

A successful project utilizing various techniques for presenting material in CD-ROM format is described and problem areas are identified.

# Toward Multimedia



Academic conferences are a long-standing and effective form of multimedia communication. Conference participants can transmit and receive information through sight and sound, that is, by viewing individuals, text, and graphics and by hearing the spoken word. This same-time, same-place communication is sufficiently valuable to justify large investments in time and travel funds. Traditional printed conference proceedings and session videotapes are attempts to recapture the value of a live conference, but they are limited both by their delivery medium and by the significant differences from the conference presentation. Video- and audiotapes require as much time as the recorded events, while printed proceedings cannot convey the dynamic nature of a conference. But more importantly, they are not integrated with each other.

**Matthew Cheyney**

**Peter Gloor**

**Donald B. Johnson**

**Fillia Makedon**

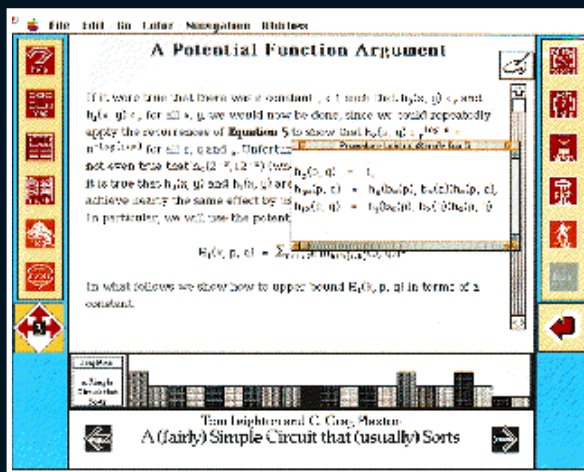
**James Matthews**

**Panagiotis Takis Metaxas**

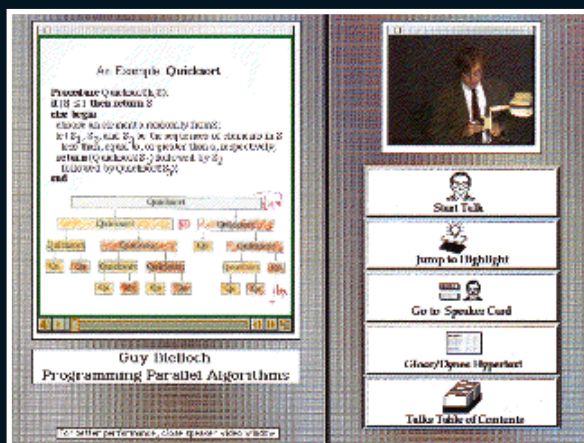


# Conference **Proceedings**

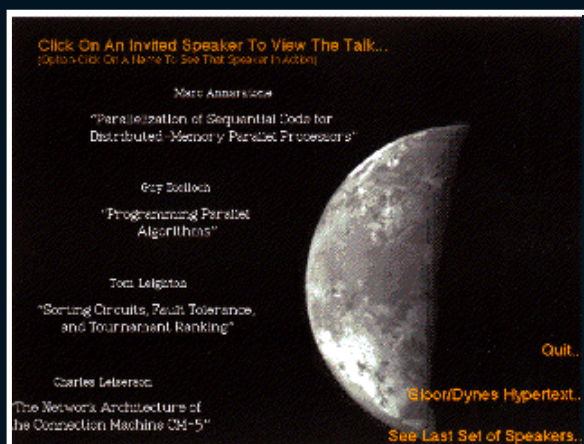
Photos: LEFT: Graphistock / David O'Connor, RIGHT: Tony Stone Worldwide / Paul Cherilis



**Figure 1.** The Article screen of the navigation shell



**Figure 2.** The Talk screen. On the left is the transparencies movie, and on the top right of the screen is the speaker movie.



**Figure 3.** The Contents screen

In June 1992, Dartmouth College hosted the first annual Dartmouth Institute for Advanced Graduate Studies (DAGS) symposium on the topic of "Issues and Obstacles in the Practical Implementation of Parallel Algorithms and the Use of Parallel Machines." In preparing the DAGS '92 Multimedia Conference Proceedings we aimed to address the shortcomings of the traditional means of presentation, and therefore deliver more of the value of an academic conference to our audience. Recently, conference organizers have realized the possibilities multimedia productions can have, and some first efforts at multimedia proceedings on CD-ROM have already been published [6, 10]. The results, however, are a collection of papers in two or three different formats that can be used mainly as "digital microfiche," along with a few digitized movies. Their main problems are a lack of interaction among the components and a lack of fast searching mechanisms. They were produced to serve as an alternative to printed proceedings, yet printing the articles reveals the absence of many pictures or graphics. Even though these are efforts in the right direction, we believe they fall short of the capabilities of multimedia.

The DAGS '92 CD-ROM on Parallel Computation [4] delivers text, graphic, audio, and video information as an integrated whole, with extensive provisions for random access and hypermedia linking. We believe this project provides a model for future conference publications and highlights some of the research issues that must be resolved before similar publications can be quickly and inexpensively produced. The experience gained from this effort will be applied not only to multimedia conference proceedings of the future, but also to multimedia textbooks and learning environments (as it already has, in [8]). We view this as the strongest point of our experiment and research and a direction that should be pursued further.

## Background

The DAGS symposium program consisted of eight talks by invited speakers plus 13 contributed talks presenting a total of 22 papers (one invited talk spanned two papers). The topic is a central problem in the area of parallel computation today. It is therefore appropriate to make the results as widely available and accessible as possible to the parallel computing research and teaching communities. So, in addition to the usual printed proceedings [5], it was decided to publish multimedia proceedings that would capture as much as possible of the conference atmosphere.

## Description of the System

The multimedia proceedings on CD-ROM contain the 22 papers in hypertext format, integrated with the eight invited talks in Apple's QuickTime movie format.



Hypermedia links and random access are used extensively in this integrated system. The CD-ROM includes:

- a navigation shell (Figure 1) that facilitates hierarchical navigation in hyperspace;
- hypertext versions of the 22 papers presented;
- digitized movies of the eight invited speakers delivering the conference talks;
- their transparencies, including marks that were added during the talks;
- hyperlinks connecting relevant parts of the proceedings; and
- a bibliography

In the future, we plan to include animations of some of the algorithms presented. Furthermore, the system can be extended in the sense that users are able to create their own hyperlinks among objects they deem relevant, do searches on keywords, and keep notes on the documents.

### A Typical Session

In a typical session, users can first get a quick overview of the contents of the CD-ROM. Then they can follow a talk on a particular topic that seems interesting by watching the movies of the transparencies and of the speaker (Figure 2). They can also get an overview of the talk by using the pop-up menu containing the section titles of the talk (by pressing the button "Jump to Highlights") or by skimming through the transparencies. If, at some particular point, the speaker mentions a theorem without proof, as is usually the case in conferences, users can jump into the hypertext to read the omitted proof in detail. Then they can go back to the talk or continue reading the hypertext, bringing onto the screen several windows containing relevant information. They can also do a search on some keyword to find out who else mentioned this keyword during the conference. Presuming that the search brought up several candidate sections of papers, users can jump into another paper and continue reading from there or even jump into the video movie of the second speaker and see how the material was presented during that talk. That way they can, for example, compare opinions on the speaker's view on a particular topic or simply see different presentation techniques on the same subject. In every place, users can make their own remarks on the subject being read and keep them for later examination or for filling in background gaps. They can also send these remarks to others who use the same systems by emailing them the file; the remarks will appear in context on the receiver's screen.

Considerable redundancy and flexibility is built into our system; evaluating which interface and functionality parts are crucial and should be kept and enhanced and which could be dispensed without affecting the overall performance of the multimedia proceedings will be part of future research and evaluation.

### Guided Tour Through "Parallel Computation"

In this section we will briefly describe a possible tour through the hypermedia document to give a first impression of the features of our system.

The system starts up with a title animation and then presents a table of contents (Figure 3) that gives the users direct access to the eight invited talks. In this example, the user decides to study the talk by Charles Leiserson, "The Network Architecture of the Connection Machine CM-5." The user then clicks on the title of this talk on the Contents screen, and the system jumps into the Speaker screen.

On the Speaker screen (Figure 4) users can get background information about a particular talk by reading a short biography of the speaker, obtain a list of the talk's highlights, read the hypertext version of the paper related to the talk, or actually watch the talk. Another button can take users back to the previous Table of Contents screen. The user here decides to view the talk and clicks on the View Talk button.

The system then switches to the Talk screen (Figure 5) where the users can watch the actual talk from the beginning by clicking the "Start Talk" button or go back to the Speaker or Contents screen by clicking the appropriate button. Users can also jump to any position within the talk using the slider controls on the left panel underneath the transparencies.

Another useful feature is the option to jump directly to the beginning of a subsection within the talk using the "Jump to Highlight" button shown in Figure 5. Upon pressing that button, the user is presented with an ordered pop-up list of the highlights of the talk (Figure 6). In this example, the user chooses to watch Charles Leiserson explain the CM-5 Data Network.

When the user releases the mouse button, the system jumps directly to the CM-5 Data Network section within Charles Leiserson's talk (Figure 7). After following the talk for a while, the user decides to read the paper associated with the talk, and does so by clicking on the "Gloor/Dynes Hypertext" button.

The system now jumps directly to the beginning of the paper associated with the talk. On the left of the Article screen (Figure 8) is a palette of tools that provide (from top to bottom) help, a map of the hypertext contents, a table of contents, bookmark facilities, a notebook, and marginal notes. Below is the "navigation diamond," which marks the hypertext node level and its connections to other node levels. On the right side is another palette, with buttons that provide (from top to bottom) a list of topics related to the current text node, the ability to add hyperlinks between text nodes, a history of the nodes visited by the user in this session, the ability to create a path through the nodes (or to use a predefined one that






Figure 4. The Speaker screen

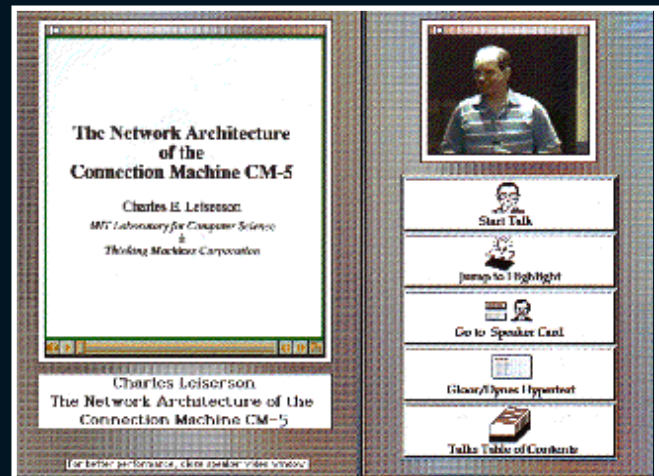


Figure 5. The Talk screen

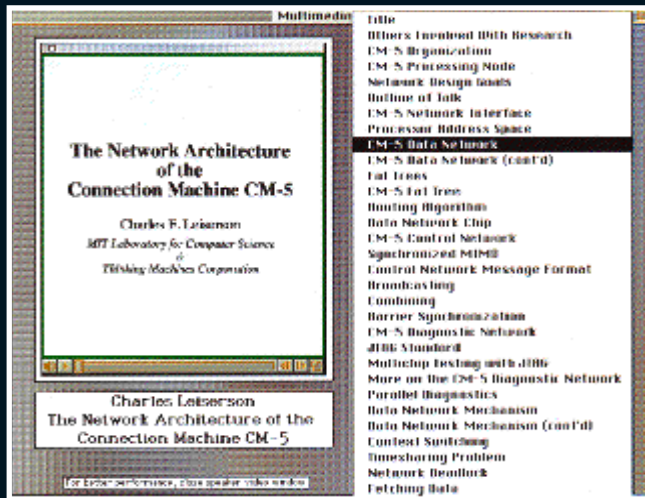


Figure 6. Highlights of the talk

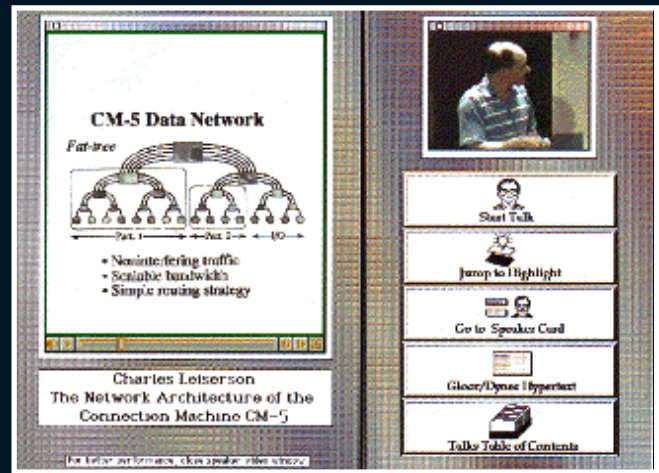


Figure 7. Dynamic access of a random section of the talk is provided

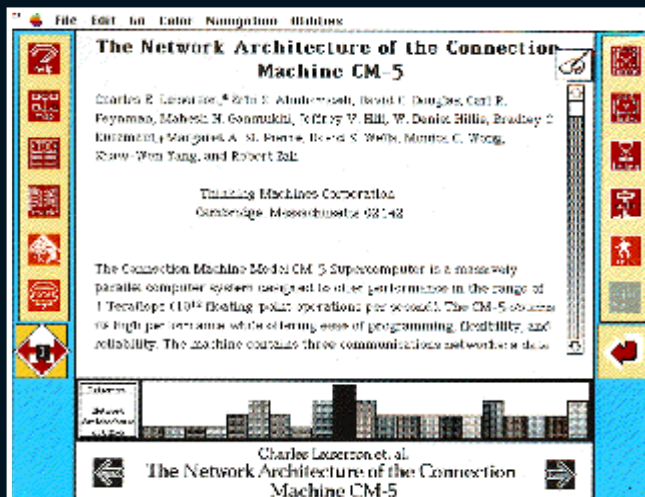


Figure 8. The hypertext version of the paper

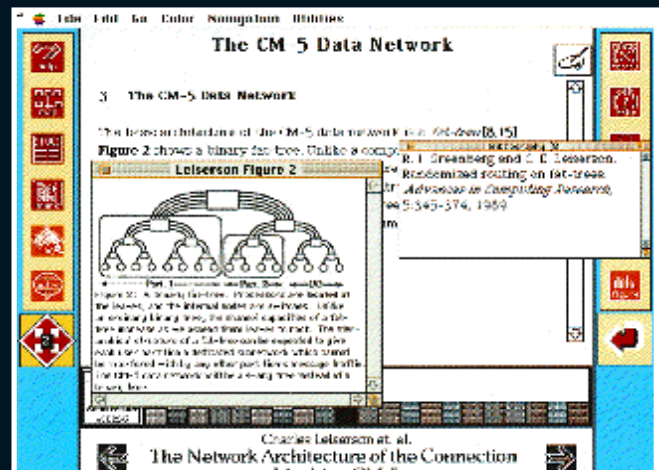
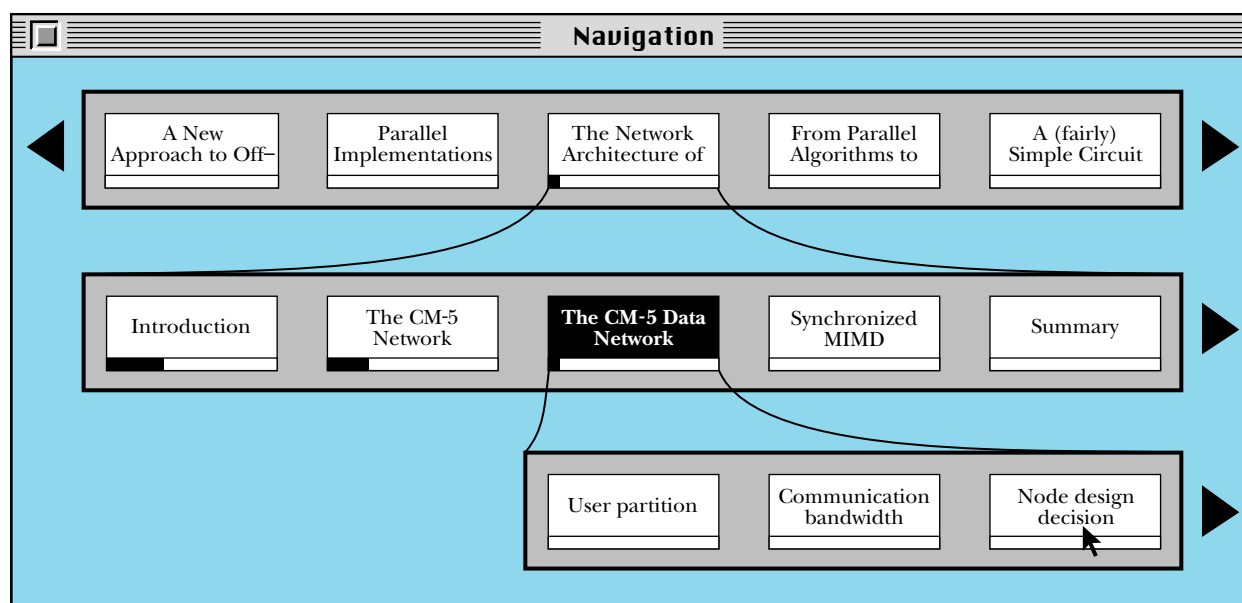


Figure 9. Figures, bibliographies, and footnotes appear in separate windows



takes them through the proceedings in a linear fashion), a link to the Talk screen, and a button to the figures related to the current node (dehighlighted when there is none). Under all that is a button that returns the user to the node last visited.

Following the level 2 nodes by clicking on the navigation diamond, the user can get an idea of the contents of the paper. A bibliographical reference or a figure mentioned in the paper can be opened by clicking on the word (Figure 9). To get a feeling of the structure of this paper, the user decides to bring up the overview map by clicking the “map” button.

The system brings up the overview map, which shows a graphical structure of the whole hypertext document in relation to the point where the user is (Figure 10). This shows that the user is on the section entitled “The CM-5 Data Network” of the paper “The Network Architecture of...”. Information about previous and next sections also appears, along with information about the subsections of the current section. The user decides to jump directly into the Node Design Decision subsection by “option-clicking” this box on the map. There are numerous other features in our system that cannot be shown in this brief guided tour. The reader is referred to the user’s manual of the Parallel Computation CD-ROM [4] for a more detailed description of the user interface.

### Development Stages

The first step in producing the multimedia proceedings was to collect the raw material (video, sound, text) and edit it to produce the videos and the hypertext.

### Producing the Digital Talks

During the conference, the speakers’ presentations

were videotaped, their overhead transparencies were copied, and their papers were collected. To make the synchronization easier, we used two video cameras for the videotaping: one focusing on the speaker and one on the projected transparencies (Figure 11, top). It turned out that the second was by far the more useful. This came as a little surprise to us, probably because we were dealing with video and had the typical videotapes in mind. After the collection, all these materials were converted to digital form for computer-based processing using a commercial video card. Due to the space constraints of CD-ROM, our intended delivery medium, we decided to deliver only the eight invited talks in a full audiovisual form.

Despite the huge storage provided by the CD technology, we knew that we could not fit even the eight 90-minute video tracks on a CD-ROM. However, most speakers at scientific conferences are usually standing next to the projector so that they can change transparencies, mark them, and point to them. The audience gets most of the information by looking at the transparencies. (This is usually not true with nontechnical talks). Therefore, we decided to display a short video loop of each speaker by selecting a piece of the video with similar beginning and ending frames so that the loop transition would not be distracting to the viewers. This was certainly the right decision, because little valuable information was lost this way; it turns out that about two minutes of video is sufficient to convey a sense of a speaker’s appearance and mannerisms and is less distracting to the person studying the talk. Given the

**Figure 10.**  
The map and the  
hypertext version  
of the paper

small size of the video capture screen, one can see the speaker's movement and overall appearance, but not tiny details like the movement of the lips. So, with some surprise, we found that a loop of a few minutes of the speaker video could not be distinguished from the whole video.

This fact helped us with the video editing, because breaking the synchronization between the audio and video tracks allowed us to edit the audio without introducing video skipping. Editing both the audio and the video data of the speakers in a way that preserves synchronization and smooth transition between edited frames seemed a formidable task for which we had neither tools nor time. Moreover, it does not seem to contribute much to the viewing of the talk.

The video loops were digitized with a low-end video capture board and compressed to provide effi-

cient playback from a CD-ROM. The loops were kept small to enhance playback performance and to keep the lack of synchronization between the speaker's lips and words from becoming a distraction. Due to the low resolution of the 120-by-180-pixel QuickTime video window, the user would have to look very closely to notice the absence of synchronization between the speaker's lip movement and the voice track.

many topics and more sensitive linking between the papers and the talks.

Given the variety of the playback speeds of the commercial CD-ROM drives available, some computer configurations will have a hard time displaying both the transparencies video and the speaker video loops on the screen at a comfortable speed. For that reason, we have given the user the option of stopping the speaker video loop and replacing it with a static color picture of the speaker.

### Producing the Hypertext

By far the most time-consuming stage was to convert the 22 papers into hypertext format, as it took about two-thirds of the person-year developing time. Wanting to produce a true hypermedia product and not to sacrifice its usability, we decided to present the papers in hypertext form using an advanced hypertext engine.

**Given the variety of the playback speeds of the commercial CD-ROM drives available,** *some computer configurations will have a hard time displaying both the transparencies video and the speaker video loops on the screen at a comfortable speed.*

cient playback from a CD-ROM. The loops were kept small to enhance playback performance and to keep the lack of synchronization between the speaker's lips and words from becoming a distraction. Due to the low resolution of the 120-by-180-pixel QuickTime video window, the user would have to look very closely to notice the absence of synchronization between the speaker's lip movement and the voice track.

The audio track of each talk was digitized and then edited to remove pauses and noise words such as "umm" and "ahh." The edited talks were roughly half as long as the originals, and much more listenable. To improve the quality, we amplified most of them using a commercial sound processing application.

The overhead transparencies were scanned and edited for clarity and contrast; they were also made smaller to fit in the appropriate window of the interface program. One of the problems we encountered in this process was the fact that after diminution, transparencies that were poorly handwritten had low readability. So some of the transparencies were retyped.


Then, with the original videotapes as guides, we synchronized the transparencies to the edited audio tracks using a commercial video editor. The resulting "movie" reproduces the most important features of a talk, the speaker's words and the transparencies, and preserves their temporal connection. These "movies" were indexed to allow random access to a list of pri-

We used the Gloor/Dynes hypertext system [3] that was developed at MIT for the CD-ROM version of the *Introduction to Algorithms* textbook [1] and is based on Apple's HyperCard software.

We chose this hypertext system because it is a full-featured state-of-the-art hypertext engine on the target platform. Since it was co-developed by one of the authors of the DAGS '92 Multimedia Conference Proceedings, knowledge about the internal workings of the engine was readily available for this project. One of the most distinguishing features of the Gloor/Dynes hypertext engine is the ability to traverse the contents of the hypertext document at different conceptual levels. Thus, a hyperdocument can be structured into at most four hierarchical levels, ranging from broad overviews covering large numbers of concepts in a very superficial fashion to the focused, detailed level at which the majority of text is written (Figure 12). Besides the usual hypertext facilities, such as enhanced full-text search, history, a means for annotating nodes, bookmarks, and a navigation map, the system also includes level-specific navigation tools—the navigation diamond and viewfinder. For a detailed description of the system see [2].

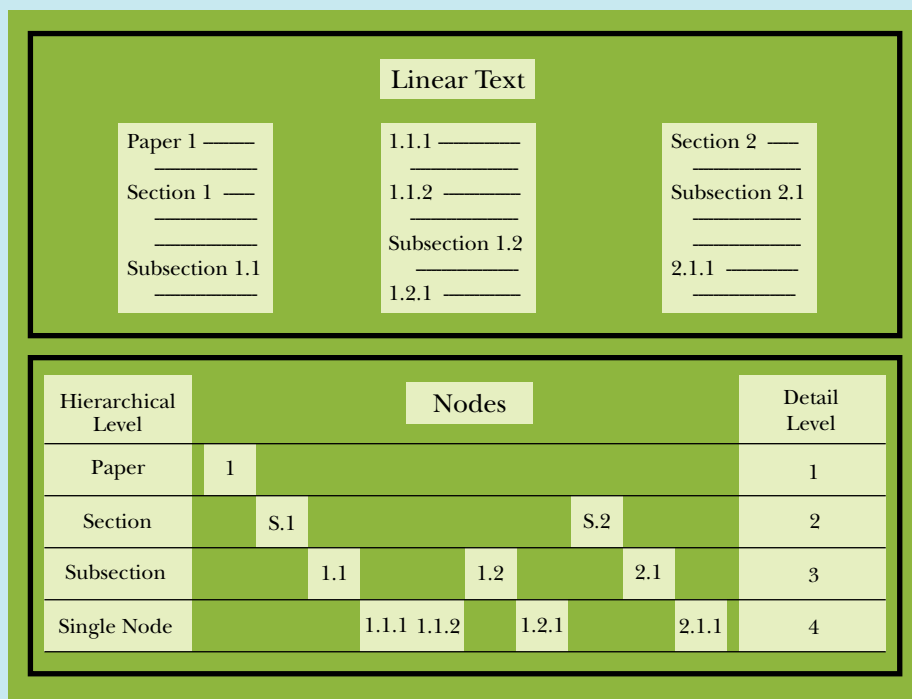
We first broke the 22 papers into hypertext "nodes" and assigned each node a "node level" that reflected its degree of generality. For example, the abstracts of all the papers are on level 1. This way, a





**Figure 11.**

Description of the developmental stages. The speaker and the transparencies were videotaped and digitized to produce the speaker loop and the transparencies movie (top). The papers were digitized and chunked into nodes to produce hypertext (bottom). Hyperlinks were finally added between the two components before the pressing of the CD-ROM (right). Extensive testing followed on machines with a variety of configurations.



**Figure 12.**

Converting linear text into hierarchical hypertext



user can quickly become familiar with the themes of the presented papers by visiting all the first-level nodes of the proceedings (Figure 12). All the beginnings of sections within a paper are on level 2, so a user can get an overview of the contents by following these nodes.

This “chunking” process was handled manually by computer scientists with expertise in the subject area. It appears that one cannot automate this process unless the authors have written their papers following some carefully predefined specifications. Since we did not anticipate this problem well in advance, we did not provide such specifications during the writing of the papers. The result was, of course, that each paper was written in a way that reflected the personal style of the author, not the expectations of a node-based hypertext. We tried to have authors provide chunking information afterward, but the results were not sufficiently consistent from paper to paper, so in the end all chunking was performed by the experts of our team. However, the author response to our “chunking guidelines” revealed flaws in the description<sup>1</sup> and helped us devise better specifications for this task.

Then we converted all the papers to HyperCard format. Every author provided Postscript versions of the papers, and half of them provided electronic versions in TeX or LaTeX format; the latter were converted to HyperCard form with homemade utilities and manual cleanup. Unfortunately, we could not use PostScript files for the hypertext engine, so the remaining papers were scanned, processed by optical character recognition (OCR) software, and then manually corrected. A great number of errors were introduced by the scanning and recognition process, and some pages were simply retyped. A number of text features required special treatment. HyperCard does not support mathematical text, not even subscripts or superscripts, so special fonts were used in their place. Uncommon symbols were similarly provided by custom fonts. This still leaves out very complex equations, which were scanned and displayed as graphics using custom software.

Figures were scanned and edited for clarity. The text was manually marked to provide links to citations, tables, and figures. Each hyperlink in the text appears as a boldfaced word. The system supports multiple windows containing scanned figures, tables, or bibliographic information. Finally, hyperlinks leading to referenced sections were introduced.

## Testing

The final step was to integrate all these elements into a single user interface and to test the resulting sys-

tem. The interface was designed to be simple, usable, and attractive. Extensive color graphics and online help facilities were built into the navigation shell. The system was tested on a number of machines with different capabilities and configurations. Special care was taken to optimize the transfer data rate so it would perform well on a variety of commercially available CD-ROMs.

We spent about one person-month testing the system on a variety of machines with different configurations at Dartmouth and Wellesley. In addition to the development machine (a Quadra 950 with a double-speed CD-ROM drive), we tested the system on another Quadra 950, a Quadra 700, a Centris 610, a Mac IIvx, and a Mac IIsi. Testing machines were equipped with double-speed CD-ROM drives. The last configuration represents the testing of the minimum system requirements.

**T**he most serious problem we encountered was conflicts between system extensions, which caused HyperCard to freeze at times. Experimenting in a trial-and-error fashion, we were able to recognize and fix the conflicts. We also encountered several minor problems with screen size (positioning the windows to show as much as possible on a 13-inch monitor), and some performance issues (a Mac IIsi with a single-speed CD-ROM was not fast enough to feed the two video windows). We produced three testing CDs that we repeatedly tested before achieving satisfactory results.

## Research Issues

The production of the DAGS '92 Multimedia Proceedings has presented us with two subjects for future research: improving the creation process of multimedia proceedings and evaluating the effectiveness of multimedia proceedings.

## Multimedia Authoring Systems

How can such proceedings be assembled in a short time with less manual labor? Even though, in the beginning, some production delays were caused by evaluation of the best alternative manual solution to each problem we faced, the experience gained has resolved most of these problems, and we now need to focus on the automation of the manual steps.

Some of the time-consuming steps, such as digitizing papers and transparencies, could be eliminated if the source material were available in electronic form. Removing pauses and noise words from audio sources automatically should also be possible with sufficiently sophisticated techniques. We are currently experimenting with more advanced software that can help in that direction [7]. Our goal is the design and implementation of a generally usable multimedia

<sup>1</sup>For example, “node” is a term that should be backed by examples to be clearly understood.



authoring system specifically targeted toward the creation of multimedia proceedings.

When working with mathematical text, the ability to use TeX or LaTeX input is important. We recently developed a LaTeX to HyperCard translator [9] that can automate the task of manual inserting mathematical input. Converting linear text to a hierarchy of hypertext nodes may be the most intractable problem; it is difficult to see how to replace the human expert, in the near term. It is clear that projects such as this will not be undertaken if they require the person-year of labor we invested; so progress must be made in automating this process. We are currently working on a carefully specified list of rules and instructions that could help authors divide an article into hypernodes. Ultimately, we hope they will use this list in preparing the final version of their papers. This approach, if successful, will eliminate the most difficult problem of future productions.

### Usability Testing

The second subject for research concerns the effectiveness of the final result. We believe the multimedia proceedings provide most of the content of live talks and linear papers, with the advantage of random talk access and hypertext linking. Of particular interest is the use of this technology in the production of hypermedia books. It is often the case that a good textbook author is also a good speaker; hypermedia textbooks, based on the results of our experiment, are expected to have a positive impact on the writing of future textbooks. We intend to produce hypermedia textbooks that are based on a combination of multimedia versions of lectures and hypertext textbooks using the same technology as in the DAGS '92 Multimedia Proceedings. But this proposition needs careful evaluation, and the cost and benefit of certain features (such as video loops and hypertext) need to be further scrutinized in order to justify the substantial investment in production resources.

### Conclusions

We believe the production of the DAGS '92 Multimedia Proceedings is a step toward academic publications that more fully reproduce the experience of a live conference or classroom. Even though it was published on CD-ROM, it can be viewed as a first step toward publishing conference proceedings and journals on the Internet, an issue currently being addressed by the academic community, scientific societies, and publishers. Currently such publications are one of a kind and expensive to produce, and with clear but unmeasured advantages over their traditional counterparts. Our experience highlights these shortcomings, but it also suggests that, with further research, improved systems can be built with less effort and greater rewards. ■

### References

1. Cormen, T., Leiserson, C., and Rivest, R. *Introduction to Algorithms*. MIT Press, Cambridge, Mass., 1990.
2. Dynes, S.B. and Gloor, P.A. *Using hierarchical knowledge representation for an animated algorithm hypertext learning environment*. Tech. Rep. TNS group, MIT Laboratory for Computer Science, 1992.
3. Gloor, P., Dynes, S., and Lee, I. *Animated algorithms*. MIT Press, Cambridge, Mass., 1993.
4. Gloor, P., Makedon, F., and Matthews, J., Eds.; Johnson, D.B. and Metaxas, P.T., Contr. Eds. *Parallel Computation: Practical Implementation of Algorithms and Machines*. TELOS/Springer-Verlag, Santa Clara, Calif., 1993.
5. Johnson, D.B., Makedon, F., and Metaxas P.T., Eds. *Proceedings of the 1992 DAGS Conference*. (June 23–27 Hanover, NH), 1992.
6. MacSciTech. *Proceedings CD-ROM of the 1992 MacSciTech Conference on Scientific and Engineering Applications for the Macintosh*. (January 15–17, San Francisco, Calif.) 1992.
7. Matthews, J., Gloor, P., and Makedon, F. VideoScheme: A Programmable Video Editing System for Automation and Media Recognition. *Proceedings of the ACM Multimedia '93*. (Anaheim, Calif.), 1993.
8. Metaxas, P.T. and Makedon, F. Eds. *Introductory Lectures on Data-Parallel Computing*. A.K. Peters, Wellesley, Mass., 1995.
9. Metaxas, P.T. and Zhao, K. *A LaTeX to HyperCard Translator*. Tech. Report CSD-TR11-1995, Wellesley College, Wellesley, Mass., 1995.
10. Rada, R. *Proceedings CD-ROM of the First ACM International Conference on Multimedia*. (August 1–6, Anaheim, Calif.), 1993.

### About the Authors:

**MATTHEW CHEYNEY** is an undergraduate student at Dartmouth College. email: mcheyney@dartmouth.edu

**PETER GLOOR** is an assistant vice president at Union Bank of Switzerland, Zurich, and an adjunct assistant professor at Dartmouth College. **Author's Present Address:** Union Bank of Switzerland Bahnhofstrasse 45, 8021 Zurich, Switzerland; email: gloor@lcs.mit.edu

**DONALD B. JOHNSON** was Professor and Chair of the Computer Science department at Dartmouth College.

**FILLIA MAKEDON** is Associate Professor at Dartmouth College and Director of the DAGS Institute. email: makedon@dartmouth.edu

**JAMES MATTHEWS** is Chief Programmer at Computing Services, Dartmouth College. email: James.W.Matthews@dartmouth.edu

**Authors' Present Address:** Department of Computer Science, Dartmouth College, 6211 Sudikoff Laboratory, Hanover, NH 03755.

**PANAGIOTIS TAKIS METAXAS** is Assistant Professor in the Computer Science department at Wellesley College and an adjunct assistant professor at Dartmouth College. **Author's Present Address:** Department of Computer Science, Wellesley College, 106 Central Street, Wellesley, MA 02181; email: pmetaxas@wellesley.edu

---

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.