



GRiNS: A GRaphical INTERface for creating and playing SMIL documents

Dick C.A. Bulterman^{*,1}, Lynda Hardman¹, Jack Jansen¹, K. Sjoerd Mullender¹,
Lloyd Rutledge¹

CWI, Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands

Abstract

The W3C working group on synchronized multimedia has developed a language for Web-based Multimedia presentations called SMIL: the Synchronized Multimedia Integration Language. This paper presents GRiNS, an authoring and presentation environment that can be used to create SMIL-compliant documents and to play SMIL documents created with GRiNS or by hand. © 1998 Published by Elsevier Science B.V. All rights reserved.

Keywords: SMIL; Multimedia; Web authoring; Streaming applications

1. Introduction

While the World Wide Web is generally seen as *the* embodiment of the infrastructure of today's information age, it currently cannot handle documents containing continuous media such as audio and video very elegantly. HTML documents cannot express the synchronization primitives required to coordinate independent pieces of time-based data, and the HTTP protocol cannot provide the streamed delivery of time-based media objects required for continuous media data. The development of Java extensions to HTML, known as *Dynamic HTML* [5], provide one approach to introducing synchronization support into Web documents. This approach has the advantage that the author is given all of the control offered by a programming language in defining in-

teractions within a document; this is similar to the use of the scripting language Lingo in CD-ROM authoring packages like Director [6]. It has the disadvantage that defining even simple synchronization relationships becomes a relatively difficult task for the vast majority of Web users who have little or no programming skills.

In early 1997, SYMM — a W3C working group on SYNchronized MultiMedia — was established to study the definition of a declarative multimedia format for the Web [8,12]. In such a format, the control interactions required for multimedia applications are encoded in a text file as a structured set of object relations. A declarative specification is often easier to edit and maintain than a program-based specification, and it can potentially provide a greater degree of accessibility to the network infrastructure by reducing the amount of programming required for creating any particular presentation. The first system to propose such a format was CMIF [2–4]. Other more recent examples are MADEUS [7] and RTSL

* Corresponding author.

¹ E-mail: {Dick.Bulterman,Lynda.Hardman,Jack.Jansen,Sjoerd.Mullender,Lloyd.Rutledge}@cwi.nl




Fig. 1. Media objects for use in (part of) a Web-based newscast.

[9]. The SYMM working group ultimately developed SMIL, the Synchronized Multimedia Integration Language [13] (which is pronounced *smile*).² In developing SMIL, the W3C SYMM group has restricted its attention to the development of the base language, without specifying any particular playback or authoring environment functionality.

This paper presents a new authoring and runtime environment called GRiNS: a GRaphical INterface for SMIL. GRiNS consists of an authoring interface and a runtime player which can be used together (or separately) to create/play SMIL-compliant hypermedia documents. The authoring part of GRiNS can be used to encode SMIL presentations for any SMIL compliant browser or stand-alone player; the GRiNS player can take any SMIL-compliant document and render it using a stand-alone player. GRiNS is based largely on earlier experience with the CMIFed authoring system and the CMIF encoding format, both of which strongly influenced the development of SMIL.

In Section 3, we give an overview of the GRiNS authoring and playback interfaces. Section 4 then discusses the availability of the GRiNS environment as part of the CHAMELEON suite of authoring and presentation tools. We begin the paper with a review of the SMIL language and a brief discussion of typical applications and runtime environments for which SMIL was developed.

2. Declarative Web-based hypermedia

This section briefly reviews the nature of Web-based multi-/hypermedia documents and the basic

philosophy of the SMIL language. We also briefly sketch the type of runtime protocol support that is expected for processing SMIL documents.

2.1. A typical example SMIL document

In order to focus our attention on the class of applications that this paper concerns itself with, we present an example of a generic Web application: that of a network newscast.³ The basic premise of this newscast is a story on the explosive growth of the WWW. Several media objects have been defined that together make-up such a newscast. (Note that the selection of objects is arbitrary; we have selected a moderate level of complexity to illustrate the features of GRiNS.)

Let us assume that the objects shown in Fig. 1 have been selected to make up the presentation. The newscast consists of an opening segment, a segment in which the local host gives background information, a segment in which a remote reporter gives an update, a segment in which the local host gives a wrap up, which leads into the trailing theme music. Finally, each correspondent also has a Web homepage that can be accessed from the story.

Figure 2 shows two views of the newscast example, taken at different times in the presentation. Figure 2a shows a portion of the introduction of a story on the growth of the World Wide Web. In this portion, the local host is describing how sales of authoring software are expected to rise sharply in the next six months.

Figure 2b shows a point later in the presentation, when the host is chatting with a remote correspon-

² As SMIL-V1.0 is a W3C Draft Standard, its details may change at any time up to its acceptance.

³ In 1991, as part of the first public CMIF paper, a network newscast was also used as a prototypical example. Sometimes the electronic world moves as slowly as the real one!

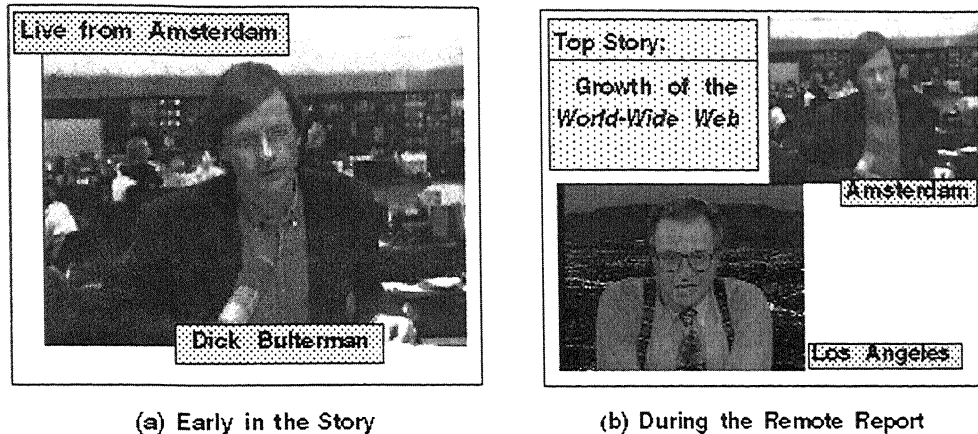


Fig. 2. Two views of the Web newscast.




Fig. 3. Augmenting information during the news.

dent in Los Angeles, who describes that many Hollywood stars are already planning their own SMIL pages on the Web.

The ability to link to various homepages makes the semantic content of the document dynamic. As shown in Fig. 3, the information content can be augmented during the story depending on the viewer's interests. Links are not restricted to homepages, of course: any addressable object or document deemed relevant by the author should be available. Such behaviour comes at a cost, however: we must be able to specify what happens to the base presentation when the link to the homepage is followed: should it pause, should it continue, or should it be replaced by the linked pages. (We discuss this in more detail below.)

A complete description of the implementation details of the Web newscast is beyond the scope of this paper. We will, however, refer to it as a running example in the sections below.

2.2. A brief SMIL overview

SMIL is a declarative language for describing Web-based multimedia documents that can be played on a wide range of SMIL browsers. Such browsers may be stand-alone presentation systems that are tailored to a particular user community or they could be integrated into general purpose browsers. In SMIL-V1.0 (the version that we will consider in this paper), language primitives have been defined that allow for early experimentation and (relatively) easy implementation; this has been done to gain experience with the language while the protocol infrastructure required to optimally support SMIL-type documents is being developed and deployed.

Architecturally, SMIL is an integrating format. That is, it does not describe the contents of any part of a hypermedia presentation, but rather describes how various components are to be combined temporally and spatially to create a presentation. (It also defines how presentation resources can be managed and it allows anchors and links to be defined within the presentation.) Note that SMIL is not a replacement for individual formats (such as HTML for text, AIFF for audio or MPEG for video); it takes information objects encoded using these formats and combines them into a presentation.

Appendix A contains a simplified SMIL description of the newscast example sketched in Section 2.1. SMIL describes four fundamental aspects of a multimedia presentation:

- *temporal specifications*: primitives to encode the temporal structure of the application and the refinement of the (relative) start and end times of events;
- *spatial specifications*: primitives provided to support simple document layout;
- *alternative behaviour specification*: primitives to express the various optional encodings within a document based on systems or user requirements; and
- *hypermedia support*: mechanisms for linking parts of a presentation.

We describe how SMIL encodes these specifications in the following paragraphs.

2.2.1. SMIL temporal specifications

SMIL provides coarse-grain and fine-grain declarative temporal structuring of an application. These are placed between the `<body> ... </body>` tags of a document (lines 15 through 52 of Appendix A). SMIL also provides the general attribute `SYNC` (currently with values `HARD` and `SOFT`) that can be used over a whole document or a document part to indicate how strictly synchronization must be supported by the player. Line 1 of the example says, in effect: do your best to meet the specification, but keep going even if not all relationships can be met. Coarse grain temporal information is given in terms of two structuring elements:

- `<seq> ... </seq>`: A set of objects that occur in sequence (e.g., lines 16–51 and 31–43).
- `<par> ... </par>`: A collection of objects that occur in parallel (lines 17–24 or 36–44).

Elements defined within a `<seq>` group have the semantics that a successor is guaranteed to start after the completion of a predecessor element. Elements within a `<par>` group have the semantics that, by default, they all start at the same time. Once started, all elements are active for the time determined by their encoding or for an explicitly defined duration. Elements within a `<par>` group can also be defined to end at the same time, either based on the length of the longest or shortest component or on the end time of an explicit master element. Note that if objects

within a `<par>` group are of unequal length, they will either start or end at different times, depending on the attributes of the group.

Fine grain synchronization control is specified in each of the object references through a number of timing control relationships:

- *explicit durations*: a `DUR="length"` attribute can be used to state the presentation time of the object (line 34);
- *absolute offsets*: the start time of an object can be given as an absolute offset from the start time of the enclosing structural element by using a `BEGIN="time"` attribute (line 33);
- *relative offsets*: the start time of an object can be given in terms of the start time of another sibling object using a `BEGIN="object_id+time"` attribute (line 41).

(Unless otherwise specified, all objects are displayed for their implicit durations — defined by the object encoding or the length of the enclosing `<par>` group.) The specification of a relative start time is a restricted version of CMIF's *sync_arcs* [4] to define fine-grain timing within a document. At present, only explicit time offsets into objects are supported, but a natural extension is to allow content markers, which provide content-based tags into a media object.

2.2.2. Layout specifications

In order to guarantee interoperability among various players, SMIL-V1.0 supports basic primitives for layout control that must be supported on all SMIL players. This structure uses an indirect format, in which each media object reference contains the name of an associated drawing area that describes where objects are to be presented. A separate layout resolution section in the SMIL `<head>` section maps these areas to output resources (screen space or audio). An example of a visible object reference is shown in line 29, which is resolved by the definition in line 6. Non-visible objects (such as audio) can also be defined (see lines 33 and 7). Each player/browser is responsible for mapping the logical output areas to physical devices. A priority attribute is being considered to aid in resolving conflicts during resource allocation. (As of this writing the name channel will probably be used as an abstract grouping mechanism [2] but the deadline for text submission came before this issue was resolved.)

2.2.3. Alternate behaviour specifications

SMIL provides two means for defining alternate behaviour within a document: via the `<switch>` construct and via the `<channel>`. The `<switch>` allows an author to specify a number of semantically equivalent encodings, one of which can be selected at runtime by the player/browser. This selection could take place based on profiles, user preference, or environmental characteristics. In lines 19–22 of Appendix A, a specification is given that says: play either the video or still image defined, depending on system or presentation constraints active at runtime. The `<channel>` statement is intended to provide a higher-level selection abstraction, in which presentation binding is delayed until runtime. It is a simplified version of the *shadow channel* concept in CMIF.

2.2.4. Hypermedia and SMIL

HTML supports hypertext linking functionality via a straight-forward process: each document has a single focus (the browser window or frame) and anchors and links can be easily placed within the document text. When the link is followed, the source text is replaced by the destination. In an SMIL presentation, the situation is much more difficult. First, the location of a given anchor may move over time as the entity associated with the anchor moves (such as following a bouncing ball in a video) — and even if it does not move, it still may be visible for only part of the object's duration. Second, since SMIL is an integrating format, conflicts may arise on ownership of anchors and the semantics of following any given link.

In SMIL-V1.0, a pragmatic approach to linking is followed. Anchors and links within media objects are followed within the context of that media object. An anchor can also be defined at the SMIL level, which affects the presentation of the whole SMIL document. This effect depends on the value of a `SHOW` attribute, which may have values:

- **REPLACE**: the presentation of the destination resource replaces the complete, current presentation (this is the default);
- **NEW**: the presentation of the destination resource starts in a new context (perhaps a new window) not affecting the source presentation; or
- **PAUSE**: the link is followed and a new context is created, but the source context is suspended.

In our example, two links are defined: one on lines 28–30 and the other on lines 37–39. If the link associated with line 29 is followed, the current presentation's audio and video streams keep going, and a new SMIL presentation containing an HTML homepage is added. If the link at line 38 is followed, the current video object pauses while the homepage is shown.

General support for hypermedia is a complex task. Readers are invited to study the Amsterdam Hypermedia Model [3], which serves as the basis for the current SMIL hypermedia proposal.

3. GRiNS: authoring and presentation for SMIL

GRiNS is an authoring and presentation system for SMIL documents. It is a part of the CHAMELEON multimedia document processing suite. Where the CHAMELEON suite provides tools for the authoring of adaptive documents — that is, documents that can be converted to a variety of presentation formats based on the (dynamic) needs of the document authors — GRiNS provides dedicated tools for creating and presenting SMIL documents. Both GRiNS and CHAMELEON are based on technology developed for CWI's CMIF environment.

The GRiNS authoring tool is based on a structured approach to document creation [2]. The GRiNS presentation tool provides a flexible presentation interface that supports all of the semantics of the SMIL V1.0 specification. In the following sections, we give examples of how GRiNS can be used to create the WebNews example. We then discuss some of the issues involved in providing support for the GRiNS presenter.

3.1. The GRiNS authoring environment

The GRiNS authoring environment supports creation of a SMIL document in terms of three views:

- the *logical structure* view, where coarse-grain definition of the temporal structure takes place;
- the *virtual timeline* view, where fine-grain temporal interactions are defined; and
- the *playout view*, where layout projections are defined.




Fig. 4. The structure of the Web Growth story. (Time flows from top to bottom.)

Hyperlink definition and specification of alternative data objects can occur in either the logical structure view or the virtual timeline view.

3.1.1. The logical structure view

If we were to define the Web Growth story in terms of its overall structure, we would wind up with a representation similar to that in Fig. 4. Here we see that the story starts with an opening sequence (containing a logo and a title), and ends with a closing jingle. In between is the “meat” of the story. It contains an introduction by the local host, followed by a report by the remote correspondent, and then concluded by a wrap-up by the local host. This “table of contents” view defines the basic structure of the story. It may be reusable.

The GRiNS logical structure view allows an author to construct a SMIL document in terms of a similar nested hierarchy of media objects. The interface provides a scalable view of the document being created. Note that only structure is being edited here: much of this creation takes place before (or while) actual media objects have been created.

Figure 5a shows a part of the WebNews hierarchy in terms of the logical structure view. (The green box in the middle shows two alternatives in the presentation: a video or a still image, one of which will

be selected at runtime. Figure 5b shows a typical dialog box that is used to provide details of how a particular element is to be included in the presentation. The logical structure view has facilities cutting and pasting parts of the presentation, and it allows individual or sub-structured objects to be previewed without having to play the entire application.

During design and specification, the values of object attributes — including location, default or express duration or synchronization on composite objects — can be entered by the author. In practice, duration of an object or a group will be based on the enclosing structure, will which be calculated automatically. Note that while the basic paradigm of the logical structure view is a hierarchical structure, the author can also specify loop constructs which give (sub-)parts of the presentation a cyclic character.

3.1.2. The virtual timeline view

The logical structure view is intended to represent the coarse timing relationships reflected in the `<par>` and `<seq>` constructs. While the attributes associated with an element (either an object or a composite structure definition) can be used to define more fine-grained relationships — such as DURATION or the desire to REPEAT an object during its activation period — these relationships are often difficult to visualize with only a logical structure view. For this reason, GRiNS also supports a timeline projection of an application. Unlike other timeline systems, which use a timeline as the initial view of the application, the GRiNS timeline is virtual: it displays the logical timing relationships as calculated from the logical structure view. This means that the user is not tied to a particular clock or frame rate, or to a particular architecture. (The actual timing relationships will only be known at execution time.) Note that the virtual timeline view is a generated projection, rather than a direct authoring interface. It is used as a visualization aid, since it reflects the view of the GRiNS scheduler on the behaviour of the document.

Figure 6 shows a virtual timeline of the Web Growth story. Rather than illustrating structure, this view shows each of the components and their relative start and end times. The timeline view provides an insight into the actual temporal relationships among the elements in the presentation. Not only does it




Fig. 5. The hierarchy view (a) and attributes of a node (b).




Fig. 6. The virtual timeline view of the WebNews. (Time flows from left to right.)

show the declared length and start times of objects whose duration or start offset is predefined, it also shows the derived length of objects whose duration depends on the structure of the document. The timing view shows predefined durations as solid blocks and derived durations as blocks with a diagonal line.

Events on the timeline are partitioned by the various output channels defined by the user. In the figure, channels which are darkened are turned off during this (part of) the presentation. This allows the author to see the effect of various <switch> settings during the presentation.

When working with the virtual timeline view, the user can define exact start and end offsets within `<par>` groups using declarative mechanisms (the sync arc), shown as an arrow in the figure. Sync arcs are meant to provide declarative specifications of timing relationships which can be evaluated at run-time or by a scheduling pre-processor.

If changes are made to the application or to any of the attributes of the media objects, these are immediately reflected by in the virtual timeline view. As with the logical structure view, the user can select any group of objects and preview that part of the document. When the presentation view is active (see next section), the virtual timeline view also displays how the playout scheduler activates, arms and pre-fetches data during the presentation.

Note that both the logical structure and the virtual timeline views are authoring views; they are not available when a document is viewed via the playout engine only. Both views isolate the user from the syntax of the SMIL language.

3.1.3. The presentation view

The presentation view has two purposes: first, during authoring, it provides a WYSIWYG view of the document under development, and second, it provides a mechanism to interactively layout the output channels associated with the document. Figure 7a shows part of the presentation being developed, while Fig. 7b illustrates the channel map associated with the part of the document. During playout,

the values of the layout attributes can be dynamically changed if required.

3.2. The GRiNS playout engine

The GRiNS authoring interface is actually an augmented GRiNS playout engine. It is a playout engine that also provides a structure and timeline editor. GRiNS also provides a “stand-alone” playout engine which can play any compliant SMIL-V1.0 document.

The playout engine consists of a document interpreter that, guided by user interaction, steps through a SMIL document. The document interacts with a link processor (which implements the NEW, PAUSE, REPLACE semantics of the link), and a channel scheduler which builds a Petri-net based representation of the document. Actual media display is handled by a set of channel threads (or their architectural equivalent on various platforms), which interact with presentation devices. The channel scheduler and the various channel threads implement the application and media dependent parts of the RTP/RTCP protocol stacks.

As part of the general SMIL language, an author can specify alternative objects that can be scheduled based on resources available at the player, the server or the network. The GRiNS player also allows users to explicitly state which channels they wish to have active. This is useful for supporting multi-lingual applications or applications in which users are given

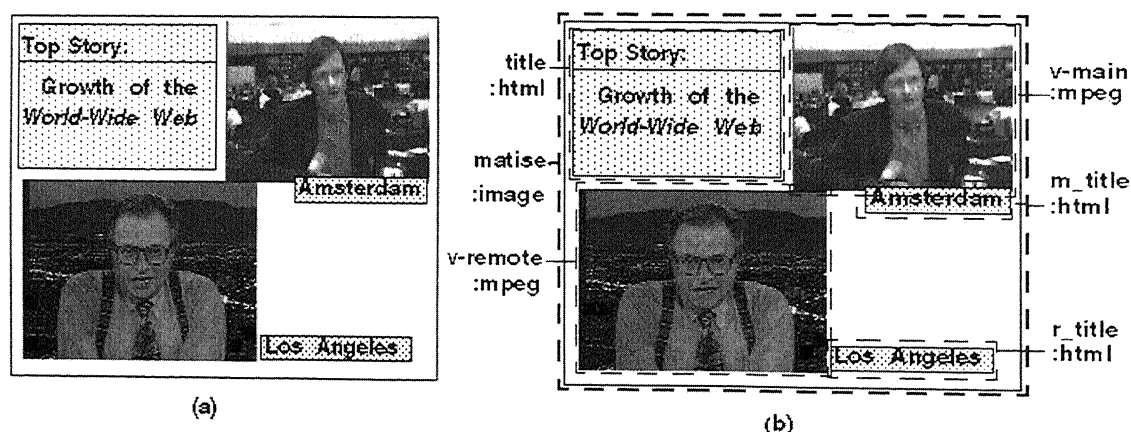


Fig. 7. The (a) presentation view and (b) associated channel map.

an active choice on the formats displayed. While the generic use of Channels as an adjunct to the switch for selection is still under debate, we have found it to be useful for delaying many semantic decisions to play-out time.

The GRiNS playout engine is implemented as a stand-alone application, and is not integrated directly into a full-function Web browser. This makes it suitable for separate distribution (such as on a CD-ROM), or for occasions when only limited ability to alter the presentation state of a document is required. (For example, if you are making a document to display a certain type of information, the stand-alone player can restrict the user's ability to enter random URLs and to browse other documents during presentation.) The use of a separate playout engine also allows us to experiment with new schedulers, protocols or data types much more easily than in a full-feature browser.

4. Current status and availability

The GRiNS authoring environment and playout engine have been implemented on all major presentation platforms (Windows-95/NT, Macintosh and UNIX). They are intended to provide a reference architecture for playing out SMIL documents and to integrate structured authoring concepts in SMIL document creation. While a full function description of all of the features of the environment are beyond the scope of this paper, we do have documentation available for interested parties. General information is available at GRiNS Web page (<http://www.cwi.nl/Chameleon/GRiNS>).

The present version of the GRiNS environment

is implemented as a combination of Python-based machine independent and dependent code, and a collection of media output drivers. The quality and availability of drivers for any particular media type vary slightly across platform, but we are making a concerted effort to provide a high degree of cross-platform compatibility. (This is an on-going process.)

We feel the SMIL can play a significant role in providing a common document representation that can be interpreted by many players. Each player can make its own contribution in terms of quality of service or implementation efficiency. In this respect, we feel that our main focus will be on the development of GRiNS authoring tools that will work with anyone's playout engines. Until these engines become prevalent, we will also provide our own environment to assist in SMIL experimentation.

The GRiNS environment was developed in part within the ESPRIT-IV CHAMELEON project. Plans exist for the free or low-cost distribution of the playout environment and the authoring interface. Interested parties should refer to the GRiNS Web site for more details.

Acknowledgments

The work of the W3C's SYMM working group was coordinated by Philipp Hoschka of W3C/INRIA. A list of contributors is available at [12]. CWI's activity in the GRiNS project is funded in part through the ESPRIT-IV project CHAMELEON [1] of the European Union. Additional sources of funding have been the ACTS SEMPER [10] and the Telematics STEM [11] projects.

Appendix A. SMIL code for the Web news example

(Note: Line numbers have been added for clarity and reference from body text.)

```

1 <smil sync=nc="soft">
2   <head>
3     <layout type="text/smil-basic">
4       <channel id="matise"/>
5       <channel id="m_title" left="4%" top="4%" width="47%" height="22%"/>
6       <channel id="v-main" left="52%" top="5%" width="45%" height="42%"/>
7       <channel id="a_uk_main"/>
8       <channel id="music"/>

```


```

9     <channel id="pix" left="5%" top="28%" width="46%" height="34%" />
10    <channel id="v-remote" left="3%" top="44%" width="46%" height="40%" />
11    <channel id="r_uk" />
12    <channel id="r_title" left="52%" top="58%" width="42%" height="22%" />
13  </layout>
14  </head>
15  <body>
16    <seq id="WebGrowth">
17      <par id="opening_segment">
18        <text id="leader_title" channel="m_title" src="the.news/html/title.html" />
19        <switch id="news_leader">
20          <video channel="v-main" src="mpeg/logo1.mpv" />
21          
22        </switch>
23        <audio id="leader_music" channel="music" src="audio/logo1.aiff" />
24      </par>
25      <seq id="story_1_web_growth">
26        <par id="node_22">
27          <text id="dcab_intro" channel="m_title" src="html/dcab_intro.html"
28            dur="8.0s" />
29          <a href="archives-dcab.smi#1" show="new">
30            <video channel="v-main" src="mpeg/dcab1.mpv" />
31          </a>
32          <seq id="sequences">
33            <par id="background_info">
34              <audio id="a_1a" channel="a_uk_main" src="audio/uk/dcab1.aiff"
35                begin="0.9s" />
36              
38            </par>
39            <par id="live_link-up">
40              <a href="archives-larry.smi#1" show="pause">
41                <video id=r_larry" channel="v-remote" src="mpeg/larry.mpv" />
42              </a>
43              <text id="remote_title" channel="r_title" src="html/r_title.html"
44                dur="6.0s" begin="id(r_larry) (begin)+1.8s" />
45              <audio id="remote_voiceover" channel="r_uk" src="audio/uk/larry1.aiff"
46                begin="id(r_larry) (begin)+1.7s" />
47            </par>
48          </seq>
49        </par>
50      </seq>
51    </seq>
52  </body>
53 </smil>

```

References

- [1] CHAMELEON: An authoring environment for adaptive multimedia presentations, ESPRIT-IV Project 20597, see <http://www.cwi.nl/Chameleon/>; see also <http://www.cwi.nl/~dcab/>
- [2] D.C.A. Bulterman, R. van Liere and G. van Rossum. A structure for transportable, dynamic multimedia documents, in: *Proceedings of 1991 Usenix Spring Conference on Multimedia Systems*, Nashville, TN, 1991, pp. 137–155; see also <http://www.cwi.nl/~dcab/>
- [3] L. Hardman, D.C.A. Bulterman and G. van Rossum. The Amsterdam hypermedia model: adding time and context to the Dexter model, *Comm. ACM* 37(2): 50–62, Feb. 1994; see also <http://www.cwi.nl/~dcab/>
- [4] D.C.A. Bulterman, Synchronization of multi-sourced multimedia data for heterogeneous target systems, in: P. Venkat Rangan (Ed.), *Network and Operating Systems Support for Digital Audio and Video*, LNCS-712, Springer-Verlag, 1993, pp. 119–129; see also <http://www.cwi.nl/~dcab/>
- [5] W3C, HTML 4.0, W3C's next version of HTML, <http://www.w3.org/MarkUp/Cougar/>
- [6] Macromedia, Inc., Director 6.0, <http://www.macromedia.com/software/director/>
- [7] C. Roisin, M. Jourdan, N. Layaïda and L. Sabry-Ismail. Authoring environment for interactive multimedia documents, in: *Proc. MMM'97*, Singapore; see also <http://opera.inrialpes.fr/OPERA/multimedia-eng.html>
- [8] P. Hoschka, Towards synchronized multimedia on the Web, *World Wide Web Journal*, Spring 1997; see <http://www.w3journal.com/6/s2.hoschka.html>
- [9] RTSL: Real-Time Streaming Language, see <http://www.real.com/server/intranet/index.html>
- [10] SEMPER: Secure Electronic Marketplace for Europe, EU ACTS Project 0042, 1995–1998, see <http://www.sempier.org/>
- [11] STEM: Sustainable Telematics for the Environment, EU Telematics Project EN-1014, 1995–1996.
- [12] W3C Synchronized Multimedia Working Group, <http://www.w3.org/AudioVideo/Group/>
- [13] W3C SMIL Draft Specification, see <http://www.w3.org/TR/WD-smil>



Dick Bulterman is founder and head of the Multimedia and Human-Computer Interaction group at CWI in Amsterdam, an academic research laboratory funded by the government of The Netherlands. Prior to joining CWI, he was assistant professor of engineering at Brown University in Providence, RI. He holds Masters and Ph.D. degrees in computer science (1977, 1981) from Brown and a degree in political economics from

Hope College (1973). In addition to CWI and Brown, he has held visiting positions at Delft, Utrecht and Leiden Universities in Holland, and has consulted frequently on operating systems and multimedia architectures. He is technical manager of the ES-

PRIT-IV project CHAMELEON and has been involved in several European Community projects.




Lynda Hardman is a researcher at CWI, where she has worked since 1990. Prior to joining CWI, she held research and development positions at a variety of universities and companies in the United Kingdom, most recently at Office Workstations, Ltd. (OWL), where she became interested in Hypertext architectures. She holds a Ph.D. from the University of Amsterdam (1998) and undergraduate degrees from Herriot-Watt University

in Scotland. She has been involved in numerous national and European research projects.



Jack Jansen joined CWI in 1985 as a member of the Amoeba distributed systems group. He studied mathematics and computer science at the Free University in Amsterdam. He is currently one of the principal implementers of the GRINS architecture and is also responsible for development of GRINS support for the Macintosh. He has gotten a haircut and contact lenses since this picture was taken.



Sjoerd Mullender is a research programmer at CWI, where he has worked on various projects since 1989. He received a masters degree in mathematics and computer science from the Free University in Amsterdam in 1988. Since joining CWI, he has worked on projects ranging from operating systems to user interface design. He is a principal implementer of the GRINS system.



Lloyd Rutledge is a Post-Doc at CWI specializing in the study of hypermedia architectures. He holds a Ph.D. in computer science from the University of Massachusetts in Lowell (1996) and an undergraduate degree in computer science from the University of Massachusetts in Amherst. He has published widely on the topic of structured multimedia systems, with a special interest in

HyTime and MHEG document architectures. At CWI, he has worked on the CHAMELEON project, with a particular interest in the support of adaptive projections of applications in multiple document formats.