch from one medium to another will make the use of multimedia much more effective. With only one projector screen, the high-resolution computer images and normal video images will be displayed interactively with high quality sound. An instructor conducts a class by advancing through the sequence of previously planned audio, visual, and student response events. A document recognition database for digitized audio and video is also designed for an integrated digital system.

(D) Instructor's panel

To help the instructor to view and change of the planned sequence, an instructor's panel that provides valuable presentation data and information as well as interactive control of presentation is critical. The instructor can see at the panel information about current presentation title, prompting data, upcoming media events. He can alter the presentation sequence and operate the media through soft keys and the stylus interactively. In addition, the stylus allows him to draw lines, figures and free-form writings in an electronic blackboard mode.

6. *The First Prototype for the Singapore 2000 Exhibition*

IBM Singapore and ISS jointly developed a prototype that integrated existing advanced technologies to demonstrate the concept of the multimedia classroom for various types of skill training. This effort was specially designed for the Singapore 2000 exhibition held from June 6 to June 23, 1990 as part of the Singapore's 25th anniversary celebration. It was one of the major exhibits held in the manpower development pavilion, where the emphasis was on the people that made Singapore's success a reality, and upon whom Singapore's global future depends.

The prototype we developed used off-the-shelf hardware. The Advanced Technology Classroom (ATC), a computer-based education and executive presentation system, was used. It aimed at substantially reducing the classroom time and improving learning significantly.

The system comprises five major components: 1) A smart lectern: this is the main center of control of the system. The unit uses a PC/AT to control the functions and the audio system. It includes the entire speech system, plus the input mixing and amplification stages of the program audio system. The top of the lectern houses a flat plasma panel display and the laser sensor unit serving as a touch panel that provides input to a computer with a special stylus with a reflective tip. It facilitates various pointing and drawing functions. 2) A student response unit: interactive response keypads are used to obtain student responses to questions posed by the instructor. 3) A presentation unit: all visual material is projected on a large, high-gain screen through a video projector or a large projection TV screen. 4) An audio/video equipment unit: this houses the multimedia assembly of audio and video equipments. These can be controlled from the smart lectern, by the hand-held remote control or the podium with the panel. 5) Authoring and command software tools: these provide a software environment to assist the instructor with course development. Its features include structured course development, multi-phase support from design to final verification of individual presentation media, format suggestions and built-in verification and cross-checking.

Software and Classware

We had produced ten lessons given at the exhibition by instructors from various organizations in Singapore with the participation of students from audiences. We have created ten lessons such as History of the Singapore River, Basic Mandarin, The Making of Micromouse, Computer Numerical Control Lathe, and Vehicle Testing, etc.

The major effort for this project is to select a set of lesson modules that relate to Singapore's manpower development. We have to work with instructors to create the well-designed course materials and prepare using the ATC technology effectively. We hope this technology will play an important role on manpower development to fulfill the Singapore's vision and aspiration to become a developed nation by the year 2000.

The video materials for the courses are pressed into laser discs after the material is well designed and created. In the design of each module, ten steps of the instructional design process are followed: (1) define the audience, (2) research the content, (3) clarify instructional objectives, (4) explore teaching strategies, (5) outline the module, (6) detail the plan, (7) conduct paper walk-through, (8) produce audio-visuals and video, (9) pressing video disc, (10) pilot test. In the course of this process, a decision has to be made on the media to use and how the instructor will use the keypads. Video production is very expensive and time consuming. Existing video material will be used if possible. Some existing softwares such as Lotus Freelance, IBM Linkway, IBM ATC authoring program, and programming tools for IBM motion video adaptor are used for this task. With the same materials, we also produce a hypermedia system using Linkway for after-classroom self learning.

7. *An Affordable Prototype to Encourage Experiments*

To address today's practical needs, we are designing two new multimedia classroom systems. One of them uses advanced ideas and off-the-shelf A/V equipments. The major objective of this project is to produce a very low-cost solution with similar functionalities and keep the quality as much as possible. We will be able to have the schools to try the system immediately and produce further requirements through experiments. The system is being developed on IBM PS/2 and PC/AT environments with 640KB of RAM, 40MB of hard disk, and a VGA card and the monitor. The VGA display adaptor can be transformed into composite video and display at a 28 inch multisystem TV or projector. The system provides interfaces to the video disk player, video tape player, as well as audio tape and compact disk player, which can be controlled using the mouse. The total hardware cost for such a system could be around S$10,000 to S$15,000 dollars.

The software is used as an aid to quickly search a particular segment of video to play, or a series of still slide frames to play from any existing video tape or videodiscs, thereby allowing teachers to reenforce their lessons with audio and video materials easily.

8. *Research Issues*

The total educational software revenue of 1989 was just $350 million, whereas the video game revenue was $3.5 million. How to build an interactive instruction delivery system that can stimulate thoughts and learning will be the key to the success of education software industry. Our research focuses on the development of softwares on CommuteWare, TeamWare, GroupWare and CompeteWare to support the training of communication, collaboration, coordination and leadership.

CommuteWare for the communication training

Communication is one of the key social skills to be trained at the school. Traditional face-to-face conversation is the predominant method of today's school. Among them, some are mediated in the classroom by the teachers and some are among the students to share their own information. Report writing and weekly diary are the asynchronous communication channels between the teacher and the student. Multimedia computers networked in schools and at homes present a wider spectrum of opportunity for the communication through electronic mail, computer conferencing, voice conversation and multimedia presentation. The multimedia classroom presents the opportunity for a much wider range of communication skill and expression training. The challenges arise from how to make a remote interaction as effective as face-to-face interaction by taking advantages of multimedia technology and a more relaxed environment.

TeamWare for Collaboration Training

Collaboration is the cornerstone of the school activity. The multimedia classroom presents the facilities for the exploration of unstructured problem solving with team collaboration. It encourages shy students to have equal opportunities in participation and to provide an equally-weighted input. Effective collaboration requires information sharing, idea brain storm, decision consensus and structured conclusion. The system needs to allow simultaneous access to the same information and automatic notification to each other's action.

GroupWare for Coordination Training

Communication and collaboration can be more effective if a group's activities are coordinated. Without that, a team of classmates can often engage in conflicting or repetitive actions. Coordination can be viewed as a subject of learning itself. A groupware can support several users in a common task providing an interface to a shared environment in which people can participate.

CompeteWare for Leadership Training

Group activity requires a leadership be generated in order to settle for a consistent direction. However, during the process, it also generates competition among the members. The competition environment provides very good peer pressure for learning. It makes students learn much faster and more effectively. A CompeteWare can help students compete in an either open or closed environment. It also motivates each individual to explore his strength and intelligence in a more relaxed manner.

9. Conclusion

Professional development is the key to our future success. The emergence of service-oriented and high-technology industries has heightened the importance of skill transfer and training of human resources. Most people seem to agree that education is the key to the future of our society. Especially in Singapore, the people is the most important resource for the country. We hope the use of technology in the classroom will increase teaching and learning effectiveness by addressing various levels of classroom interactions and student learning rates. Our effort in the project demonstrates one way of using technology in the multimedia classroom of the future. The challenge lies on training the teachers, developing the course materials and studying the impact of teaching and training when using the multimedia classroom.

Acknowledgment

I would like to thank IBM Singapore and ISS for supporting this project. The following people had contributed to this project significantly: Teh Wei-Shoong, Sue Bowles, Eugine Chang, Andy Loh, and Rodney Tan. I also wish to acknowledge the help from Singapore Broadcasting Corp., Curriculum Dev. Inst. of Singapore, NgeeAnn Polytech., Singapore Polytech., Nanyang Tech. Institute, National Univ. of Singapore, National Productivity Board, Economic Dev. Board, Vocational Inst. Training Board, and IBM AP education.

Call for Papers

I C D T 92

International Conference on Database Theory
Berlin, Germany, October 14 - 16, 1992

ICDT 92 is the successor of two series of conferences on theoretical aspects of databases that were formed in parallel by different scientific communities in Europe. The first series, known as International Conference on Database Theory, was initiated in Rome in 1986, and continued in Bruges (1988) and Paris (1990). The second series, known as Symposium on Mathematical Fundamentals of Database Systems, was initiated in Dresden in 1987, and continued in Visegrad (1989) and Rostock (1991). The merger of these conferences should enhance the exchange of ideas and cooperation within a unified Europe and between Europe and the other continents. In the future, ICDT will be organized every two years, alternating with the more practically oriented series of conferences on Extending Database Technology (EDBT). ICDT 92 is organized by Fachausschuss 2.5 of the Gesellschaft fuer Informatik, in cooperation with EATCS and ACM (pending).

Topics We solicit papers describing original ideas and new results on theoretical aspects of all forms of databases. Major themes to be covered include, but are not limited to:

data models and design theory	distributed and heterogeneous databases	database programming languages
dependencies and constraints	active database systems	concurrency control and recovery
incomplete information	parallelism in databases	complexity and optimization
deductive databases	query languages	data structures and algorithms for databases
complex objects	updates and transactions	fundamentals of security and privacy

Papers Authors are invited to submit 6 copies of a full conference paper (which might be an extended abstract where proofs are only sketched) of at most 5 000 words before February 8, 1992, to one of the **program co-chairs:**

Joachim Biskup - ICDT 92
Institut für Informatik
Universität Hildesheim
Samelsonplatz 1
D-W-3200 Hildesheim - GERMANY
Tel: +49-5121-883-730
Fax:+49-5121-860475
e_mail: biskup@infhil.uucp
(mcsun!unido!infhil!biskup)

Richard Hull - ICDT 92
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0782 - USA

Tel: +1-213-740-4523
Fax: +1-213-740-7285
e_mail:hull@cse.usc.edu

It is expected that as before the proceedings of ICDT 92 will be published by Springer-Verlag, and will be available at the conference.

Program committee

S. Abiteboul (France)	F. Afrati (Greece)	N. Bidoit (France)	J. Biskup (Germany)
J. Demetrovics (Hungary)	G. Gottlob (Austria)	R. Hull (USA)	L. Kalinichenko (USSR)
G. Lausen (Germany)	U. Lipeck (Germany)	N. Lynch (USA)	J. Naughton (USA)
J. Paredaens (Belgium)	K.-J. Räihä (Finland)	D. Sacca (Italy)	Y. Sagiv (Israel)
A. P. Stolbushkin (USSR)	L. Tanca (Italy)	L. Tenenbaum (USSR)	B. Thalheim (Germany)
M. Yannakakis (USA)	J. Zlatuska (CSFR)		

Organizing committee

G. Vossen (RWTH Aachen): local arrangements and general organization
B. Mahr (TU Berlin): local arrangements
B. Thalheim (Universität Rostock) : MFDBS-coordination

Important dates

Deadline for submission:	February 8, 1992
Notification of acceptance or rejection	May 15, 1992
Camera-ready papers due	July 1, 1992
Conference	October 14 - 16, 1992

CALL FOR PARTICIPATION

First International Conference on Parallel and Distributed Information Systems

December 4–6, 1991
Fontainebleau Hilton Resort, Miami Beach, Florida

SPONSORS

ACM SIGMOD & SIGARCH
IEEE Computer Society
Florida International University

COMMITTEES

Steering Committee:
Sushil Jajodia (Chair), George Mason U.
David DeWitt, U. of Wisconsin
Masaru Kitsuregawa, U. of Tokyo
Shamkant B. Navathe, Georgia Tech.
Napthali Rishe, Florida International U.

Conference Chair:
Amit P. Sheth, Bellcore

Program Chairs:
Hector Garcia-Molina, Princeton U.
H.V. Jagadish, AT&T Bell Laboratories

Tutorial Chair:
Arie Segev, U. of California, Berkeley

Publicity:
M. Tamer Özsu, U. Alberta/GTE Labs. (ozsu@cs.ualberta.ca)

Treasurer:
Gene Wuu, Bellcore

Local Arrangements:
Luis Cova, Florida International U.

Registration:
Mark Weiss, Florida International U. (weiss@fiu.edu)

Panel Coordinator:
Wei Sun, Florida International U.

Asian Coordinators:
Yasushi Kiyoki, U. of Tsukuba, Japan (kiyoki@is.tsukuba.ac.jp)
Kyu-Young Whang, KAIST, Korea (kywhang@sorak.kaist.ac.kra)

European Coordinator:
Erich Neuhold, IPSI-GMD, Germany (neuhold@darmstadt.gmd.de)

FOR MORE INFORMATION

Contact:

- Regional coordinators, registration chair, publicity chair; or
- IEEE/CS Conference Department, 1730 Massachusetts Ave, N.W., Washington, DC 20036; Tel: (202) 371-1013; Fax: (202) 728-0884.

SCOPE

This conference is concerned with database and other aspects of parallel and distributed systems related to data and knowledge-intensive applications. It encompasses the topics of design, development, and utilization of information systems in a distributed environment together with all aspects of application of parallel processing for information management.

TECHNICAL PROGRAM HIGHLIGHTS

The conference program includes 29 technical papers covering topics such as shared nothing and shared memory systems, disk replication and arrays, issues of operating systems/optimization/file management in the context of parallel and distributed systems, parallel query and logic processing and text retrieval, transaction processing in federated databases, and parallelism for object-oriented systems. It also includes two keynote addresses and two panels focusing on real life data/knowledge management issues and promising technological advances, two tutorials by internationally known experts, and several project synopses reporting on work-in-progress on prototype implementations.

Tutorial 1: **Distributed Transaction Processing Monitors** by P.A. Bernstein, Digital Equipment Corporation.

This tutorial provides a technical overview of transaction processing (TP) monitors, which are software products that integrate applications, databases, and communications. Features of TP monitors (e.g., message management, process management, and transactional communications) and emerging standards are discussed. Examples are drawn from a variety of commercial TP monitor products.

Tutorial 2: **Architectures for High-End Transaction and Query Processing** by A. Thomasian, IBM T.J. Watson Research Center.

This tutorial discusses trends and issues in high-performance transaction and query processing, classifies architectures as data-sharing, data-partitioning, and hybrid systems (e.g., client/server architectures) and provides an overview of some existing implementations as well as alternatives presented in the literature. A discussion of benchmarks is also included.

CONFERENCE VENUE AND HOTEL REGISTRATION

The conference will be held at the Fontainebleau Hilton Resort and Spa, which is the most famous resort on Miami Beach with its award winning restaurants, tennis courts, water sport facilities and health spa. It is situated on 18 acres of Atlantic Ocean beachfront, approximately 20 minutes from the Miami International Airport. A special rate of US $110 + 11% tax per room (single or double occupancy) has been negotiated. To make reservations, contact the hotel directly before November 13. Outside Florida: 1-800-HILTONS, in Florida: 1-800-548-8886, outside US: 305-538-2000, fax: 305-534-7821, telex: 51-9362.

SOCIAL PROGRAM HIGHLIGHTS

In addition to the opening reception and the banquet, the social program includes an optional dinner and show at the famous Club Tropigala at la Ronde, which provides a unique combination of fine dining and award-winning floor show in an exotic Caribbean setting, and a full day cruise to Bimini, Bahamas on the Saturday following the Conference.

CONFERENCE REGISTRATION

Please circle your choice. Note that the early registration deadline is November 8.

		Members	Nonmembers	Students
Conference only	(Early)	$260	$325	$60
	(Late)	$315	$395	$75
Single Tutorial	(Early)	$140	$175	$60
(Which one:_______)	(Late)	$170	$205	$75
Both Tutorials	(Early)	$225	$280	$95
	(Late)	$270	$325	$120
Club Tropigala dinner and show		$39	Bimini cruise	$75

Name: ________________
Affiliation: ________________
Address: ________________
____
____
Tel: ________________
Email: ________________
Membership Society/No: ________________

TOTAL ENCLOSED $________ ☐ Check enclosed ☐ Credit Card; Type (Amex, Visa, MasterCard): ________

Card No: ________________ Expiration Date: ________ Signature: ________________

Return to **Dr. Mark Weiss, Florida International University, School of Computer Science, Miami, FL 33199, USA.**

CALL FOR PAPERS

2nd International Workshop on Research Issues on Data Engineering: Transaction and Query Processing (RIDE-TQP)

Feb. 2-3, 1992
Mission Palms, Arizona

General Chairperson:
Clement Yu,
Dept. of EECS
U. of Illinois at Chicago
Chicago, IL 60680
yu@uicbert.eecs.uic.edu

Program Chairman:
Philip S. Yu, IBM

Program Committee:
(Preliminary)
D. Bitton (USA)
M. Carey (USA)
F. Golshani (USA)
S. Jajodia (USA)
Y. Kambayashi (Japan)
M. Kitsuregawa (Japan)
W. Kohler (USA)
J. McPherson (USA)
E. Neuhold (Germany)
C. Pu (USA)
M. Rusinkiewicz (USA)
L. Russell (USA)
P. Scheuermann (USA)
M. Shan (USA)
A. Sheth (USA)
S. Son (USA)
C. Tomlinson (USA)
J. Wolf (USA)
O. Wolfson (USA)

Steering Committee:
A. Elmagarmid, Purdue U.
J. Urban, Arizona State U.

Finance:
J. Bell

Local Arrangements:
J. J. McCarthy

Publication:
K.-L. Wu, IBM

Publicity:
W. Sun

This is the second in a series of annual workshops to be held in conjunction with the IEEE CS International Conferences on Data Engineering. The RIDE-TQP Workshop is intended to bring together leading architects, developers, practitioners, and researchers on transaction and query processing. Current trend in research and marketplace aims not only to enhance the transaction or query processing performance, but also to support both transaction and query processings concurrently on the same system. This trend is propelled by recent advances in computer architecture, basic software, algorithms, and performance evaluation. We are particularly interested in the current state-of-art and future developments that integrate these advanced techniques. Suggested topics include but are not limited to:

- TQP architectures, multiprocessors, large main memory, RAID, datacycle;
- TQP software, database systems, OS support, distributed and heterogeneous databases;
- TQP algorithms, fault tolerance and recovery, concurrency control, query optimization, intra-query parallelism, storage management;
- TQP performance modeling and analysis, workload characterization, benchmarks, measurements.

Paper Submission:

Five (5) copies of an extended abstract (no more than 7 pages) should be mailed before July 15, 1991, to the program chairman:

Philip S. Yu
IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598
psyu@ibm.com

The selection will be based on originality and contribution to the field. A limited number of papers will be accepted for poster sessions. A proceedings will be published by IEEE CS.

Important Date:

Deadline for submission: July 15, 1991
Notification of acceptance: Oct. 21, 1991
Final paper due: Nov. 27, 1991
Conference: Feb. 2-3, 1992

Advance Announcement

DATA ENGINEERING

The Eighth International Conference on Data Engineering

February 3-7, 1992

Tempe Sheraton Mission Palms, Tempe, Arizona

Sponsored by IEEE Computer Society

With corporate support from Bull Worldwide Information Systems

General Chairperson: Nick Cercone, Simon Fraser University, Canada

Program Chairperson: Forouzan Golshani, Arizona State University, USA

SCOPE

Data Engineering is concerned with the semantics and structuring of data in information systems design, development management and use. It encompuses not only the traditional applications but the emerging technologies and issues. The purpose of the confernece is to provide a forum for the sharing of practical experiences and research advances from an engineering point of view among those interested in automated data and knowledge management. The eighth ICDE will emphasize the issue of Technology Transfer.

PHOENIX IN FEBRUARY

February is considered to be the best time for Phoenix. With clear sunshine and daytime temperatures averaging around 70 degrees, it provides an ideal alternative to the freezing climate of east coast and rainy days of the mid-western and other western areas. The site of the conference is the town of Tempe - a blend of rugged Old West with a sophisticated New West atmosphere. Attendees can travel within a few hours from deserts and canyons to the largest forest of Ponderosa Pine in the world. The Grand Canyon, the Petrified Forest, The Navajo and Hopi Indian Reservations, Lake Havasu and the London bridge are all within driving distance of Tempe.

SPECIAL ACTIVITIES

Data Engineering welcomes spouses as well as other participants to join in a special series of activities designed to enhance your visit to the Valley of the Sun.

Event 1: A half-day guided tour of landmarks and historic sites in the Valley of the Sun. Tue., Feb.3, 12:30 p.m.-4:30 p.m.

Event 2: Visit the Heard Museum, Pueblo Grande ruins and Desert Botanical Garden. Wed. Feb.4, 9:00 a.m.-4:30 p.m.

Event 3: Enjoy a BBQ steak and chicken dinner at a "cowboy" restaurant. Thur., Feb.6, 5:30 p.m.-9:30 p.m.

Event 4: A one-day tour through scenic Arizona mountains to visit the Grand Canyon. Sat., Feb.8, 7:30 a.m.-9:00 p.m.

The Tutorial Titles and Schedules:

- **No. 1: Spatial Database Systems**
  Oliver Günther, University of Ulm
  Monday, February 3, 8:30 a.m. - 12 noon
- **No. 2: Multimedia Information Systems**
  Bruce Berra, Syracuse University
  Monday, February 3, 1:30 p.m. - 5:00 p.m.
- **No. 3: Information Security In Computing Systems**
  Harold Podell, U.S. Government
  Monday, February 3, 8:30 a.m. - 5:00 p.m.
- **No. 4: Distributed Data Management**
  M. Tamer Öszu, University of Alberta
  Monday, February 3, 8:30 a.m. - 5:00 p.m.
- **No. 5: Active, Real-Time and Heterogeneous DB Systems**
  Sham Navathe, Georgia Institute of Technology and
  Sharma Chakravarthy, University of Florida
  Friday, February 7, 8:30 a.m. - 5:00 p.m.
- **No. 6: Parallel Query Processing In Database Systyems**
  T. Patrick Martin, Queens University and
  Bernhard Seeger, University of Munich
  Friday, February 7, 8:30 a.m. - 5:00 p.m.

The conference general sessions will be held on Tuesday, Wednesday and Thursday. Papers and panels relating to the following areas will be presented and discussed:

- AI and Knowledge Based Systems
- Applications and Application Systems
- Benchmarks and Performance Evaluation
- Data Engineering Tools and Techniques
- Database Design and Modeling
- Database Management and Structure
- Deductive and Extensive Databases
- Multimedia Database Systems
- Distributed Database Systems
- Integrity and Security Techniques
- Learning and Discovery in Databases
- Object-Oriented Database Systems
- Query Languages and Processing
- Scientific Databases
- Heterogeneous Systems and Interoperability

For a full Data Engineering advance program, please contact:

Oris D. Friesen
Bull Information Systems
P.O. Box 8000, MS H32
Phoenix, AZ 85066-8000, USA

Telephone: +1-602-862-5200
FAX: +1-602-862-6105
Email: friesen@system-m.phx.bull.com

Conference Registration • The Eighth International Conference on Data Engineering

Please mail or fax this registration form and payment to:
Bev Graff
Bull Information Systems
P.O. Box 8000, AZ05-B40
Phoenix, Arizona 85066-8000
USA
Phone: +1-602-862-4255; Fax: +1-602-862-4288

Conference and Tutorial Registration Fees:

Advance Registration — Until January 10, 1992

	Member	Non-Member	Student
Conference Only	$290	$365	$90
One Full-Day Tutorial	$220	$275	*
One Half-Day Tutorial	$140	$175	*

Late/On-site Registration — After January 10, 1992

	Member	Non-Member	Student
Conference Only	$350	$440	$100
One Full-Day Tutorial	$265	$330	*
One Half-Day Tutorial	$170	$215	*

* NO STUDENT DISCOUNT

Please check the tutorial(s) you wish to attend:

__ Tutorial 1: Spatial Database Systems **(Half-Day)**

__ Tutorial 2: Multimedia Information Systems **(Half-Day)**

__ Tutorial 3: Information Security in Computing Systems **(Full-Day)**

__ Tutorial 4: Distributed Data Management: State-of-the-Art, Unsolved Problems and New Issues **(Full-Day)**

__ Tutorial 5: Active, Real-Time and Heterogeneous DB Systems **(Full-Day)**

__ Tutorial 6: Parallel Query Processing in Database Systems **(Full-Day)**

Special Activities Fees:

Event	Number of Participants	Fee until January 10	Fee after January 10
Event 1: The Valley of the Sun	______	@ $25 = ______	@ $30 = ______
Event 2: Land & People of the Southwest	______	@ $40 = ______	@ $45 = ______
Event 3: Cowboy Cook-out	______	@ $40 = ______	@ $45 = ______
Event 4: The Grand Canyon	______	@ $65 = ______	@ $75 = ______
		Total =$______	Total =$______

Please Type or Print

Last/Family Name ______________________ First ______________________ Middle ______

Affiliation ______________________ IEEE CSMembership Number ______________________

Mailing Address ______________________ Mail Stop ______________________

City/State/Zip/Country ______________________

Work Phone ______________ Telex/Fax Number ______________ Email ______________

Conference Fee:$__________ **(a)**

Tutorial Fee:$__________ **(b)**

Special Activities:$__________ **(c)**

TOTAL = (a + b + c): $______________

Method of Payment

___ Check or Money Order ___ Credit Card

____ MasterCard ____ Visa ____ AmEx ____ JCB ____ Discover ____ Carte Blanche

Cardholder Name ______________________ Signature ______________________

Card Number ______________________ Expiration Date ______________________

Registration is not complete without remittance of registration fee. All payments must be made in US dollars.

NOTE: Conference registration fee includes admission to the technical sessions, one copy ofthe proceedings, break refreshments and reception. Refunds will be made if requested in writing no later than January 10, 1992.

IEEE

Worldwide Information Systems

Bull

IEEE Computer Society
1730 Massachusetts Avenue, NW
Washington, DC 20036-1903

Non-profit Org.
U.S. Postage
PAID
Silver Spring, MD
Permit 1398

Mr. David B. Lomet
9 Cherry Lane
Westford, MA 01886
USA

Produced by the IEEE Computer Society Data Engineering Technical Committee