

A Layered Approach for Computerization of a Multimedia Archive

Daryl H. Hepting

CMPT TR 93-14
School of Computing Science
Simon Fraser University
Burnaby, British Columbia
Canada V5A 1S6

December, 1993

Abstract

Merce Cunningham is widely regarded as the father of modern dance. His work, which began in the mid-1930's, is a treasure from which all the world can benefit. Much of his work, captured in several different media, is preserved in an archive maintained by the Cunningham Dance Foundation. This archive has always been a tool with which Mr. Cunningham has furthered his art. In order for it to effectively support the work of other artists, its vast collection of information must be made accessible to a wide audience, in meaningful ways. The challenge of the computerization project is to apply technology to provide that increased access, without compromising artistic integrity.

1 Introduction

With the advent of affordable digital multimedia computer technology described by Fox [10], many new applications which were only recently unthinkable to most are being developed. Much of the new technology provides

improved access to more and better information. The possibilities for this new technology are great and will involve almost every aspect of daily life.

Consider an analogy of a library with a physical card catalogue system but the goal of unrestricted information access through a computer. Computerization of the library's card catalogue system represents an important first step towards that goal. The next step, which the computer industry will soon be ready to take, is the computerization of the library's contents. Both parts represent improvements in accessibility to the library's resources. The exact means to fully achieve the stated access goal are not known, but research into hypermedia systems [31] has already provided important lessons. Irven *et al.* [15] provide an excellent example, excerpted below, from their experience with the *Telesophy* system:

A family must relocate 2,000 miles to the New York area. They have not decided whether to live in Manhattan, New Jersey, Westchester, Long Island or Connecticut. Three hundred communities are potentially within reach of their new job locations. Before they leave their home town, they access, through information terminals, real estate services in each of the areas they are considering, which show pictures and other information of houses currently for sale. Using another service, offered by a separate provider, the couple research specific information relating to the specific communities ... Having browsed through this information, they select several towns meeting their unique needs. They then carefully look again through the color pictures and other information on houses for sale; finally, placing houses of interest in a "folder" which the real estate company will use to set up their appointments when they visit the New York area during the next several days.

In conjunction with the developments in hardware technology described by Fox [10], database technology must also develop in order to support these new applications. Cattell [5] mentions briefly how database systems are evolving in order to meet the challenges beyond the current standard of database applications used, for example, in banks. Similar functionality is required by all of the new database applications and much of it seems to mesh well with ideas taken from object-oriented systems. It is possible to

characterize the opposing positions taken in the database community by whether new database features are added based on a truly object-oriented approach or on a more traditional value-oriented approach. Aside from the database problems mentioned, research efforts are now being directed to making the notion of object pervasive in the computer system (IBM's System Object Model [8] is an example). As an illustration of the difference between old and new database systems, Silberschatz *et al.* [29] explains whereas once transactions were assumed to be on the order of fractions of a second, they may now be on the order of days. The amount of data which the systems will be asked to manage will also increase dramatically, as they take on data elements potentially several million bytes long.

As computerized information services become more prevalent, another question aside from that of scale becomes evident: how will the presentation of information be handled in these systems? Indeed, these new applications will encourage far greater exploration of the computer as a creative tool. Existing multimedia systems have engaged the services of designers to ensure that the information which the system has to offer is not lost behind the clutter of poorly designed screens. Computers have now developed sufficient capabilities to be considered as tools for other forms of artistic expression, such as painting [13].

Within the space of these technological advances, not only do current computer applications find room to expand their scale, but new areas have taken on the challenge of computerization. It is within this framework that this paper seeks to explore the issues related to the design of a multimedia information system for the archive of renowned choreographer Merce Cunningham. Much of the current experience in the development of multimedia databases comes from the areas of design and office information systems. The lessons of those projects, along with an examination of other novel technologies will guide the development of this project.

The challenges associated with the development of a computerized multimedia archive are quite new. The main issues are all related to data management:

- how can the complex relationships in the archive be modelled?
- what features in a database are required for adequate support of multimedia data?

- how can the information in the archive be accessed effectively?

The archive of American painter Andrew Wyeth is in the process of being computerized in a system developed by IBM [23]. An integral part of that system is its ability to capture the original colours of the works and display them faithfully, thereby maintaining the integrity of the art. Although Mr. Cunningham's archive is more varied by nature, this artistic element is also key.

This paper proposes a model for the computerization of the Cunningham archive, specifically in terms of its data management requirements. It focuses on the larger development issues rather than the specifics of a particular implementation. The rest of the paper is organized as follows. Section 2 of the paper describes the Cunningham archive project in terms of its goals. Broad categories for the information stored in the archive are described briefly. Based on those definitions, some of the relationships amongst the elements are discussed. Section 3 discusses the reasoning behind the selection of a data model to represent the relationships in the archive, with reference to pieces of a proposed schema. Section 4 discusses the database requirements for support of multimedia applications. The process by which digital multimedia data is created is also briefly discussed. Section 5 discusses possible interfaces to the computerized archive, for which hypermedia presents an attractive alternative. Section 6 presents a possible application of the archive and directions for further exploration. Finally, Section 7 contains some concluding remarks.

2 The Cunningham Archive Project

At work since the 1930's, Merce Cunningham is known as the father of modern dance. Throughout its existence, the Cunningham Dance Foundation (CDF) has sought to maintain the archive of Merce Cunningham as a dynamic entity, rather than as a museum exhibit.

Due to the notoriety of Mr. Cunningham and the value of the pieces in the archive, it was imperative that a unified agenda for preservation was put into place [1]. The project has a grand scale because of the breadth of Mr. Cunningham's art and the length of time during which he has pursued that art. The following is a complete list of the different types of elements in the archive:

- Print: performance programs dating back to 1938, press clippings from the late 1930's, posters and flyers from the early 1940's, photographs and colour slides, books and periodicals which include writings by Merce Cunningham or John Cage (a composer with whom Cunningham often collaborated), or features about the Merce Cunningham Dance Company, and administrative records.
- Production: a limited number of costume and set designs and a descriptive and illustrated inventory of items not physically in the archive.
- sound: recordings and scores of works no longer in the repertory, recordings of oral history, radio interviews, lectures, and symposia.
- Film and video: original camera (and sound) recordings, outtakes, rough cuts and post production material from film/video collaborations, recordings of works choreographed for the camera, recorded or adapted performances, documentaries, educational programs, master films and videotapes for duplication, rehearsal and performance recordings, interviews, documentaries, and newscasts featuring Mr. Cunningham.
- Miscellaneous: certificates, trophies, awards won by Mr. Cunningham, and other biographical effects.

With the recent construction of approximately 300 cubic meters of storage rack space, the CDF has begun to claim physical control of the archive. Michael Bloom, a director of the CDF, feels that “intellectual control” will only come as the result of efforts at computerization [3].

Bloom has described the purpose of the archive as both “regenerative and redistributive.” Regenerative because Mr. Cunningham uses the archive for his current artistic endeavours. Redistributive because the treasures in the archive are shared with scholars, critics, editors, publicists, and exhibition curators. Preservation is key to both aspects, but in this sense it means not only protecting the artifacts from destruction, but allowing Mr. Cunningham to continue his art into the future. If these two functions are performed, then the artistic integrity of the archive has been maintained, at least to a large extent.

The preservation plan involves several different agendas for information, media, choreography, music, education, and legal matters. Although the

information and media agendas might seem similar, they represent different, yet equally important, needs of the logical (information) and physical (media) aspects of the archive. All of these agendas represent appropriate concerns in developing a computerized archive containing multimedia data:

- Choreography is the key to the archive. Much of Mr. Cunningham’s choreographic notes are maintained separately from the main archive, yet rehearsal videos and computer-choreographed material provide accessible, direct examples of this aspect.
- The music associated with Mr. Cunningham’s work has always been innovative, whether it be the work of John Cage, John Tudor or one several others. The number of digital recordings of these pieces is increasing, and they make an important addition to the contents of the archive.
- Education as a function of the Cunningham Dance Foundation will only be enhanced by the implementation of a computerized information resource. In addition to the historical value of the archive, its elements could be combined to form course notes about elements of the Cunningham technique, complete with moving picture ¹ sequences.
- The CDF struggles to explore a structure within which the work of Merce Cunningham will be protected, but the computerization project brings several new issues to the fore. When dealing with digital media, there are many issues of property which remain largely untried in today’s courts [25].

The preceding statements serve both to highlight the great diversity present in this archive and propose an organization of its material along several different lines. However, there are many other relationships which are more or less subtle than the ones already mentioned. For example, a reasonable organization could characterize each dance as consisting of choreography, music, and design elements. However, Robert Rauschenberg (former company designer) has described [17] the relationship of those elements (choreography, music and design) in a completely different fashion: “it was

¹this term is intended as a generalization for the source of the sequence, whether it be film, video or animation.

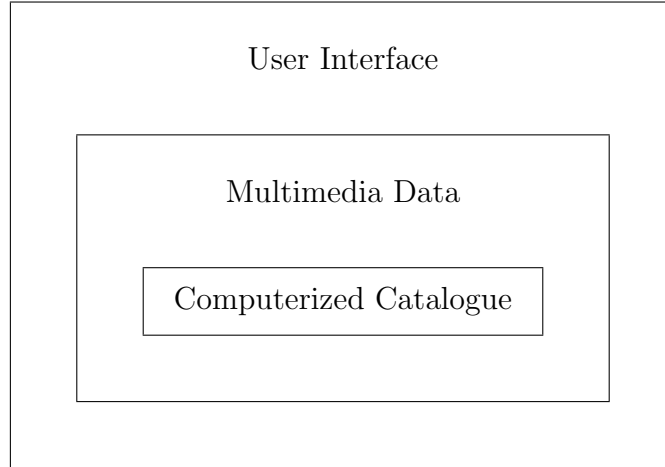


Figure 1: The different layers involved in the computerization project: computerization of the catalogue is layer 1, addition of multimedia data is layer 2 and user interface development is layer 3.

the most excruciating collaboration but it was the most exciting and most real, because nobody knew what anybody else was doing until it was too late.” In fact, the ability to capture the intent of these complex relationships within the archive database is key, and the data model selected must have sufficient expressive power to accomplish it.

Once a data model is selected, a computerized catalogue of the archive can be created. This will represent an important step towards gaining intellectual control of the archive, since the computerized catalogue will be much more than merely an analogue to its physical counterpart. It will also be the core of the proposed computerization project. Figure 1 presents an overview of the project in terms of these layers.

With the catalogue in place, the CDF will be able to concentrate on the conversion of the archive elements to a digital format. The specific format to be used must be considered carefully, since this process will span some years, and it will not be practical to do a second time. The digitization of the archive elements and their inclusion in the computerized database represents the second layer of the project.

The third and final layer is concerned with providing an effective means of public access to the archive. Hypermedia-based methods seem appropriate, at least as a point of departure. There is a great deal of ongoing research in this field, especially with regards to the potential combination of hypermedia systems with databases.

The layers which comprise the computerized archive project are conceived in a pragmatic fashion. Work towards the goals of the preservation agenda must begin immediately, yet the entire scope of the computerization may take several years to realize and undergo many stages of evolution in the interim. Along with these requirements comes the necessity of a database system which can be adapted easily as circumstances change.

3 Choosing a Data Model

Silberschatz *et al.* [29] cite the need, amongst others, for the support of complex objects in the next generation of database systems. With reference to multimedia systems specifically, Woelk *et al.* [35] consider the need for a data model which also provides the means to express complex relations amongst objects.

Joseph *et al.* [16] cite the following reasons why current, conventional database management systems cannot support the new applications:

- lack of expressive data modelling capabilities,
- the “impedance mismatch” between programming languages and database systems,
- inadequate interactive performance,
- lack of support mechanisms for long transactions, and
- lack of support for schema evolution and version management.

Of particular interest to this discussion will be the data modelling capabilities, the support for long transactions and the support for schema evolution in the “new” database management systems, although all are important. This discussion is limited to comparisons between extended relational and object-oriented database management systems (DBMS's). These are not

the only two possibilities being researched, but they serve to illustrate the division of the database community by the approach. These competing approaches may be classified based on support for either an object-oriented or a value-oriented data model. Following Joseph *et al.* [16], all relational, extended relational, and logic database systems can be grouped into the value-oriented category. This approach is typified by the fact that relationships between objects are not stored explicitly; rather they are inferred by comparing values of attributes. Although extended relational systems do not require first normal form, and can store nested relations, they do not allow sharing of relations. A positive aspect of these extended relational systems is that they do represent an incremental extension to the already familiar technology of relational databases – which may make conversion between relational and extended relational models conceptually easier.

Some of the practical difficulties with value-oriented data models are seen in the Human Genome Initiative (HGI) [11], a very data intensive project which seeks to map out the human genome. Lack of object-identity concepts provide problems when manipulating data elements and attempting to express relationships in terms of order. Although many scientists involved with the HGI project would seem to prefer the object-oriented model, they worry about the current lack of efficient query-based access mechanisms available to them.

Another view of the problem with the value-oriented approach comes from the *Intermedia* project [32], exploring hypermedia systems issues at Brown University. In this case, many difficulties were encountered when mapping between a rich programming language data model and a restrictive database model – a prime example of the “impedance mismatch” problem. Although this comparison was made against a relational DBMS, many of the arguments still hold.

Figure 2 illustrates how the notion of a “dance” might be converted into a database schema. Note the different semantics implied by each organization. Given the earlier criterion, expressed in the quote from Rauschenberg, the more expressive object-oriented data model should be selected for the archive application. Although the expressive power of the data model is the most important selection criteria, the object-oriented model will serve well in the other two key issues indicated above. Namely, an object-oriented DBMS will support schema evolution and long duration transactions. Schema evolution will become increasingly important as multimedia data is considered, since

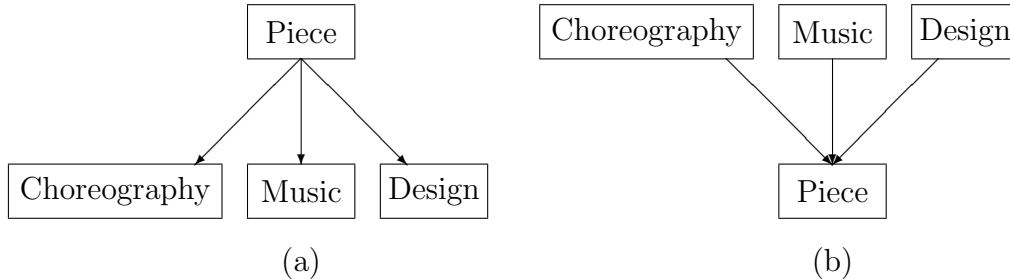


Figure 2: Different approaches to the problem of schema organization. Both (a) and (b) represent possible approaches for organizing information relating to a “Piece” (a dance, for example). Note that in (a), “Choreography,” “Music,” and “Design” elements are represented as parts of the “Piece.” In (b), “Piece” is represented as taking parts from each of the “Choreography,” “Music,” and “Design” components. The restriction that the “Music” belong to a particular “Piece” is lifted. This more expressive representation (following more closely the words of Rauschenberg) cannot be modelled using an extended relational DBMS.

neither the archive nor the relationships between elements in the archive are completely fixed or static and it is completely reasonable to expect changes as the computerization project unfolds. Unlike traditional database applications where schema modifications were rare, it seems to represent an integral portion of the new, data-intensive applications. A rather complete discussion of the issues involved in schema evolution is provided by Banerjee *et al.* [7], for example.

4 Providing Support for Multimedia Data

Multimedia systems will create an enormous amount of digital data with which the DBMS must deal. Aside from the sheer size of the database, multimedia systems will require the DBMS to handle real-time constraints in the management of transactions on the database. In terms of the Cunningham archive project, the process by which this digital format data will be generated is still unresolved. Certainly there are many issues at stake.

Both the data conversion and data management aspects of the project will be discussed in this section.

4.1 Data Conversion

The contents of the Cunningham archive are incredibly diverse. Film and video samples exist in a myriad of formats, with varying degrees of stability. This effect of the age of the archived materials is most obvious in pieces which of film, video, and audio, which were made at a time when the technologies were still new, meaning that the physical media on which they are stored is especially fragile. This also holds for printed material, although to a lesser degree. Transcription of archive items to digital form will allow the original stock to remain under ideal storage conditions while the digital copies provide masters for further reproduction. To ensure maximum flexibility, results of the digitization process should be stored in a way which maintains fidelity to the original source. Industry standards for data formats are still bound to fluctuate, so the most prudent strategy might be to maintain the digital data in some neutral format from which the actual format used by the current system could be derived. This would, however, contribute considerably to the storage requirements of the database, since an extra, larger copy of each digitized element would be required. A related issue is the resolution at which archive elements should be digitized. Transcription with maximum fidelity may not be reasonable, especially without access to unlimited storage capacity.

Rights to reproduction must also be considered at this stage, since there are items in the current physical archive for which the CDF does not own reproduction rights. Once a single digital copy of an item is made, control of that item is lost to a large extent. This becomes a very real issue as remote access to the archive is considered, and it may serve to limit the contents of the electronic archive to be significantly different from the physical one. Samuelson [25] discusses some of the pitfalls of digital media in relation to questions of copyright and ownership.

Each type of source media requires selection of a storage format. The wise selection of those data formats will mean an easier progression as the system grows and evolves. As many elements of the archive are of historical interest and value, it will be important to store information about the form and the content (when the two differ, as in the case of press clippings) of each element.

In the case of the elements with textual components, their conversion must be done in two steps of scanning: one to create a visual likeness to the element and one to capture its contents. The content portion of textual items would most likely be included first into the computerized archive, since it does not require specialized display devices. Maintaining the structure of documents at this level in the archive is important. Standards like SGML (Standard Generalized Markup Language) [6] and ODA (Office Document Architecture) [4] have been developed to describe this document structure. Hypermedia aspects of the system might also benefit from encoding in a specification language like HyTime [24]. As hypermedia is considered, the general problem of segmentation must be addressed. In essence, it represents the decomposition of long elements like film, video, or lengthy text into forms which can be indexed and more easily utilized. The process of segmenting the information, forming categories of information and creating indices for those categories is described by Smoliar [33].

4.2 Data Management

In one of the first papers to consider multimedia database systems, Woelk *et al.* [35] stated that, aside from data modelling capabilities, mechanisms which allow the sharing and manipulation of data were most important.

Given the wide variety of requirements which multimedia systems impose, the ORION database management system [19] attempted to deal with three general requirement areas:

- extensibility: by modifying or adding new types of devices and protocols which support multimedia information.
- flexibility: in controlling the capture, presentation, and transfer of multimedia information in the application environment.
- efficiency: by limiting the number of copies of multimedia objects and the number of transfers associated with them. Woelk *et al.* [35] discuss a scheme for limiting copies of multimedia data items. Christodoulakis [9] describes techniques for limiting the amount of data transferred, by allowing the system to make a simple initial test before accessing the whole item. Other strategies are possible, but will not be presented here.

Although not specifically a requirement of multimedia database systems, management of long-duration transactions [19] is a definite concern. In data-intensive transactions in the multimedia system, the size of each data file presents two, non-exclusive difficulties:

- when the file is the result of a query, it is important to provide some of the results as soon as possible [9].
- when the file contains time-dependent data, there are associated real-time constraints which require that the whole file be retrieved (and presented) on a very strict schedule [12].

It is also possible to consider computational-intensive long transactions for the Cunningham archive. Mr. Cunningham has used the *LifeForms* program from Simon Fraser University [27] to create some of his recent choreography. Ultimately, the archive would include the option of invoking the *LifeForms* program to replay some of these computerized choreographic notes. Since the *LifeForms* program maintains the choreography in a compact notation, processing of that notation to generate an animated sequence demonstrates the property of database amplification as described by Smith [30]. In this case, playback would be dependent on the processor, rather than the data retrieval.

The support for a variety of presentation and capture devices, which are required to keep the system “open,” necessitate some facility for schema evolution [19, 28]. The direct support for multimedia capture devices in the Cunningham archive system highlights another possible interpretation of the “living archive,” as new events could be captured, annotated, and stored directly into the archive.

Additional requirements are imposed by remote access. One may attempt to provide remote access at a variety of levels: to the archive catalogue, to the time-independent database, or to the full multimedia database. The *Telesophy* system [26] was conceived to provide knowledge at a distance (as the name literally implies). Indeed, the combination of information systems and ISDN [22] technology will provide that in a very real sense. As in the case for workstation based systems, the synchronization of media provides greater challenges when it is distributed. Little and Ghafoor [20] discuss some of the issues involved and present a technique based on the idea doing

as much processing as possible at the data source in order to reduce network traffic.

5 Providing an Interface to the Archive

Despite the discussions of data models and basic support for multimedia data, these issues are of little concern to the artists who will use the system most. To determine whether the criteria of artistic integrity is being maintained, one must return to the goals of the archive. In order for the archive to fulfill its role of being regenerative and redistributive, it must facilitate the artistic process. While a database system is a necessary condition for the success of this project, it is not sufficient and the project will ultimately fail if access to the archive is not provided in a facilitating manner.

An interface which is truly appropriate for the artistic integrity constraints may be some time in coming. However, exploration of the computer as a creative instrument by means of electronic publishing will help to move the whole process along the path towards better interfaces. One novel approach taken by Bier and Goodisman [2] is that documents *are* interfaces and interfaces *are* documents.

Window-based systems have gained considerable popularity as a means to manage the visual complexity on a display of limited size. As sound becomes more prevalent in applications, methods for managing the auditory complexity of workstation presentations (similar to visual windowing systems) are being developed [21].

There are many factors which will contribute to the interface of the future. However, the current technology related to hypermedia systems seems most appropriate. Several prototype hypermedia systems have been developed over the past several years, and a partial list includes NoteCards [14], Intermedia [36], and Telesophy [26]. Although hypermedia in general represents an advanced capability, Halasz [14] points out seven limitations:

- search and query: hypermedia systems provide navigational access, but query-based access is often the preferred method. For full functionality, users should be able to specify queries on both content (the contents at a node) and structure (how a node is linked to others, for example).
- composite items: a means to deal with groups of nodes as single entities.

- virtual structures for dealing with changing information: since the data model available in most hypermedia systems is static, users of those systems are forced to subject the data to “premature organization.”
- computation in networks: rather than allow the hypermedia network to remain passive, assign computations to the nodes.
- versioning: users should have the ability to maintain and manipulate a history of changes to the hypermedia network.
- support for collaborative work: creation of annotations for networks and multiple organizations of single sets of data are examples of the capabilities which would allow for productive collaborative work.
- extensibility and tailorability: casual users should have the capability to make simple changes in the hypermedia network structure, without programmer intervention.

Not coincidentally, this list shares much in common with the desired features for the next generation of database systems. The parallels with object-oriented data model issues [7] are apparent.

Also not surprising is the movement in research to add structure to hypermedia systems in order to create hypermedia databases. As seen earlier in the competition between object-oriented and value-oriented data models for the next generation of database applications, there seems to be a dichotomy in approaches towards development of these hypermedia databases. Tompa [34], for example, tries to add database functionality by extending the hypermedia data model. The alternative approach, to extend (object-oriented) database technology, is taken in the *Harmony* system [28].

In the *Harmony* system, nodes are objects and the links between nodes are given extra power to include temporal constraints and handle messages. A product of this approach means that an event occurring within a segment of moving pictures (at a particular time) can be used as the source of a link. If this link provides some audio annotation for the event, it will be played when the event occurs, and the link is navigated. In order to gain a uniform interface, each object at the interface level (or layer) is treated as a media object. The notion of layer here corresponds to descriptions of links between the generic media objects. The media object interface is reminiscent of token

objects described by Woelk *et al.* [35]. Layers, similar to Intermedia [36] webs, support the idea of collaborative work by allowing annotations to exist separate from the actual information in the nodes or objects. They could also correspond to a view on a database, depending on the context.

6 A Model for Use of the Archive

The paper by Irven *et al.* [15] describes three levels at which a user might wish to interact with an information system. These levels are now described in the context of the Cunningham archive. The following is an example of a query which a user might present to the system and how the system might handle it.

- query stage: a query is initiated in order to define the topic and to limit the scope of the search. For example, from an initial display of a page of choreographic notes, the user selects a particular sequence and requests all moving image segments associated with it.

The user might initiate this query by clicking on the area of the display showing the desired sequence. This would require that the image be segmented according to the sequences in the dance and that “buttons” are associated with each sequence on the screen. In this case, the button associated with the image segment could include the query language specification to be processed by the database manager. One model for queries in object-oriented database systems is presented by Kim [18]. Additionally, the user’s action might initiate a dialog to further restrict the query.

- browsing stage: upon return from the query, the user might see a list of the moving picture segments which illustrate the chosen choreographic sequence. The user could then examine textual summaries of the moving picture clips in order to determine which ones might be of particular interest. This textual summary could include information such as the names of the dancers shown in the clip and aspects of the recording. Once some clips of potential interest are found and selected, the associated links could be navigated (and played). Those clips which are ultimately desired could be saved, for inclusion in new multimedia archive elements.

- follow-up stage: new archive elements could be created with regards to the career of a particular dancer, the history of a particular dance or any number of other aspects of information in the archive. At this stage, a new multimedia document could be created for access by others. This may require a separate authoring or programming mode to accomplish. In this context, the concept of layers is important, since individual users may wish to create their own organizations or views of elements.

7 Conclusion

In closing, this paper has endeavoured to sketch an implementation of a computerized multimedia archive, which might be plausible in the near future.

Under this system, information would be made accessible to people who would not normally make use of a computer system. This indicates that the challenge is focused on the user interface provided by the archive information system. However, the database manager will still form the backbone of the system and cannot be neglected.

The hope for the system is that it will encourage others to use the archive of Mr. Cunningham to regenerate and redistribute the treasure of his art. In doing so, the system will likely continue to grow with new types of information, which come as the product of elements already existing in the archive.

References

- [1] A. Becofsky and M. E. Bloom. Archives in motion: A unified preservation agenda. Proposal, June 1992.
- [2] E. A. Bier and A. Goodisman. Documents as user interfaces. In *EP90: Proceedings of the International Conference on Electronic Publishing, Document Manipulation and Typography*, 1990.
- [3] M. E. Bloom, November 1992. Personal communication.
- [4] I. R. Campbell-Grant and P. J. Robinson. An introduction to iso dis 8613, “office document architecture,” and its application to computer graphic. *Computers and Graphics*, 11(4):325–341, 1987.

- [5] R. C. G. Cattell. Introduction to next generation database systems. *Communications of the ACM*, 34(10):30–33, 1991.
- [6] D. D. Chamberlin and C. F. Goldfarb. Graphic applications of the standard generalized markup language (sgml). *Computers and Graphics*, 11(4):343–358, 1987.
- [7] J. Banerjee et al. Data model issues for object-oriented applications. *ACM Transactions on Office Information Systems*, 5(1):3–26, 1987.
- [8] M. Conner et al. Developing language neutral class libraries with the system object model (som). Poster session at OOPSLA '92, 1992.
- [9] S. Christodoulakis et al. An object oriented architecture for multimedia information systems. *Data Engineering*, 14(3):4–15, 1991.
- [10] E. A. Fox. Advances in interactive digital multimedia systems. *IEEE Computer*, 24(10):9–21, 1991.
- [11] K. A. Frenkel. The human genome project and informatics. *Communications of the ACM*, 34(11):40–51, 1991.
- [12] D. J. Gemmill and S. Christodoulakis. Principles of delay sensitive sensitive multimedia data retrieval. In *Proceedings of the International Conference on Multimedia Information Systems*, pages 147–158, New York, 1991. McGraw-Hill.
- [13] M. Haggerty. Computer painting in a different light. *IEEE Computer Graphics and Applications*, 12(6):4–6, 1992.
- [14] F. G. Halasz. Reflections on notecards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, 31(7):836–852, 1988.
- [15] J. H. Irven, M. E. Nilson, T. H. Judd, J. F. Patterson, and Y. Shibata. Multi-media information services: A laboratory study. *IEEE Communications Magazine*, 26, 1988.
- [16] J. V. Joseph, S. M. Thatte, C. W. Thompson, and D. L. Wells. Object-oriented databases: Design and implementation. *Proceedings of the IEEE*, 79(1):42–64, 1991.

- [17] E. Kaplan. *Cunningham/cage*. Film, 1991. 95 minutes.
- [18] W. Kim. A model of queries for object-oriented databases. In *Proceedings of the Fifteenth International Conference on Very Large Data Bases*, pages 423–432, 1989.
- [19] W. Kim, N. Ballou, H.-T. Chou, J. Garza, and D. Woelk. Features of the orion object-oriented database system. In *Object-Oriented Concepts, Databases, and Applications*. ACM Press, 1989.
- [20] T. D. C. Little and A. Ghafoor. Spatio-temporal composition of distributed multimedia objects for value-added networks. *Computer*, 24(10):42–50, 1991.
- [21] L. F. Ludwig, N. Pincever, and M. Cohen. Extending the notion of a window system to audio. *Computer*, pages 66–72, August 1990.
- [22] M. Malek. Integrated voice and data communications overview. *IEEE Communications Magazine*, 26(6):5–15, 1988.
- [23] F. Mintzer, Y. L. Yao, and J. D. McFall. A computer system for scanning and cataloging the art of andrew wyeth. Research Report RC 17772, IBM Research Division, 1992.
- [24] S. R. Newcomb, N. A. Kipp, and V. T. Newcomb. The “hytime” hypermedia/time-based document structuring language. *Communications of the ACM*, 34(11):67–83, 1991.
- [25] P. Samuelson. Legally speaking: Digital media and the law. *Communications of the ACM*, 34(10):23–28, 1991.
- [26] B. Schatz. Telesophy: A system for manipulating the knowledge of a community. In *GLOBECOM*, 1987.
- [27] T. Schiphorst, T. Calvert, C. Lee, C. Welman, and S. Gaudet. Tools for interaction with the creative process of choreography. In *CHI '90 Conference Proceedings*, pages 167–174, 1991.
- [28] S. Shimojo, T. Matsuura, K. Fujikawa, S. Nishio, and H. Miyahara. A new hyperobject system *harmony*: Its design and implementation. In

Proceedings of the International Conference on Multimedia Information Systems, pages 243–257, New York, 1991. McGraw-Hill.

- [29] A. Silberschatz, M. Stonebraker, and J. Ullman. Database systems: Achievements and opportunities. *Communications of the ACM*, 34(10):110–120, 1991.
- [30] A. R. Smith. Plants, fractals, and formal languages. *Computer Graphics*, 18(3):1–10, 1984.
- [31] J. B. Smith and S. F. Weiss. Hypertext. *Communications of the ACM*, 31(7):816–819, 1988.
- [32] K. Smith and S. B. Zdonik. Intermedia: A case study of the differences between relational and object-oriented database systems. In *OOPSLA 87*, 1987.
- [33] S. Smoliar. Video classification and understanding. Research note, Institute of Systems Science, National University of Singapore, August 1992.
- [34] F. Tompa. A data model for flexible hypertext database systems. *ACM Transactions on Information Systems*, 7(1):85–100, 1989.
- [35] D. Woelk, W. Kim, and W. Luther. An object-oriented approach to multimedia databases. In *ACM Proceedings of SIGMOD*, 1986.
- [36] N. Yankelovich, B. J. Haan, N. K. Meyrowitz, and S. M. Drucker. Intermedia: The concept and the construction of a seamless information environment. *Computer*, pages 81–96, January 1988.