11-1994

Building Multimedia Proceedings: The Roles of Video in Interactive Electronic Conference Proceedings

Samuel A. Rebelsky
Dartmouth College

Fillia Makedon
Dartmouth College

James Matthews
Dartmouth College

Charles Owen
Dartmouth College

Laura Bright
Dartmouth College

See next page for additional authors

BUILDING MULTIMEDIA PROCEEDINGS
THE ROLES OF VIDEO IN INTERACTIVE ELECTRONIC CONFERENCE PROCEEDINGS

Samuel A. Rebelsky
Fillia Makedon
James Matthews
Charles Owen
Laura Bright
Kenneth Harker
Nancy Toth
P. Takis Metaxas

Technical Report PCS-TR94-241

11/94

Building Multimedia Proceedings
The Roles of Video in Interactive Electronic Conference Proceedings

SAMUEL A. REBELSKY
FILLIA MAKEDON, JAMES MATTHEWS, CHARLES OWEN
LAURA BRIGHT, KENNETH HARKER, NANCY TOTH
Dartmouth College Department of Computer Science
Dartmouth Experimental Visualization Laboratory
6211 Sudikoff Laboratories, Hanover, NH 03755
Email: samr@cs.dartmouth.edu

P. TAKIS METAXAS
Wellesley College Department of Computer Science
106 Central Street, Wellesley, MA 02181

Modern computer systems have changed the way that conference proceedings can be presented and archived. No longer are researchers limited by printed text; electronic proceedings allow one to search the proceedings, add and share annotations, and create paths of related concepts through the proceedings. These additional capabilities extend the opportunities and benefit the thought processes of actual conference participants and the new *virtual participants* who experience the conference through the electronic proceedings.

In this paper, we discuss the construction of electronic conference proceedings, highlighting the role of talks and other presentations (and, particularly, the audio and video of these talks and presentations) in electronic proceedings. In particular, we discuss the benefits of incorporating video and audio in proceedings; describe the interface that guides the interaction between the text of a paper, the audio of the conference speaker, the video of the speaker, and the slides the speaker uses; detail experiences using limited video in proceedings; highlight the significance of the interface and careful editing by experts in the field, and recommend strategies and mechanisms for designers of proceedings and other multimedia documents which incorporate and link large amounts of text and video.

1 Introduction

Scientific information is disseminated in many forms, both formal and informal, including journal publications, word-of-mouth, the popular press, newsletters, and conferences. Given the high speed at which technology is evolving, conferences have an increasingly important role in presenting new results. Conferences provide a key mechanism for presenting scientific information because they are frequent enough to provide the latest results, and formal enough that results are usually refereed. At the same time, they are also informal enough that results in progress can be presented.

Because of these benefits, the number of conferences has increased dramatically. Hence, it is next-to-impossible for scientists to follow all the conferences in their areas of interest. Most conferences produce printed proceedings to allow those who could not attend the conference to learn about the results presented at the conference and to provide a permanent record of these results and ideas.

While printed proceedings provide a great deal of information, much more can be done with text-based electronic conference proceedings: (a) allow for greater access (as not every library or researcher gets every proceedings, but most researchers are networked), (b) employ modern searching and synthesis algorithms, and (c) provide tools for manipulating and annotating the contents of the proceedings. These tools help researchers tame the potentially overwhelming mass of knowledge in conference proceedings and to mine the pieces of information this wealth of knowledge can provide.

Still, such electronic proceedings do not contain everything attendees take away from a conference. If the papers were the only benefit conferences provided, researchers would order proceedings instead of attending the conferences. Conferences give attendees a chance to discuss research with each other and to hear presentations on results. These presentations supplement the papers with tutorials, summaries of background material, discussions of the thought processes that led to the results, diagrams that better

explain key parts of the research, and questions from the audience. Proceedings which do not include these talks represent only a fraction of the conference; proceedings which include presentations in addition to the submitted papers provide a much closer simulation of the conference (admittedly without the informal conversations and interactions amongst participants).

Over the past few years, the Dartmouth Experimental Visualization Laboratory (DEVLAB) has developed and published interactive multimedia conference proceedings [GMM93,MMR94] that incorporate audio, video, and graphics from conference presentations, along with electronic hypertext versions of the papers from the proceedings. These presentations significantly extend the power and usefulness of a text-based electronic proceedings.

An interactive multimedia proceedings that incorporates both talks and papers presents many opportunities for use. Such proceedings maintain the explanatory and reference roles of traditional proceedings, but add many new roles. Proceedings with audio and video materials can play an extended educational role, help introduce students and researchers to both basic and front-line topics in a field, and teach new students much about the field and about the dissemination mechanisms of science in general. Talks give new students a sense of the people behind the key research in the field, introduce students to modes of presentation and slide design, and help humanize research. Electronic proceedings that incorporate presentations can further enhance and encourage electronic collaboration among researchers. Sophisticated proceedings more closely involve the readers of the proceedings in the conference materials, creating what amounts to a *virtual participant.*

In this paper, we present a comprehensive overview of electronic proceedings. In particular, we

- summarize our experiences in producing electronic proceedings that include papers, notes, talks, and tutorials;
- present benefits and drawbacks of incorporating video and audio in proceedings;
- describe essential and useful features for an interactive multimedia proceedings;
- summarize reactions to our use of video; and
- recommend strategies and mechanisms for designers of proceedings and other multimedia documents which link large amounts of text and video.

2 Electronic Conference Proceedings

The wealth of information presented in scientific conferences comes in a wide variety of forms. However, the proceedings from a conference traditionally contain only the papers presented at the conference (and, perhaps, short written summaries of poster sessions and panels). This material is traditionally prepared in advance of the conference. Missing are the talks in which researchers present results, the interactive and noninteractive demonstrations that exemplify and explain results, the panels, the question and answer sessions, the informal discussions, the workshops, and other components.

A few conferences have supplemented their printed proceedings with video and audio *tapes* of the sessions which add presentations, a key element of conferences. Often, such hybrid proceedings are geared toward a broader audience with more tutorial and explanatory material and may contain post-talk questions and answers.

While both printed and taped proceedings provide a significant amount of vital information, the format in which they provide this information is less than ideal as they fail to meet the needs of the virtual participants interested in perusing the results presented at the conference.

Typical uses of proceedings include scanning to determine what results and ideas are current in the field (e.g., "What are the people saying about ——— these days?"); traditional searches for information relevant to a topic (e.g., "Which of these papers mention ———?"); and searches used for recall (e.g., "I remember that there was a paper that included a good research procedure.; it also mentioned ———. Which paper was that?"). While printed proceedings can support the first use, most printed proceedings lack an index or other mechanism that supports the search-related uses. Additionally, printed proceedings also fail to convey the *dynamic* nature of a conference.

Taped proceedings share many of these drawbacks. While taped proceedings help convey some of this dynamic nature, a major drawback is that they require the virtual participant to experience the material *sequentially*, with the viewing of the tape often using as much time as the recorded events. Although features not possible in the "live experience" can be added, such as the ability to review and scan ahead, many of the features that researchers would want are lacking. In particular, they provide no searching or segmenting (e.g., "where does the speaker describe the key results?").

In addition, the video and text in a hybrid proceedings are rarely, if ever, tied to each other. When listening to a talk, one would like to be able to quickly jump to the corresponding part of the paper (perhaps to see the results described more formally). Similarly, when reading a paper, one might want to hear the informal description of that part of the paper. Furthermore, since papers at a conference are often related to each other, it would be useful, when scanning one paper, to find a reference to another paper and be able to "bring up" that other paper.

Computer technology, particularly information retrieval mechanisms combined with multimedia systems, can solve many of these problems by providing the features that virtual participants need and want. These technologies make the proceedings more usable and therefore more useful. However, they are still limited in their capabilities [Buf94, CR93, DDF+90, Dum90, Dum91].

Electronic proceedings may also include other materials that are not traditionally included in proceedings. There are many times that authors want or need to distribute additional materials related to their subject matter. Such materials include program listings, demonstrations, applications, and video. In the past, such distribution was done outside of the conference proceedings (e.g., on floppy disk, on videotape, or as files retrievable over the network) through special efforts of the authors. Electronic proceedings allow such materials to be incorporated directly in the proceedings.

In recent years, conference proceedings from a variety of conferences have been made available electronically, usually on CD-ROM or available over the network [All94, GMM93, IEE93, MMR94, MST92, Rada93]. Given the value of talks, surprisingly few of these proceedings include the audio or video of talks. The majority are simply renderings of the papers in electronic format, with only a small video component (if any at all). Those that use video tend to present only videotapes of results which have been prepared in advance of the conference. Although some also include software-based demonstrations, prepackaged results and demonstrations do not provide the benefits of talks and therefore cannot replace the talks.

The remainder of this paper details the issues involved in the creation of an interactive multimedia conference proceedings that includes both submitted papers and presentations at the conference. These issues include the possible forms of materials that go into the proceedings, the steps involved in building a proceedings, the role of expert editors in building and designing the proceedings, mechanisms for automating the construction of proceedings, hindrances to publishing electronic proceedings, and the features a sophisticated interactive multimedia proceedings should provide. These discussions highlight some of the reasons that so few proceedings include audio and video: including audio and video *appropriately* is far from a trivial task, and requires a great deal of time and effort. One pays a high cost in time for including extensive materials: proceedings necessarily come out long after the conference.

2.1 Components of Interactive Multimedia Conference Proceedings

Since electronic proceedings can include a much broader range of materials than traditional proceedings, when designing interactive multimedia conference proceedings, the editors of a conference proceedings must first consider the nature and types of materials that might be included in the proceedings, then determine the roles these materials can or will play, and, finally, select what to include and what not to include. Many factors influence these decisions and choices, including (a) the needs, expectations, and capabilities of the users of the proceedings, (b) the costs (time, labor, funds, and other resources) of including more sophisticated content, and (c) the benefits additional media provide.

Let us now review the potential components of an electronic proceedings to illustrate the vast array of potential materials the designers of an electronic proceedings must consider. These materials are summarized in figure 1.

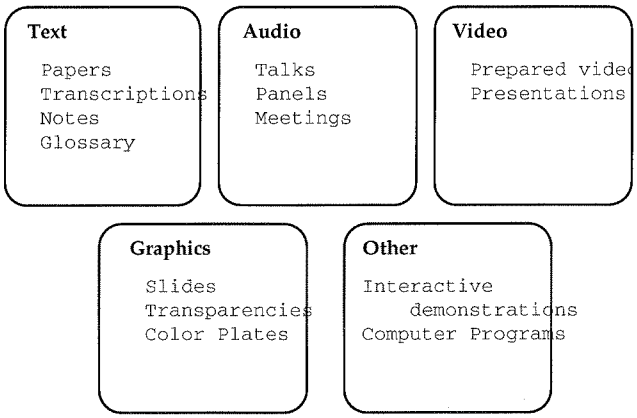

Figure 1 Potential components of interactive multimedia conference proceedings.

- **Papers.**

  Although electronic proceedings extend on traditional proceedings, there is still a large role for papers in an electronic proceedings.

  In an electronic proceedings, papers can be in ASCII format, in PostScript™ format, in a proceedings-specific format, in one of the many other formats for representing formatted text, or any combination of formats. The ASCII format makes it easier to search the papers and to reuse portions of the papers, but it usually sacrifices formatting, equations, and figures. Formats, such as PostScript™, that maintain the authors' formatting and illustrations and that make it easier to print papers are quite valuable. However, such formats require that the onscreen version of the paper follow the conventions of the printed page.

  There are both advantages and disadvantages to having a close relationship between the printed and the electronic versions of the proceedings. The main advantage is that those who have both versions can use the electronic version as a guide to the printed proceedings. A disadvantage of a close relationship is that the computer screen provides a very different reading mechanism than the printed page, and what is appropriate for one may not work for another. For example, almost no one reads long papers on screen, especially when reading formatted documents that require one to scroll through a virtual "page."

  Since neither ASCII nor PostScript provides an adequate onscreen representation, it is worthwhile, though time consuming, to recreate the documents in a more appropriate onscreen format. This format should include appropriate formatting, equations, and figures, and should be optimized for viewing on the computer screen. Of course, putting papers into this format may require additional effort (e.g., papers converted by "automatic" translators often need some post-processing).

  Space and cost often limit the size of papers in printed proceedings. Theoretically, the electronic proceedings may also contain longer versions of the papers. However, electronic proceedings should maintain page restrictions to encourage authors to present results succinctly. Authors interested in providing more extensive papers may present such papers in addition to the shorter papers.

- **The audio track from talks and panels.**

  Just as a proceedings may present talks using their audio, an electronic proceedings may incorporate the audio tracks of presentations as a separate component which needs

to be processed. The audio tracks should be edited to remove pauses, verbal ticks (e.g., "um"), and unnecessary phrases. Such editing is time consuming, but greatly improves the final product.[1]

Additionally, to make the audio more than a simple "tape," each audio track should be segmented, annotated, and indexed so that a virtual participant can quickly jump to or search for key points in the talk.

- **The video track from talks and panels.**

  While the audio track conveys most of the presentation, it does miss some important components that video can better reproduce. The video from a talk captures gestures and facial expressions that reveal more than the simple text of the speech.

  Video also gives a better sense of the speaker as a person; many of those with whom we have discussed conferences have commented that they find it important to know what speakers look like and how they teach.

  Finally, video (even very simple video) adds a level of comfort to the proceedings. Without video, the audio track is disembodied; with video, the virtual participant is made to feel like the speaker is actually lecturing.

- **Slides and transparencies.**

  Although talks are often viewed as having only an audio component and a video component, they also have a third component: the slides and transparencies that accompany the talk. These slides are key to the experience of a talk and have value of their own. Even without accompanying audio and video, slides summarize and illustrate the key concepts in the presentation.

  In an electronic proceedings, slides can be included in both graphic and text format so that they can be viewed and searched. Slides in an electronic proceedings can also be *animated* to mimic the annotations that speakers add during talks. Such animations include moving cursors and overlays that can be added and removed. Of course, these animated slides make it harder to produce the proceedings as they must be scripted and synchronized with the audio and video tracks.

- **Edited and annotated transcriptions of the presentations.**

  A multimedia proceedings should present materials in a variety of forms so that they can meet the varying needs of virtual participants. While it may seem that transcriptions add little to audio, they have many benefits.

  One can frequently read a transcription in less time than it takes to listen to or watch a talk. Transcriptions also accommodate the hearing-impaired. Additionally, transcriptions that are tied to the audio track make it possible to search the talks as well as the papers.

  As with all components, care should be taken in the preparation of transcriptions. They should be edited for readability and should not reproduce verbal ticks. Additionally, they should be annotated to describe that which cannot be inferred from text alone (e.g., to reflect the speakers inflection or gestures).

- **Speakers' notes.**

  Speaker notes prepared in conjunction with presentations provide another way for virtual participants to experience and understand presentations. For example, they may read the slides and notes instead of listening to the talk. They may also follow the notes while listening to the talk to get further information on the materials.

  These notes can also serve as tutorials on how to prepare a talk, which adds a new dimension to the proceedings.

- **Interactive demonstrations.**

  Interactive software-based demonstrations serve as supplements to papers and talks and give a visual demonstration of the results and ideas. Many ideas can be presented

[1]When preparing the DAGS proceedings, we found that such editing made the audio track up to 50% shorter.

more clearly and succinctly using such interactive demonstrations. Interactive demonstrations also have a very practical benefit: they can make the proceedings more interactive and enjoyable.

However, because each group traditionally prepares the demonstration software, these demonstrations frequently have inconsistent interfaces and will not work on all platforms. This lack of consistency contributes to learning overhead by the virtual participant who may not be able to immediately understand the abstractions and goals embodied by the demonstration.

There is a further drawback to interactive demonstrations: authors may worry about distributing their software as part of a proceedings and the licensing issues involved in such distribution.

- **Noninteractive videos.**

  Like interactive demonstrations, videos can help describe ideas succinctly and clearly. Noninteractive videos are easier to incorporate than interactive demonstrations, as videos eliminate problems of platform dependence and inconsistent user interface. Noninteractive video also requires less work than video of presentations, as noninteractive videos do not require synchronization among a variety of media (that is, the video may be stored and replayed as one object, and not as separate audio, video, graphic, and text components).

  However, videos also add to the labor and time required for electronic proceedings production as they may require a separate review process to ensure the significance and quality of the submitted video.

- **Tutorial materials.**

  A proceedings may include supplementary materials to aid in the comprehension of the results. Tutorial materials may include a glossary of key terms and a unified bibliography. Proceedings with such tutorial materials become more than a simple record of results; they may also provide an introduction to the field, become part of a course, and help conference participants (both real and virtual) to understand and use materials outside of their major subfield.

- **Records of informal discussions.**

  It is clearly impossible as well as undesirable to record every conversation that occurs at a conference. However, for some scientific conferences, it might be useful to collect summaries of attendee comments or "roundtable" discussion transcriptions. These records help preserve and present the ideas and materials that are usually lost to those who could not attend. These records can also serve as aids to those who did participate in the discussions.

- **A list of attendees with contact information.**

  This list is useful for both attendees and those who could not attend. It is especially useful for virtual participants who wish to know more about new results and need to contact someone whose address would otherwise be less readily available.

  While papers usually include a contact, the virtual participant may also be interested in topics and ideas presented outside of the main papers, talks, and panels. For example, the virtual participant may hear a researcher mention his or her own work during question and answer sessions. The list of attendees makes it possible to obtain more information on such work.

- **Informal pictures or video.**

  A record of the events and participants can further personalize the conference and highlight the difference between the proceedings and the conference itself.

- **Shared comments and annotations.**

  Science is not static, and the results presented at a conference may inspire followup ideas (both simple and complex). By including comments, the proceedings can change and grow. If these comments can be shared between participants, the comments of one

participant can further explain materials to other participants. Shared comments also encourage collaboration between attendees (and between virtual participants after the conference).

We do not recommend that every electronic conference proceedings include all of these materials. The preceding list is intended to suggest some of the sources that provide useful materials for electronic conferences. It is the editors' decision to evaluate their audience and identify which materials should be included in a particular conference proceedings.

2.3 Building an Electronic Proceedings

The steps involved in the construction of an electronic proceedings vary greatly depending on the types of materials to be included. For example, a proceedings that includes only text requires fewer steps and less work than one that includes well-linked audio and video. The content/work trade-off significantly affects the date by which the proceedings can be distributed; while a simple electronic proceedings may be ready at the time of the conference, a more sophisticated proceedings may not be ready until long after the conference.[2]

[2]We opt for months later because this time allows us to incorporate the happenings of the conference, to include a sophisticated interface, and to add annotations.

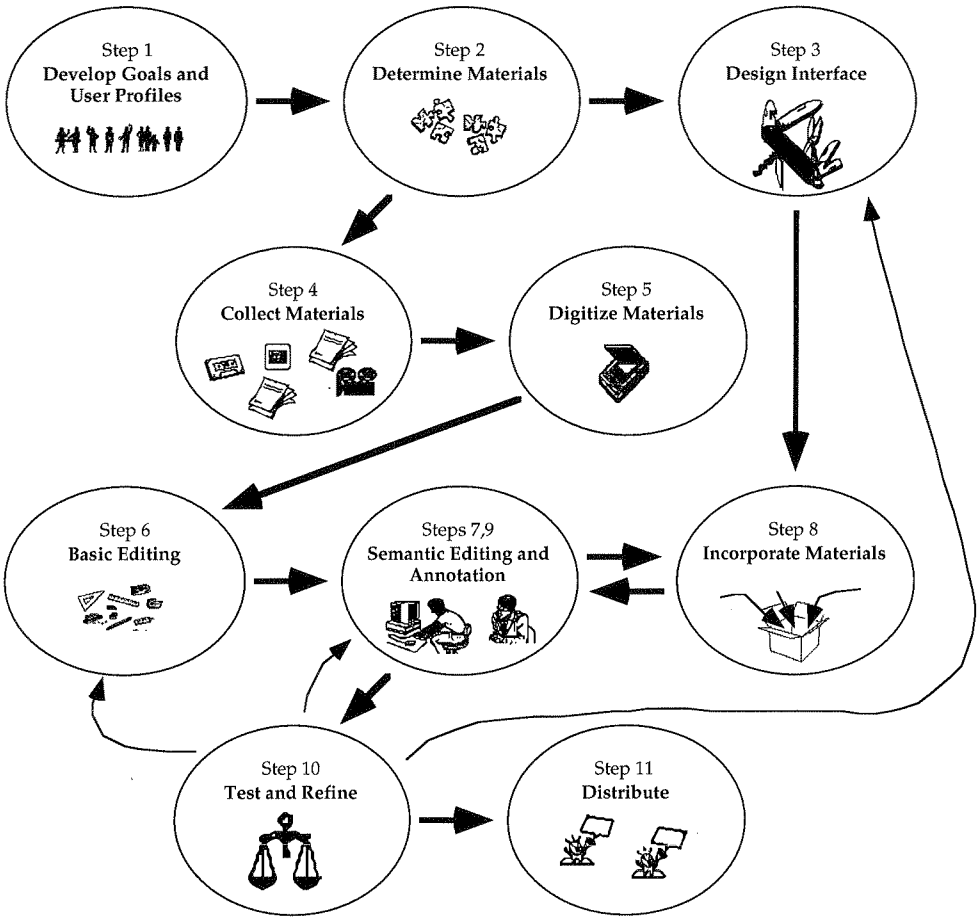

Figure 2 The steps involved in building an electronic proceedings. The goals, interface, and semantic editing are especially key to this process,. Although the process is presented as a linear series of steps, there is significant feedback between the individual steps. For example, the materials may affect the interface, and the editing processes may suggest other materials to collect.

The trade-off is not only time and cost but whether the proceedings is delivered at the time of the conference or months later. We have identified a set of common steps in the construction of any electronic proceedings, which are summarized in figure 2 and discussed further in the following paragraphs. Although the steps are presented linearly, many steps can be performed in parallel. Additionally, there is feedback between steps so that the observations of one step can affect previous steps.

For the interested reader, a more general description of the issues and processes in multimedia authoring can be found in [MRC+94]. The steps involved in the creation of the DAGS'92 electronic proceedings are detailed in [CGJ+94]. The steps involved in the creation of the IWANNT'93 electronic proceedings are detailed in [All94].

Step 1: Analysis. The development should begin with an analysis of the goals of the proceedings which includes an evaluation of the resources (manpower, finances) available and a profile of the expected and potential audience. The results of this evaluation form the basis of the further steps in the development of the proceedings.

Step 2: Determine Materials. The next step is deciding what components belong in the proceedings. Many issues can come into play in this decision process. Authors and speakers must be consulted to

determine what materials they are willing to have included.[3] The costs and benefits of including nontraditional media must be weighed. For example, transcription is costly and time-consuming, but transcriptions allow virtual participants to search talks and present the talks in a more efficient (albeit less comprehensive) format. The expected time frame for producing the proceedings also affects these decisions: an electronic proceedings that must be available at the time of the conference cannot include materials from the conference itself; an electronic proceedings that must be available within a short time after the conference cannot include materials that require significant editing.

Step 3: Design and Build Interface. The next step is to develop the interface for the proceedings. It is important to design and implement as much of the interface as possible before incorporating the materials so that one can better determine the format to record the materials and so that the proceedings can be created more quickly. However, creation of the interface may be done in parallel with collection, digitization, and editing of the materials (steps 4 through 7). The developers and designers of the interface need to consider ways to present materials, features to include and exclude (e.g., what types of searching should the interface include), the hardware platforms on which the proceedings will be made available, and the software platform used to implement the interface. The intended dissemination mechanism will also influence the design of the interface (e.g., Will it be made available on CD-ROM or networked? Will it be accompanied by a printed copy of the proceedings?). It is also important that the conference proceedings reflect the protocol, assumptions and audience sensitivities of a particular conference.

Eventually, it will be preferable to use a predefined proceedings interface. Unfortunately, as yet, there are no authoring packages available that support all the features that sophisticated interactive multimedia proceedings require. Hence, development of the interface is a necessary step in the construction of the proceedings.

Step 4: Collect Materials. The next step is to collect the materials that will make up the content of the proceedings. These materials are often collected in stages: papers, software, videos, and slides (when possible) before the conference; talks and discussions at the conference; comments, semantic links, and extended versions of papers after the conference. The ways in which these materials are collected will significantly affect the time and effort involved in the creation of the proceedings and the quality of the final product. It is preferable to obtain the materials in both electronic and hardcopy format. The electronic format eases transition to digital form; the hardcopy format provides an accurate master record to use as reference (and, when necessary, a source to be digitized).

Authors are key to the collection of appropriate materials. If the interface benefits from having papers in a specific format, authors may be asked to prepare material in that format (understandably, not all authors will be willing or able to use any particular format). Authors should also be asked to segment and annotate their papers.

Even if the design of the proceedings only calls for a few materials, it is important to obtain as many materials as possible as the design may change or unused materials may make it easier to correct materials used in the proceedings. For example, if the slides and not the video of a talk are included, the video may still aid the preparation of the slides as it will reveal changes the speaker makes and reinforce the ordering of the slides. The video also tells when slides are changed, which can affect audio editing and the synchronization of slides to audio..

Additional authors and editors may be employed to produce or synthesize special materials, including a glossary of key concepts, the table of contents, or an introduction to the field.

Step 5: Digitize and Convert Materials. Once the materials have been collected (or as they are being collected), they should be converted to a standard electronic format. Some materials will need to be digitized. While other materials may have been obtained in electronic format, they may need to be converted to another format or formats. For example, TEX documents (the format employed my many

[3]Some authors may prefer not to have their faces made public. Authors may request that their materials not be an format, such as ASCII or LATEX that permits easy copying. If authors are too restrictive of their materials, the proceedings will be both incomplete and inconsistent. Hence, authors should be encouraged to allow all materials to be included and the benefits of such inclusion should be clearly explained to the authors.

computer scientists and mathematicians) may be converted to PostScript, HTML (the HyperText Markup Language which is used for networked hypertext documents in the World-Wide-Web), and ASCII.

Steps 6 and 7: Basic and Semantic Editing. Once in digital form, the materials must be edited and broken into appropriate, self-contained segments. Traditionally, two types of editing are done: basic, format-based editing, and more sophisticated, semantic editing.

Non-experts may do the basic editing. Most of this editing is simple "clean up" required by the basic materials, by the recording process, or by the conversion process.. For example, it is useful to remove "um"s and "ah"s, pauses, and similar verbal "ticks" from the audio. This makes the audio much more pleasant to listen to.[4] One may also need to retype or redraw text and figures that were not adequately digitized and reformat some documents to fit the requirements of the computer screen. Some of these simple editing tasks may be performed automatically, but many must still be performed by hand to ensure a quality product.

Still other editing tasks must be performed by experts in the field. Such *semantic editing* includes segmenting the materials into coherent self-contained "chunks," checking the content of the materials, identifying key components of the video and audio tracks, and annotating individual materials. Semantic editing creates a new position in the creation of electronic documents: the expert electronic editor. Comprehensive editing of a talk or tutorial is very important in the production of useful proceedings. Careful semantic editing can make the difference between a useful, successful publication and a useless, boring one.

Preferably, the semantic authoring should be done by the authors of the materials, as they are the experts on the materials they are presenting and the way they wish to present the materials. However, audio and video editing are not nearly as straightforward as text editing, as they require the mastery of new software. Additionally, different authors may edit materials differently, so the proceedings editors will need to both set standards and do some additional editing.

Steps 8 and 9: Incorporation and Annotation. After the materials have been edited and segmented, they are incorporated into the interface. While this brings the proceedings closer to being a finished product, there are still more steps to be completed. Once in the interface, the materials can undergo further semantic editing. Experts can determine links between different parts of the proceedings (and between segments of a related talk and paper); annotate materials (with text and audio); synchronize the slides, audio, and video; and add similar semantic links to the materials. Experts may also create new "paths" through the proceedings, so that a virtual participant can follow selected topics through the whole proceedings. The semantic annotations add coherency to the collection of documents, making them seem more like a unified whole than a disjoint collection; make the proceedings more usable and useful; and generally enhance the experience that the proceedings provide. These benefits are not without cost; it requires significant expert time to add this semantic information.

Step 10: Testing and Refinement. At this point, the interface should be further tested and refined. Once the materials have been incorporated in the proceedings, the materials will suggest new features or changes to the interface. At this stage, the proceedings may also be distributed to a few "beta testers" who can suggest improvements to the interface, content, and semantic links. The proceedings should also be distributed to the authors who can verify that their results are presented correctly. Early dissemination to authors is very important, as it reassures the authors that they have control over their materials and because comments from authors help to repair incorrect semantic links and to add new semantic links. This testing and evaluation process should not be skipped as it invariably catches many errors.

Step 11: Distribution. Finally, the proceedings can be released (a) as a retrievable software package on the network, (b) as a remote "proceedings server" on the network, (c) on CD-ROM (often with commercial distribution), or (d) on a related medium.

---

[4]In most cases, this editing also significantly reduces the amount of audio material.

2.4 Automating the Construction of an Electronic Proceedings

Among the greatest challenges in the construction of interactive multimedia conference proceedings is that of automating the process of integrating and organizing the materials. In our experience, creating high-quality electronic proceedings with a large number of useful hyperlinks is a very time-consuming process. The production time for both [GMM93] and [MMR94] was over one year each. It is clear that in order for this effort to be successful in the future and to be accept and followed, the turnaround time has to be significantly shortened.

Some of the steps discussed above can be made shorter, automated, or eliminated. For example, once a comprehensive electronic proceedings interface[5] has been created, it can be reused with little or no change.

Some of the time-consuming steps, such as digitizing papers and transparencies, can be eliminated if the source material is available in appropriate electronic form. Removing pauses and verbal ticks 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 [MGM93]. We expect that after automated editing, the audio will require a pass by non-experts to further clean up the sound, and then a final pass by an expert to introduce hyperlinks, remove uninteresting parts, and synchronize the audio with the video. We have found that a carefully designed procedure requires two hours of basic editing and six hours of semantic editing[6] for each hour of source material. For a proceedings that includes twelve hours of video and audio, this means 96 hours of editing, or slightly over two weeks of full-time work.

Although receiving material in electronic format speeds the production process, it does not eliminate all the basic editing steps. Because authors will continue to produce papers in a variety of formats, sophisticated and comprehensive conversion programs must be developed that can flawlessly extract and convert content. While papers that contain only text might be converted transparently, at present, most conversions require significant re-editing or retyping because most papers include formulae and tables.

Semantic editing—particularly segmenting, annotating, and determining links between papers—may be the most intractable problem; it is difficult to see how to replace the human expert, in the near term. We are currently working on a carefully specified list of rules and instructions that could help the authors provide their own segmentation, links, and annotations. Ultimately, we hope that they would use this list in preparing the final version of the paper.

It is clear that the regular development of sophisticated multimedia proceedings will not be undertaken if such proceedings require multiple man-years of labor; so progress must be made in automating this process. Our goal is the design and implementation of a generally usable multimedia authoring system specifically targeted towards the creation of multimedia proceedings. This approach, if successful, will alleviate the most difficult problem of future productions. It is our belief that with the software and procedures we have introduces, a multimedia conference proceedings can be produced in one or two months by a small team.

2.5 Marketing Multimedia Proceedings: Obstacles and Evolving Strategies

Multimedia conference proceedings are software packages as well as sophisticated technical documents. Because they draw on these two dissimilar realms, they are hard to market and distribute. Traditional bookstores are reluctant to stock objects they view as "software" and not books. But most software stores are not an appropriate venue for electronic proceedings as the typical customers in a software store are not the intended audience of most proceedings.

Electronic proceedings necessarily have a smaller market than traditional printed proceedings. One of the obstacles is that not everyone who would be interested in a proceedings has ready access to a

[5]One that can meet the needs of a variety of types of proceedings for many different fields.

[6]Because pauses and other unnecessary materials have been removed, the material has been shortened and the editor's job is easier and quicker than it would be otherwise. Our estimate of six hours is therefore on the high side.

computer and CD-ROM drive. However, this is changing as more and more conferences find it worthwhile to produce electronic publications. Another obstacle is that of the choice of platform: a proceedings that is available only on select platforms (or that works less well on some platforms) has limited appeal. Again, this will change as more standard, cross-platform, multimedia development environments become available.

Even when it is possible to get electronic proceedings into the bookstore environment, they are difficult to sell. Most customers prefer to scan through a book or proceedings before purchasing it. Since few stores have the hardware necessary to demonstrate interactive multimedia publications, potential purchasers have no way of evaluating whether or not the package is appropriate for their needs. Given that many publishers attach traditional software shrinkwrap licensing restrictions to electronic books, potential purchasers may feel they cannot take the book home to try out.[7] Most researchers believe that an electronic book should be treated as a regular book in all respects, including "fair use" of the contents in educational and research situations. Software licensing restrictions often seem to disallow such use.

Additionally, many publishers are taking a very cautious stance with regard to electronic publishing. In part, this is because they are worried about the issues mentioned above. In part, this is because they do not yet have the infrastructure to support a new means of publication. We expect that in a few years the more forward-thinking publishers will not only increase publication of multimedia books, both in CD-ROM and networked form, but will also set up cooperative agreements with bookstores to ensure that computers that can be used to demonstrate and test the electronic books will be available in bookstores that sell such books.

3 The User Interface

An electronic proceedings requires more than useful content and excellent editing to succeed. It also requires a sophisticated interface that (a) is easy to understand and use; (b) follows appropriate aesthetic standards; (c) meets the needs of a variety of use, readers, researchers, and other virtual participants; and (d) provides a broad range of features. A bad interface (whether it be too simple, too complex, or too confusing) encourages researchers to return to traditional, paper-based research mechanisms. A good interface not only encourages researchers to use the proceedings, but also encourages new avenues of exploration and use.

There are many aspects to the interface for an electronic proceedings, including the search mechanisms it includes, the other features it provides, and the way it presents talks, papers, and the features for using those talks and papers. A typical proceedings will include separate windows for different papers and talks and for the features, as in figure 3. The individual features are described in the following sections.

[7] In fact, some potential customers of the DAGS proceedings have told us that the licensing agreement for the CD dissuaded them from purchasing the proceedings.

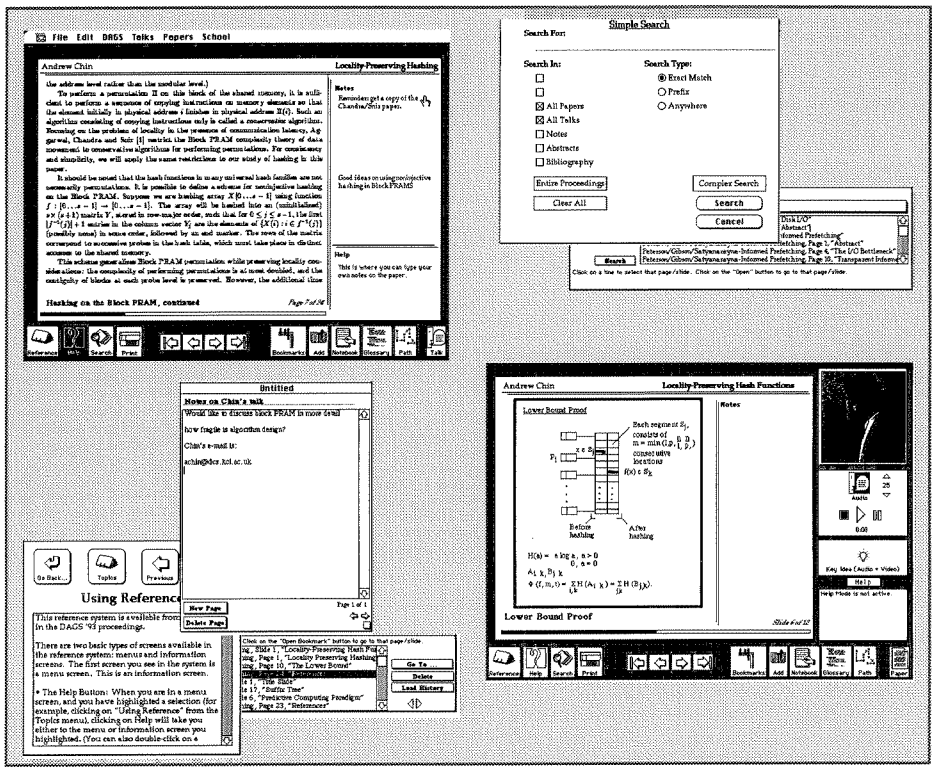

Figure 3 A complete proceedings interface. Most modern interactive proceedings provide a variety of windows, including talks, papers, a sharable notebook, a list of bookmarks, and a table of contents.

3.1 Search Mechanisms

The search mechanism is often the first complex feature that most users of electronic documents employ, and it is also one of the most commonly used features. Hence, careful attention should be paid to the ways the virtual participants can search the proceedings. There are a surprisingly large variety of search mechanisms that proceedings may provide, although most searches will be simple, keyword-based searches. Potential searching mechanisms for proceedings include:

- **Simple keyword-based searching**. In this searching method, one requests all objects that include a key word or phrase. This is the most common searching mechanism that people employ and corresponds closely to using an index in a printed proceedings. To allow for word variants, keyword-based searching might also allow one to extend searches to word prefixes (e.g. "geograph" can be used to find both "geography" and "geographical") or to limit searches to exact matches. Figure 4 presents the interface for a simple search feature.

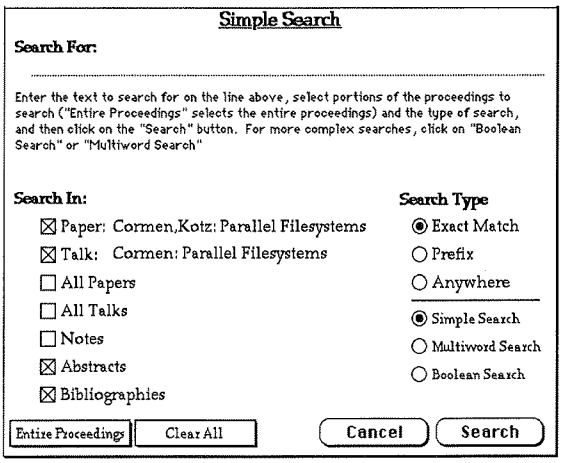

Figure 4 A simple searching mechanism. The user may search for a key phrase in a selected part of the proceedings. Because the user may not know the exact word used (or may prefer to search for only precise matches), a variety of search types are included.

- **Complex keyword-based searching.** The simple keyword-based searching can be extended in many ways. One might search for a subset of a collection of words (e.g., "all pages that include three out of the following words"). One might search for boolean combinations of words (e.g., "all pages that include ——— and ——— or ——— but not ———"). Often, boolean searches are restricted to particular forms (conjunctive or disjunctive normal form). Keyword-based searching may also be extended to include distance metrics (e.g., "all pages that include ——— and ——— within ——— sentences of each other."). Figure 5 illustrates two of these more complex keyword-based searching mechanisms.

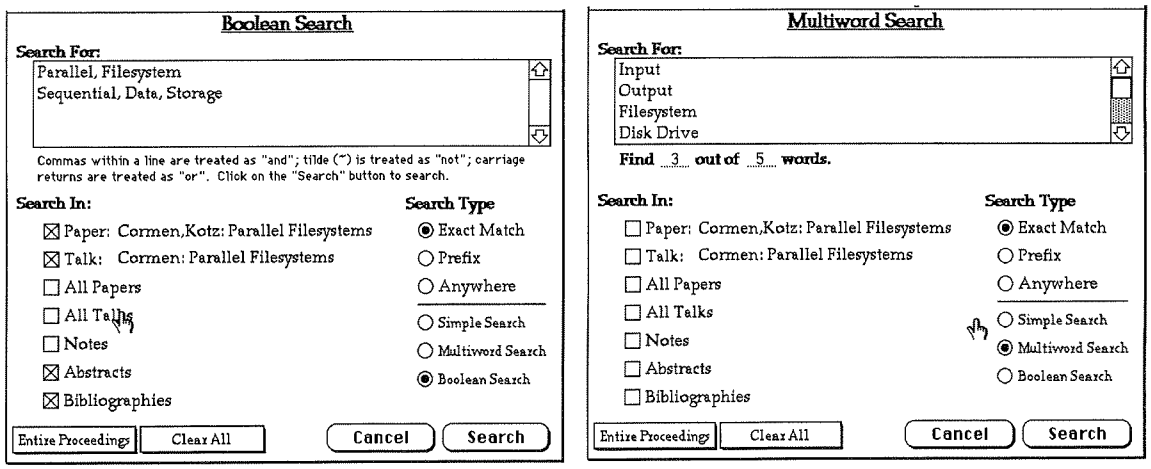

Figure 5 More complex keyword-based searching mechanisms. In the multiword search, one may search for parts of the proceedings that contain some subset of a set of words. In the boolean search, one may search for more complex boolean expressions. To simplify entry, commas on a line are treated as "and"s, tildes are treated as "not"s, and carriage returns are treated as "or"s.

- **Property-based searching.** Rather than searching by content, one may search by properties or restrict keyword searches to particular areas of the document. For example, one might search for articles with a particular author or that contain a keyword within

the abstract. Recent work suggests that layout properties should be included in such searches (e.g., "All papers with a chart on page 3") as many people recall papers visually. Property-based searching can also encompass reference chaining (e.g., "give me all papers that cite this paper"). However, this is still an open area of research.

- **Semantic retrieval**. More complex searching mechanisms can allow one to find similar papers (or similar chunks of other papers). In this type of searching, one first identifies an object of interest and then asks for similar objects. There are various mechanisms available for identifying objects that are similar to a given object, including Flash [CR93] and singular value decomposition [DDF+90].
- **Link annotations**. While semantic retrieval techniques can identify many related documents, they don't suffice for all purposes. Hence, an author or editor might provide links to related objects (both within and without a document).

Once the search has identified key chunks that meet the specified requirements, the interface should present these chunks in a clear and usable fashion, such as that of figure 6. One can then navigate through the results by selecting an object. One may also scroll through the found objects (both by using the arrows in the results window and by selecting some form of "Next" button on the individual pages or slides). Visual mechanisms (such as italics) may be employed to show which objects have been recently visited.

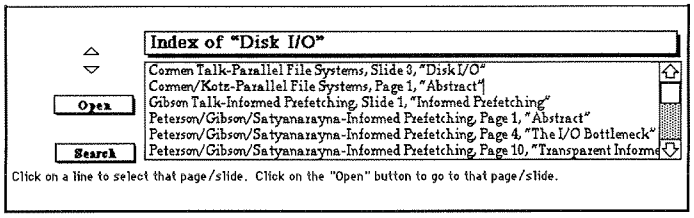

Figure 6 Presenting the results of a search as a list. The results of a search may be presented in a separate window that acts as a hypertext index. One can visit a source by selecting it. One may also scroll through sources by clicking on the up and down arrows.

There are, of course, many other mechanisms for presenting search results other than simple lists. Gloor's *Cybermap* interface [Glo91] for presenting related documents as a graph is especially promising.

3.2 Additional Features

While searching is clearly a key feature for any electronic proceedings, proceedings must include more than search features. They must include both basic and complex features that accommodate the needs of the virtual participants. Some of these features mimic the "features" a printed proceedings can provide, some are appropriate only in electronic form.

Among the most basic features in an electronic proceedings are the **standard navigation utilities**. A virtual participant should be able to quickly jump between pages: to the next page, the previous page, the first page (in current document, the current session, or the whole proceedings), or the last page. The proceedings should also include appropriate hierarchical tables of contents that link the virtual participant to the components of the proceedings and that can present the content of the proceedings in various levels of detail. For example, one might begin by looking at an overview of the titles of all the sessions in the conference, select and expand the contents of selected sessions, select and expand papers within those sessions, and select and expand sections within those papers.

Proceedings should also include some forms of **hypertext navigation**. Basic forms of hypertext navigation include a table of contents with links to the papers and a glossary that the virtual participant can access by selecting an unknown term. Proceedings may also incorporate *semantic links* to related

ideas and results. Finally, a proceedings may allow virtual participants to add and share their own links (including, possibly, links to papers not in the proceedings).

Electronic proceedings may also include more **sophisticated navigation mechanisms.** For example, a proceedings might include *paths* that trace a topic through the proceedings, stopping at individual sections of papers, talks, and notes. Figure 7 presents a simple path interface. Although paths present information in a linear order, they allow one to reorder and select material from the original proceedings. Paths also interact well with other navigation features. For example, the history of pages and slides visited in a proceedings can be treated as a path. Similarly, one might turn selected elements of a path into *bookmarks* that provide quick access to a portion of the proceedings. An electronic proceedings might also allow virtual participants to create and share their own paths. Such shared paths allow researchers to extend the proceedings and allow educators to present the materials to their students whatever order they deem appropriate.

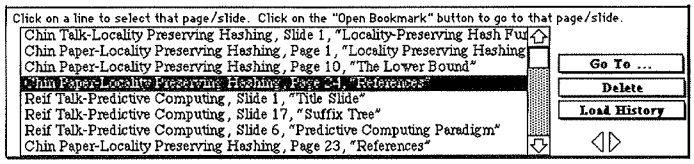

Figure 7 Creating new paths through the proceedings. Paths can be predefined by the editors of the proceedings, created by virtual participants to share with others, or derived automatically as a "history" of a session.

An additional navigation mechanism worth considering is *hierarchical navigation*, in which the document is broken up into hierarchical segments (e.g., chapters, sections, subsections, paragraphs) and, in addition to being able to move forward in the text, one can navigate "up" to the enclosing section, "down" into a subsection, or "forward" to the next segment at the same level (skipping the details in the subsections). Figure 8 presents the Gloor-Dynes hierarchical navigation tool [DG93] which provides both a "navigation diamond" that presents each of these directional features and a "context description" in which the current segment is shown in relation to the surrounding segments. Such navigation requires tat papers be segmented and annotated with their "node levels." Additionally, it presumes that authors organize their materials and write in a sufficiently hierarchical fashion that one may comprehend portions of the paper without reading all the details.

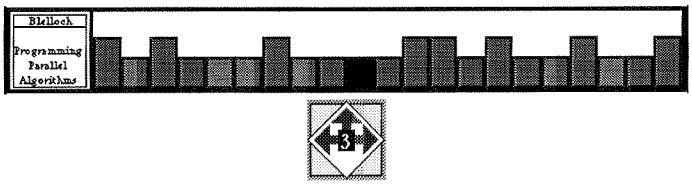

Figure 8 The Gloor-Dynes hierarchical navigation tool. The relationship of the current "node" (i.e., section) in the proceedings to the rest of the proceedings is shown by rectangles at the bottom of the screen. The navigation diamond allows one to navigate the node hierarchy, quickly jumping between nodes of the same level, returning to a higher level, or advancing into a lower level.

Because many people find it easier to read printed text, an electronic proceedings may also provide **printing facilities**. For example, a proceedings might allow the virtual participant to print a whole paper,

selected pages in the paper, the abstract of a paper, or the bibliography of a paper[8], slides from a talk, definitions from the glossary, or annotations and notes. However, there are also compelling reasons to restrict printing. In particular, if the electronic proceedings provides a networked supplement to a printed proceedings, the publishers of the printed proceedings may object to allowing virtual participants to print papers without paying royalties.

To encourage virtual participants to use the materials in new ways, and to aid those synthesizing information from the proceedings, one might include **copy and paste functions.** Such functionality would allow a researcher to quickly prepare a new presentation by copying and modifying slides from other presentations. There are also disadvantages to such functions. They might be seen as encouraging plagiarism (even though any reuse should be cited). Copy and paste may also affect creativity, as they allow one to incorporate ideas without rewriting the presentation of those ideas and therefore rethinking and reevaluating the ideas. Many authors have suggested that they do not want their text in format that allows one to copying text for both these reasons; they worry that people will use their text without citing it and they worry that those who cite their text may not understand their results.

Proceedings should also provide mechanisms for adding and sharing **annotations.** There are a variety of appropriate annotation mechanisms, many of which correspond to the types of annotation mechanisms available with printed proceedings. In particular, a reader may add *marginal notes* to pages in the proceedings (and even to slides), take notes in a separate *notebook* that is not tied to any particular portion of the proceedings, and include *bookmarks* to make it easier to return to particular parts of the proceedings. By including all three types of annotations, an electronic proceedings allows virtual participants to bring their normal work habits to the electronic realm.

By allowing virtual participants to view **multiple onscreen documents** the electronic proceedings presents the opportunity for these participants to compare ideas and results from multiple papers (or from a paper and the corresponding talk) side-by-side without jumping back and forth in the proceedings. Given these, surprisingly few electronic proceedings accommodate multiple windows.

Finally, proceedings should include a variety of **help mechanisms**, including (a) onscreen information to remind the virtual participant of what it is possible to do and how to do it, (b) context-sensitive pop-up help, that describes each object the virtual participant points to, (c) a larger online hypertext-based help system that gives more extensive documentation for the proceedings as a whole and for individual features, and (d) printed documentation.

3.3 Presenting Papers

Although interactive multimedia proceedings incorporate many materials and features in addition to the text of papers, papers remain a key component of proceedings. It is therefore especially important that the papers and the features for manipulating papers be presented coherently. When putting a paper on the screen, one must consider the relationship between the text of the paper, information about the paper, and the basic and advanced features used for manipulating the proceedings. Figures 9 and 10 demonstrate two possible designs for an interface to papers in the proceedings.

It is important to always present the context of the paper along with the content of the paper. In particular, the author and title should be available on every "page" of the electronic paper, usually at the top or the bottom. In addition, the interface can present "positional" information on the current page (the relation of the current page to the complete paper). This information may be presented as an explicit page number, as a progress bar, or using hierarchical context (as in the Gloor-Dynes interface).

Features can be presented in a variety of ways, including buttons on the margins of the window, menu items, and function keys. When possible, related functions should be grouped together. Some functions can have multiple interfaces, but care should be taken to ensure that the use is not overwhelmed with related redundant possibilities.

[8]Or a combined bibliography from multiple selected papers.

The interface to a proceedings should avoid including many similar but distinct features. For example, there should only be one way to go "forward" in the proceedings (to continue reading after the current page). Proceedings that present full printed pages on screen fall prey to this problem; to continue reading in a document in such proceedings, one might scroll the page down, scroll the page up to move to the second column, or switch to a new page, each of which requires a separate gesture. The Gloor-Dynes Hypertext interface shown in figure 9 also succumbs to this problem: it includes a scroll bar to continue reading text in the current node,[9] a standard right arrow to move to the following node, and a right arrow in the navigation diamond to move to the next node at the same level (which may be distinct from the following node).

Figure 9 presents the Gloor-Dynes Hypertext interface [DG93], which was used as the basis [GMM93]. In this interface, the main text is presented at the center of the screen and links in the text (e.g., "Equation 5") are indicated with a bold text style. The title and author of the paper appear at the bottom of the screen and the name of the current node appears at the top of the screen. The context of the current node is illustrated by the rectangles at the bottom of the screen. The buttons on the left margin provide help, a map of the proceedings, a table of contents, a collection of modifiable bookmarks, a log (notebook), and marginal notes. The marginal notes button is highlighted when marginal notes for the current node are available. The buttons on the right margin provide links to related topics (added by editors), links defined by the user, a history of nodes visited, and a link to the corresponding part of the talk.

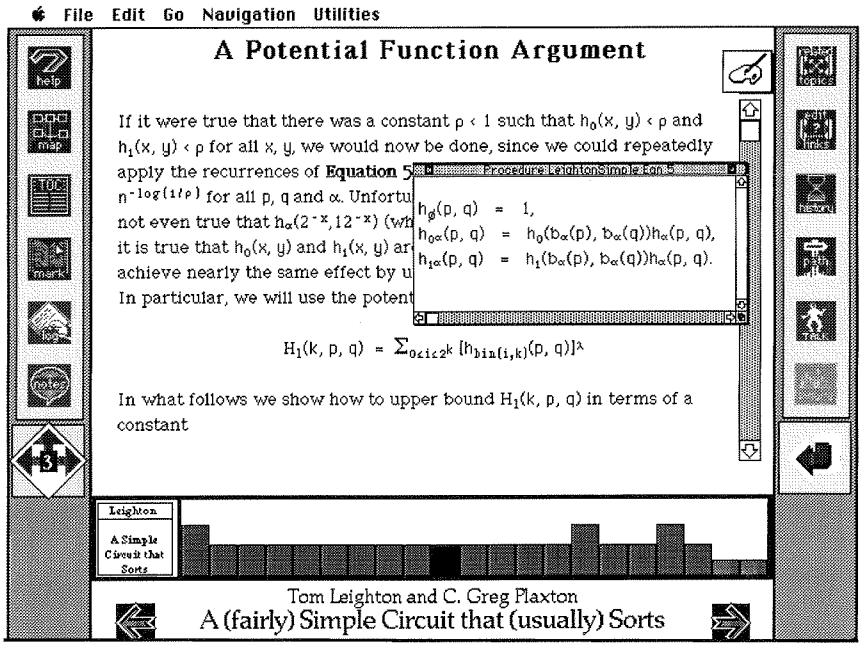

Figure 9 The Gloor-Dynes Hypertext interface. The text of the paper appears at the center of the screen. The key features are available in palettes at the side of the screen. The boldface objects in the text are links to other parts of the proceedings, to bibliography entries, to equations, or to figures.

Figure 10 presents the revised DAGS hypertext interface, which was used in [MMR94]. Again, the text is presented at the center of the screen. This new interface uses a somewhat simpler navigation mechanism; only one page is shown at a time, and one may jump to the next page using the right arrow. To move quickly to another portion of the document, one may also click on the progress bar that appears at the bottom of the screen. To better mimic the paper-based proceedings, this interface allows the virtual

[9] A "node" in the Gloor-Dynes interface is a coherent section of the document. Nodes are assigned levels in a hierarchy (higher levels are more general, lower levels more specific); this hierarchy helps guide navigation and provide context.

participant to take notes on the side of the screen. All the main features are available at the bottom of the screen. These include a guide to the interface, context-sensitive onscreen help, searching, printing, simple navigation, bookmarks, the notebook, and glossary, a path tool, and a link to the corresponding portion of the talk.

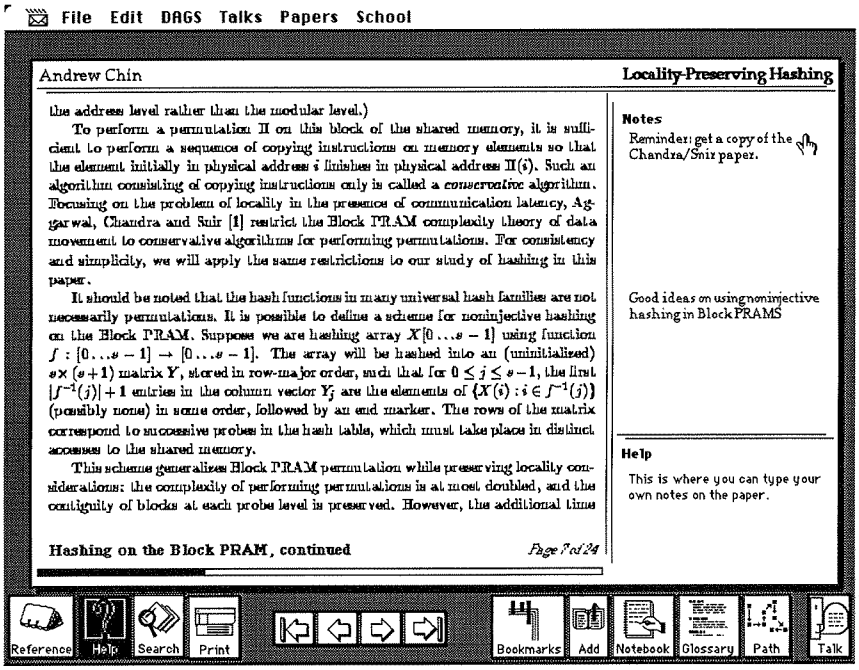

Figure 10 The revised DAGS HyperText interface. Note that the navigation mechanisms have been unified and simplified. Additionally, the marginal notes have been moved onscreen, so that they more closely resemble what one does in printed proceedings.

There are, of course, many other ways to present papers. These two interfaces illustrate some of the choices one might make when designing an interface, such has how to show a paper, what size chunks to segment a paper, what context to provide on each page, and what features to make most readily accessible.

3.4 Benefits of Sophisticated Proceedings

The interface to an interactive multimedia proceedings can extend the benefits and uses of a proceedings. In addition to the traditional proceedings roles of disseminating and preserving results, electronic proceedings can be used to encourage collaboration, introduce new fields, teach materials in a new manner, and initiate students into the craft of presentation.

Sophisticated electronic proceedings can stimulate collaboration among researchers. At one level, viewing and annotating a presentation frequently motivates discussions between members of a research group. But proceedings also encourage long-distance collaboration as researchers may share annotations, paths, links, and notes with each other. Proceedings that can "grow" further encourage collaboration as ideas that are generated by questions from the audience can be linked back into the proceedings (and the proceedings itself can have links to these new extensions).

The talks that are preserved in presented in interactive multimedia proceedings much such proceedings a potential introduction to a field. As discussed earlier, most talks include tutorial portions that are designed to accommodate newcomers to a field. Additionally, because electronic proceedings

provide search and link features, a student confused by a term in one paper or presentation may search through other papers or a glossary for further explanation of the term.

Comments from colleagues who have used our proceedings in teaching parallel computing to both undergraduates and graduate students have been very favorable. Typically, an instructor assigns the students to review a paper and talk over a period of one would. Students then have to answer several questions that are discussed in the talk or papers. Some students follow the talk first, then study the paper, and then write the homework. Others try using the searching mechanisms provided in the proceedings to identify relevant pieces of the paper and then jump into the corresponding parts of the talk and listen to them. Undergraduates tend to use the first approach, while graduate students are more likely to use the second. Both are, of course, completely acceptable.

However, this is only a simple use of electronic proceedings; multimedia proceedings provide many new tools for teaching and learning. Educators wishing to use the electronic proceedings to convey the information contained in the papers and talks can take advantage of the sophisticated features not found in more traditional proceedings or educational materials. One advantage that electronic proceedings have is their strong searching capabilities. Effective use of these searching tools allows the educator to pinpoint portions of the proceedings that concern a specific issue or topic, across the entirety of the proceedings. Once the educator has located a cohesive set of information to present, additional notes can be typed in as annotations, and a comprehensive set of notes can be created in a separate notebook. Thus, the educator can create an overall lecture or study plan, and provide thought-provoking questions and comments in conjunction with the source material.

Among the most innovative tools a sophisticated electronic proceedings can provide is the ability to create a logical path through the proceedings. With a path feature, an educator can create a path that highlights a topic in the proceedings, add annotations next to the material, and provide notes to guide the student's overall experience. The student can then navigate through the entire proceedings, viewing portions of talks and sections of papers that relate with a particular topic, rather than having to absorb the entire proceedings.

Since these annotations and paths may be easily shared, all the educator must do is provide access to the CD-ROM or network host on which the proceedings reside, and a copy of the notes, path, and annotations files. If other educators have access to the same electronics proceedings, these electronic lectures can be shared for further use or additional comment.

Because comprehensive multimedia proceedings include slides, papers, presentations, and notes, they may also be used to teach new scientists how to present (and how not to present). To a novice, the tasks of turning an idea into a paper, turning a paper into a presentation, and of making a presentation, can be particularly intimidating and confusing. Because an electronic proceedings includes links between papers and talks, one may learn how more experienced scientists turned papers into talks. The audio and video can teach about presentation techniques such as how to use transparencies and overlays, when pauses are appropriate, how and when to present summaries. The proceedings can also help answer basic questions that novices often ask, such as "how many slides does one need for a half-hour talk?"

4 Using Video in Electronic Proceedings

As indicated above, one of the most compelling reasons to include video in an electronic conference proceedings is to better reproduce the original conference. An electronic proceedings with video allows the virtual participant to experience the presentations at the conference "as they happened" with the accompanying gestures, facial expressions, and the many other things that make a talk more than a recitation of a paper.

Although there are many benefits to including video, it should be simply "dropped into" an electronic conference proceedings. Unless care is taken in the design and editing of the audio and video from presentations, they serve as little more than computerized video tapes. However, a well-edited and annotated presentation with video and audio greatly enhances both the talk and the corresponding paper.

Such editing may require significant time investment from either the author or from an editor who understands the field well.

Video can be quite expensive to include in proceedings: it requires significant disk storage. While a proceedings with text and with audio and slides from selected talks can easily fit on a single CD-ROM,[10] a proceedings that incorporates more than a small amount of video will require many more CD-ROM disks, which will complicate the use and searching of the proceedings. Additionally, the processing power required to uncompress compressed video, synchronize audio and video, and show digital video may also adversely affect the performance of the proceedings. While many computers have the power and storage necessary to display digital video, and many more will soon gain those capabilities, the average desktop computer is limited in the amount of video it can show in conjunction with audio and sequenced graphics. Because of those limitations, one may choose to use video loops or still images during talks. One might also include short video segments, both formal and informal.

Even this restricted video can benefit electronic proceedings. A short video loop or key video segment can add to the comfort level of the proceedings and can give a better sense of the researcher presenting his or her results. This last benefit is especially important for attracting a broader variety of people to the proceedings and to the field. A common complaint about science is that it is too cold and inhuman. Video from the conference (including informal videos) shows otherwise.

However, there are drawbacks to video loops. As suggested earlier, a key benefit to video is that it can show the gestures and expressions that cannot be adequately captured or conveyed by audio and text. A video loop does not present these gestures and expressions. We are exploring a hybrid mechanism that incorporates the benefits of video loops but also presents the key gestures and expressions. When the speaker is talking in his or her normal fashion, a video loop is used. When the speaker does something "special," a separate clip is used. This hybrid interface requires a significant amount of expert work to identify the separate clips to use and when to use them.

4.1 Using Video in the Interface to Presentations

There are also many design decisions to make when presenting video, audio, and slides from conference talks. Among other things, one must consider how to segment the talk, how to provide access to the segments of the talk, how to control the talk, how to synchronize the segments, and what features to provide. Because people are used to passively watching video, one may use a relatively simple interface to the talk, with simple controls to start and stop the talk and to jump to key points. One may also incorporate more advanced features to encourage the virtual participant to interact more actively with the talk.

Figure 11 presents a simple interface for presenting talks. This interface was used in [GMM93]. In this interface, the slides from the talk appear to the left of the screen and the video of the speaker appears in the upper-right-hand corner. The "Jump To Highlight" button allows one to quickly navigate to a key point in the talk. The scroll bar at the bottom of the slide window also allows quick scanning through the talk. Because of limitations in speed on the desktop machines used to view these proceedings, the video of the speaker is treated as optional, with the slides and audio tied together into a "movie." Rather than spend processing power synchronizing the video to the slides and audio, this interface uses a short (30 second) video loop of the speaker.

Reactions to the video loop have been mixed. Some virtual participants have said that they appreciate the video of the speaker for the comfort factor and because it gives a sense of how that speaker presents materials. These commentators say that they don't mind the lower-quality video on low-end machines. They also say that they mostly pay attention to the slides, and the moving figure need not be synchronized nor have a high frame rate. Others have suggested that the video is not worth including because of the low quality and because of the effect it has on the audio and slides. To these virtual participants, it is only worth including video if that video is of sufficiently high quality.

[10]For example, [GMM93] included the text of 22 papers and the audio and slides from eight one-hour presentations on one CD-ROM.

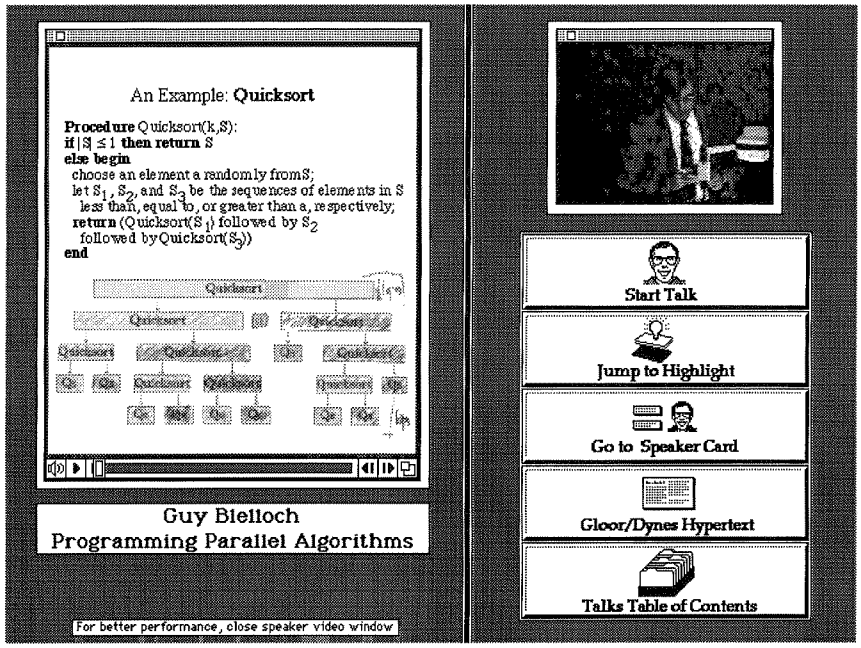

Figure 11 A simple multimedia talk interface. The video of the speaker is presented in the upper-right-hand corner while the slides appear to the left. The audio and slides are synchronized and the video is presented as a loop. The "Jump to Highlight" button allows one to quickly move around in the talk.

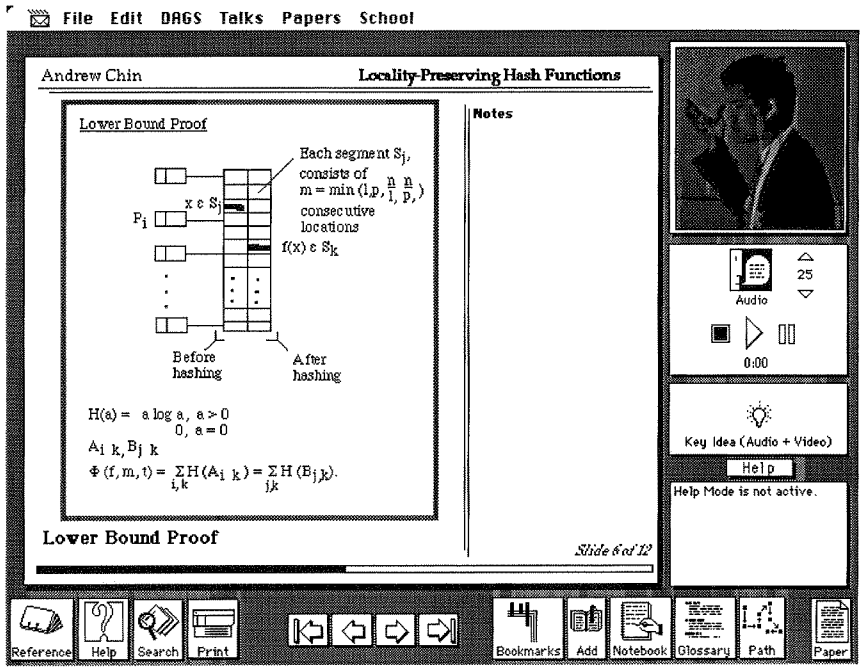

Figure 12 A more elaborate interface for presenting multimedia talks. This new interface more closely mimics the paper interface. In particular, it gives the virtual participant the opportunity to annotate slides, to add them to paths, and to make bookmarks on slides. In addition, it presents standard, "VCR-like" controls that allow the virtual participant to start, stop, and pause the talk at any point, jump to another slide, and restart the talk.

Figure 12 presents a more complex interface for talks designed to make the talk component of the proceedings more closely resemble the paper component. This interface translates many of the features of figure 10 to the domain of talks. It allows virtual participants to search for slides, put bookmarks on slides, add notes on slides, and include slides in a path through the proceedings.[11] Because of the problems of the video loop described above, no video of the speaker is included in this portion of the interface. Rather, a short segment of speaker, slides, and audio is included separately. This short segment allows the virtual participant to get a sense of the speaker (and, when the segment is chosen appropriately, to hear a summary of the key idea in the presentation), but does not interfere with the main track of the talk (the slides and audio).

There are, of course, many other components that could be added to these interfaces. At present, they do not present notes from the speakers or transcriptions of the talks. A transcription might allow one to search for a particular piece of text and then start the talk with that piece of text. Additionally, if the proceedings includes way to create new presentations by creating new paths through the talks, then it might also include a way to create and include new slides. Finally, it might be extended by animated slides or animations of algorithms described in the talk.

4.2 Simulating Video Tracks from Presentations

Because video loops do not provide a sufficient record of the conference (as they miss gestures and expressions) and because full video is still infeasible in most situations, it is preferable to create hybrid video tracks that use a video loop most of the time, but employ "key segments" when the particular motions or gestures of the speaker are significant (or out of the ordinary). We are presently looking at ways to produce and present such hybrid tracks.

The creation of hybrid video tracks can be accomplished with traditional editing tools, but it is a time consuming manual process. We are investigating the use of VideoScheme, a programmable video editing environment, to augment this process with automatic content analysis [MGM93]. The key segments are likely to stand out from their surroundings in terms of raw image continuity, object motion, camera motion, and audio volume. All of these can be computed and used as a basis for extracting candidate key segments.

Content analysis can also be used to provide more sophisticated searching and navigation interfaces for complete video tracks, by highlighting transition points and other times of likely interest.

4.3 Other Video Components

Video from presentations is only one of the many types of video that a multimedia conference proceedings can include. There are a variety of shorter pieces of video that can make conference proceedings more useful and enjoyable. Some video clips present results in additional ways; others help reproduce the environment and dynamic nature of a conference. Video clips that may be incorporated into multimedia proceedings include:

- **Key segments from presentations.** In addition to using the whole video from a presentation or a video loop, one might also choose one or two selected segments that summarize the results or illustrate a key point in the presentation. Such clips can serve as introductions to the presentation. They also provide an informal summary of the speaker's style and mannerisms.

[11] Because both slides from talks and pages from papers may be included in paths, one may produce much more sophisticated paths through the proceedings. For example, one might jump from an explanation in a paper to a corresponding figure from a talk. Additionally, one may create new presentations by creating a new path through selected slides in the proceedings.

- **Separately recorded summaries and interviews.** Presenters may also be interviewed and asked to summarize their results. As with the key segments discussed previously, such summaries serve as an introduction and overview to the presentation and paper.
- **Prepared videos.** Many authors already prepare videos that illustrate their results. Some videos are shown as part of the talk and others are shown separately. If authors know that such videos can be incorporated into the proceedings, they will be more likely to prepare these videos.
- **Conference overview.** The organizers of the conference may provide a short "guided" tour to the conference. Additionally, they may provide introductions for the individual talks.
- **Informal videos.** During the conference, one may record informal sessions and conference social events. Such videos show additional aspects of a conference that are not traditionally included in a proceedings and better involve the virtual participant in the conference.

While there are benefits to these shorter video segments, like all components of an electronic proceedings, they should be incorporated carefully and with thought to their relationship to the complete proceedings. For example, a video clip should be annotated with links to the materials that present the concepts it illustrates and to other videos that illustrate related concepts.

5 Summary and Future Work

The construction of electronic conference proceedings is an emerging technology, and not all the issues surrounding the design and development of such proceedings have been worked out. Even text-only electronic proceedings have potential for many new features and designs. Given that very few electronic proceedings have included video, and fewer still have included video from presentations, there is still much work to do on the design and implementation of interfaces for interactive multimedia conference proceedings.

Even though electronic proceedings do not yet have a stable interface that one can simply plug materials into, it is possible to create relatively sophisticated electronic proceedings for both small and large conferences. Such proceedings do require significant work from the editors and engineers of the proceedings, but the time spent has many benefits. Work spent designing and building electronic proceedings is repaid by the broader audience the proceedings accommodates and attracts and by the lessons learned that can be applied to a variety of multimedia projects and that influence the development of future proceedings. Additionally, the more proceedings are available electronically, the closer the academic community will come to building and using standard proceedings interfaces.

To ease the development of such proceedings, we are developing a standardized interface and features for electronic proceedings, as well as recommended protocols for authors and editors to follow. However, many of the needs of the developers of proceedings can and should be met by traditional multimedia authoring systems. We call on the creators of multimedia authoring systems to provide better support for significant amounts of text (including support for sophisticated searching mechanisms), for sophisticated links, and for the many features we have described above.

We are presently exploring many extensions to and variations of electronic conference proceedings, many of which are mentioned in the sections above. We have begun to investigate the use of VideoScheme [MGM93] to produce more sophisticated video loops and to automate the basic editing of audio and video. We are also investigating the use of the techniques for building an interactive multimedia conference proceedings to present an interactive multimedia class (as classes share many things in common with conference proceedings, particularly links between the video of presentations and the text of the course).

Acknowledgments

Many people contributed to the development of the proceedings for DAGS'92 and DAGS'93. These contributions influenced both the content and interface of the proceedings and led to many of the ideas presented in this paper. Contributors include: Chris Broggi, Michelle Chen, Matthew Cheyney, Mike Chong, Jennifer Collins, Thomas Cormen, Scott Dynes, Robert Frost, Peter Gloor, Mobina Hashmi, Donald Johnson, David Kotz, Naval Ravikant, Nii Sai Sai, Clifford Stein, Roberto Taboada, Chen Yang, Le Ye, Jun Zhang, Qin Zhang, Yihao Zhang, and Kanyi Zhao.

The DAGS institutes and the preparation of electronic proceedings for the DAGS institutes were supported by a variety of institutions, including the National Science Foundation (NSF grants 5-34251, 5-34294, 5-34332), the New England Consortium for Undergraduate Science Education (NECUSE), the Dartmouth Institute for Advanced Graduate Studies, and the Dartmouth College Presidential Scholar Fund.

Bibliography

[AGB93] J. Alspector, R. Goodman, and T.X. Brown (editors). *Proceedings of the International Workshop on Applications of Neural Networks to Telecommunications.* Erlbaum: Hillsdale, NY, 1993.

[All94] Robert B. Allen. "Electronic Proceedings (Eprocs) for IWANNT'93." *Human Factors in Computing Systems: CHI'94 Conference Companion*, p. 287. ACM Press: New York, NY, 1994.

[Buf94] John F. Koegel Buford (contributing editor). *Multimedia Systems.* ACM Press: New York, NY, 1994.

[Bus45] Vannevar Bush. "As We May Think," *Atlantic Monthly*, 1945.

[CR93] A. Califano and I. Rigoutsos. "FLASH: A fast look-up algorithm for string homology." *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, New York City, June 1933.

[CGJ+94] Matthew Cheyney, Peter A. Gloor, Donald B. Johnson, Fillia Makedon, James W. Matthews, and P. Metaxas. "Conference on a Disk: A Successful Experiment in Hypermedia Publishing (Extended Abstract)." *Educational Multimedia and HyperMedia, 1994.* AACE: Charlottesville, VA, 1994.

[DDF+90] Scott Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. "Indexing by latent semantic analysis." *J. Amer. Soc. Information Science* 41(6), (1990) 391-407.

[DG93] Scott B. Dynes and Peter A. Gloor. "Using hierarchical knowledge representation for an animated algorithm learning environment." MIT LCS/TNS Technical Report, 1993.

[Dum90] Susan Dumais. "Enhancing Performance in Latent Semantic Indexing (LSI) Retrieval." TM-ARH-017527 Technical Report, Bellcore, 1990.

[[Dum91] Susan Dumais. "Improving the Retrieval of Information from External Sources. Behaviour Research Methods." *Instruments and Computers* 1991, 23 (2) 229-236.

Glo91] Peter Gloor. "Cybermap, Yet Another Way of Navigation in Hyperspace." *Proc. ACM Hypertext 91.* San Antonio, Tx, Dec. 15–18, 1991. 107–121.

[GMM93] Peter A. Gloor, Fillia Makedon, and James W. Matthews (editors). *Parallel Computation: Practical Implementation of Algorithms and Machines.* Springer Verlag, 1993.

[HvRB93] Lynda Hardman, Guido van Rossum, and Dick C.A. Bulterman. "Structured Multimedia Authoring." *ACM Multimedia '93.* Anaheim, CA, 1993.

[IEEE93] Institute of Electrical and Electronics Engineers. *SuperComputing '93.* IEEE Computer Society Press, 1994.

[MBC93+] Fillia Makedon, Guy Blelloch, Tom Cormen, George Cybenko, Phil Hatcher, Donald Johnson, David Kotz, Panagiotis Metaxas, Grammati Pantziou, Eric Schwabe, and Clifford Stein (editors). *Proceedings of the DAGS'93 Symposium on Issues and Obstacles in the Practical Implementation of Parallel Algorithms and the Use of Parallel Machines: Parallel I/O and Databases.* Hanover, NH, 1993.

[MGM93] James Matthews, Peter A. Gloor, and Fillia Makedon. "VideoScheme: A Programmable Video Editing System for Automation and Media Recognition." *ACM Multimedia '93.* Anaheim, CA, 1993.

[MMR94] Fillia Makedon, P. Takis Metaxas, and Samuel A. Rebelsky (editors). *Parallel Computation and Parallel I/O.* Forthcoming.

[MRC+94] Fillia Makedon, Samuel A. Rebelsky, Matthew Cheyney, Charles Owen, and Peter Gloor. "Issues and Obstacles with Multimedia Authoring." *Educational Multimedia and HyperMedia, 1994.* AACE: Charlottesville, VA, 1994.

[MST92] MacSciTech. *Proceedings CD-ROM of the 1992 MacSciTech Conference on Scientific and Engineering Applications for the Macintosh.* 15–17 January 1992. San Francisco, California.

[Rada93] R. Rada (prod. chair). *Proceedings CD-ROM of the First ACM International Conference on Multimedia,* 1-6 August, Anaheim, California, 1993.

[Zal94] Janusz Zalewski. Review of "Parallel Computation: Practical Implementation of Algorithms and Machines." *IEEE Computer,* July 1994.

Appendix: A Short Survey of Electronic Proceedings

Electronic proceedings are also not a new idea; for nearly as long as there have been computer networks and conference papers written on the computer, researchers have found ways to make their conferences papers available on the network. Frequently, small conferences and workshops would make all the papers from the conference available for transfer in whatever format the authors wrote the papers (for theoretical computer science, TEX and LATEX have become de facto standards). As time has passed, new interfaces for retrieving and searching papers have been implemented and a variety of proceedings have been created. In the paragraphs below, we discuss some of these proceedings, commenting on the materials the proceedings include, the dissemination mechanism and time frame, the features the proceedings provide, and lessons the proceedings teach.

Proceedings of the Chicago Workshop on Circuit Complexity

As mentioned earlier, setting up a collection of conference papers via ftp is a simple method of creating and distributing an electronic conference proceedings. This type of electronic proceedings requires the reader of the proceedings to transfer the papers to his/her local machine and search with traditional text searching program (e.g., grep). Nonetheless, it is a mechanism that most computer science researchers are comfortable with and use regularly. [12]

The proceedings from the Chicago Workshop on Circuit Complexity serve as a case study for this type of electronic proceedings. Soon after the conference, the organizers of the workshop made all the papers available via anonymous ftp. Although the proceedings of the Chicago Workshop on Circuit Complexity contain many important results and valuable materials, the maintainers of the proceedings report that these proceedings received little use. No explanation was given, although possibilities include (1) the majority of the audience for the workshop attended the workshop; (2) researchers were not sufficiently aware of the availability of these proceedings; or (3) researchers did not feel an electronic copy of the proceedings merited the effort necessary to transfer the files as they could order a printed proceedings more easily.

Key lessons learned from these proceedings are that both advertisement and more than simple ftp capability are needed for a proceedings to be used and useful.

Electronic Proceedings of DAGS 1992 Institute on Parallel Computing

The Dartmouth Institute for Advanced Graduate Studies in Parallel Computing was held in the summer of 1992 and brought together many front-line researchers in both the theory and practice of parallel computing. After the conference, both the papers and talks were tied together in a sophisticated multimedia/hypermedia system.

[12] In an informal survey of the attendees at the 1994 ACM Symposium on the Theory of Computing (STOC94), many attendees said that they would prefer to be able to simply transfer files to their own machines where they could search and print them at their leisure.

The proceedings were published by Telos/Springer-Verlag on CD-ROM. In addition to being distributed through traditional channels, copies were distributed to all attendees of the DAGS 1993 Conference on Parallel Input and Output [MBC+93].

These proceedings had two separate, but linked, components: a simple interface for viewing and browsing talks, and a sophisticated hypertext interface (the Gloor-Dynes Hypertext Engine) which provides sophisticated searching and annotation of papers. The engine also allows virtual participants to find related sections of the paper (and other papers) and to create new paths through the proceedings. The presentation interface includes slides, a video loop, and complete audio for the eight invited talks. Finally, there are many links between the talks and the corresponding papers.

Because of limitations of the source materials (many were obtained only in hardcopy format) and the restrictions of the Hypertext engine, the creation of this proceedings required a significant amount of manual labor by both novices and experts. This labor, which is discussed in [CGJ+94], included segmenting the papers, retyping formulas, and determining and adding links.

Although this proceedings included a highly-sophisticated interface, presented important results and useful introductory materials, and garnered extremely favorable reviews (e.g., [Zan94]), they have received far less use than the results and interface warrant. In part, this is because of a mismatch of technology: these proceedings are only accessible from Macintoshes, yet most parallel computing researchers do their everyday work on more traditional workstations and on parallel machines and therefore do not regularly use Macintoshes. Additionally, because these proceedings are primarily marketed for their content and sales channels offer little chance to "try out" these proceedings, those interested in the interface may not have sufficient chance to evaluate the benefits of such proceedings.

Key lessons learned from this proceedings are that the development of materials for a sophisticated materials is time consuming, that the development of a sophisticated interface in conjunction with the editing of the materials complicates matters, and that marketing issues must be carefully considered.

Electronic Proceedings of 1993 ACM Multimedia Conference

The proceedings of the 1993 ACM Multimedia Conference [Rada93] were published on CD-ROM and distributed at the conference as part of the proceedings. Because of the time-frame, the proceedings could only include materials available before the conference. These materials included the papers and videos prepared by conference presenters.

Adobe's Acrobat™ software provides the primary interface for these proceedings. Its features include onscreen viewing, searching, and printing. These proceedings do not include sophisticated annotating or linking features. Although the papers are stored in PostScript™ format, not all papers are printable on all PostScript printers. Although the proceedings were accessible from PC's, Macintosh computers, and Unix workstations, the features and capabilities were not uniform across platforms..

Key lessons learned from these proceedings include that a limited time-frame severely limits the possible features and that PostScript™ is not as uniform as one may expect.

Electronic Proceedings for IWANNT'93

Although the proceedings of the 1993 International Workshop on Applications of Neural Networks to Telecommunications were published in paper format [AGB93], they were also made available electronically, on the network.

These proceedings are presented with SuperBook™, which, at the time these proceedings were created, offered only limited support for video and audio content. Hence, these proceedings included only text. These proceedings include many of the features that one would expect in a text-based proceedings, including sophisticated searching and annotation capabilities. Because the proceedings are

networked, participants can add annotations that other participants can view and extend.[13] Like the creators of the DAGS proceedings, the creators of these proceedings learned that conversion to a uniform format requires significant effort and that authors submit their work in a wide variety of formats. Because these proceedings were text-only, they could be made available both at and after the conference.

Although many people accessed these proceedings, only 37% of those using the proceedings indicated that they would pay even a small amount for networked access to electronic proceedings.

These proceedings reinforce the lesson that unless one restricts the form and format authors may use for their documents, creation of proceedings requires significant conversion and editing effort. These proceedings also show that virtual participants are interested in collaborative annotation and that an electronic proceedings can add to a printed proceedings, rather than replacing it.

Electronic Proceedings for Supercomputing '93

The proceedings for both the 1992 and 1993 IEEE Conferences on Supercomputing were published by the IEEE some months after the Supercomputing 1993. In includes both papers and color plate images. The materials are accessible with the ViewTool retrieval software which was shipped on the CD-ROM and which works on a variety of platforms, including Unix workstations, Apple Macintosh computers, and PCs running Windows.

ViewTool provides relatively sophisticated boolean searching capabilities, allows the reader to view both ASCII and the "page" in the proceedings, and accommodates linear reading. However, neither the ASCII nor the onscreen page are completely readable; the electronic proceedings presupposes a printed copy. The electronic proceedings also maintains close ties to the printed proceedings: all references correspond to the printed proceedings (including column number).

Because of problems with the initial version of these proceedings, the IEEE created a second version.

Key lessons learned from these proceedings include that it is possible to produce an electronic proceedings for a large conference in a reasonable time frame.

[13]The editors of these electronic proceedings have limited such footnotes to the attendees at the workshop.