



# Integrating the Amsterdam hypermedia model with the standard reference model for intelligent multimedia presentation systems

Lynda Hardman<sup>a</sup>, Marcel Worring<sup>b,\*</sup>, Dick C.A. Bulterman<sup>a</sup>

<sup>a</sup> CWI, Amsterdam, Netherlands

<sup>b</sup> WINS, University of Amsterdam, Amsterdam, Netherlands

---

## Abstract

The standard reference model (SRM) for intelligent multimedia presentation systems describes a framework for the automatic generation of multimedia presentations. This framework, however, lacks an explicit document model of the presentation being generated. The Amsterdam hypermedia model (AHM) describes the document features of a hypermedia presentation explicitly. We take the AHM and use it as a basis for describing in detail the stages of generating a hypermedia presentation within the SRM framework, which we summarise in a table. By doing so, the responsibilities of the individual SRM layers become more apparent. © 1997 Elsevier Science B.V.

*Keywords:* Standard reference model; Intelligent multimedia presentation system; Amsterdam hypermedia model

---

## 1. Introduction

Intelligent multimedia presentation systems (IMPS) deal with the dynamic creation of multimedia presentations optimally geared towards the needs of a user. In broad terms, the created presentation should define what is presented to the user (the content), where it is presented (the spatial layout) and when it is presented (temporal layout). Given a collection of user goals and availability of resources, these three aspects leave an immense number of possible presentations open. An IMPS is a reasoning system aimed at selecting the optimal one.

The standard reference model (SRM) defined in Ref. [1] considers the decomposition of this process into well-defined layers. In the content layer, a rough

outline of the content and structure of the presentation is made; in the design layer every part of the presentation is designed; and, in the realization layer, the design is realized into a full specification of the presentation in a platform-independent way. These are complemented with a control and a presentation layer, mainly dealing with the input and output from and to the user. A number of expert modules are defined, aiding the layers in their reasoning process.

The SRM considers the three aspects of the document in parallel. In the current version of the SRM, it is not precise which part of each of the document aspects is defined in which layer. The reason for this omission is the lack of a well-defined document model. The Amsterdam hypermedia model (AHM) [2,3] was explicitly designed to capture all of the above aspects, and hence can serve in the further specification of the model. The main goal of this contribution is to integrate the AHM with the SRM

---

\* Corresponding author.

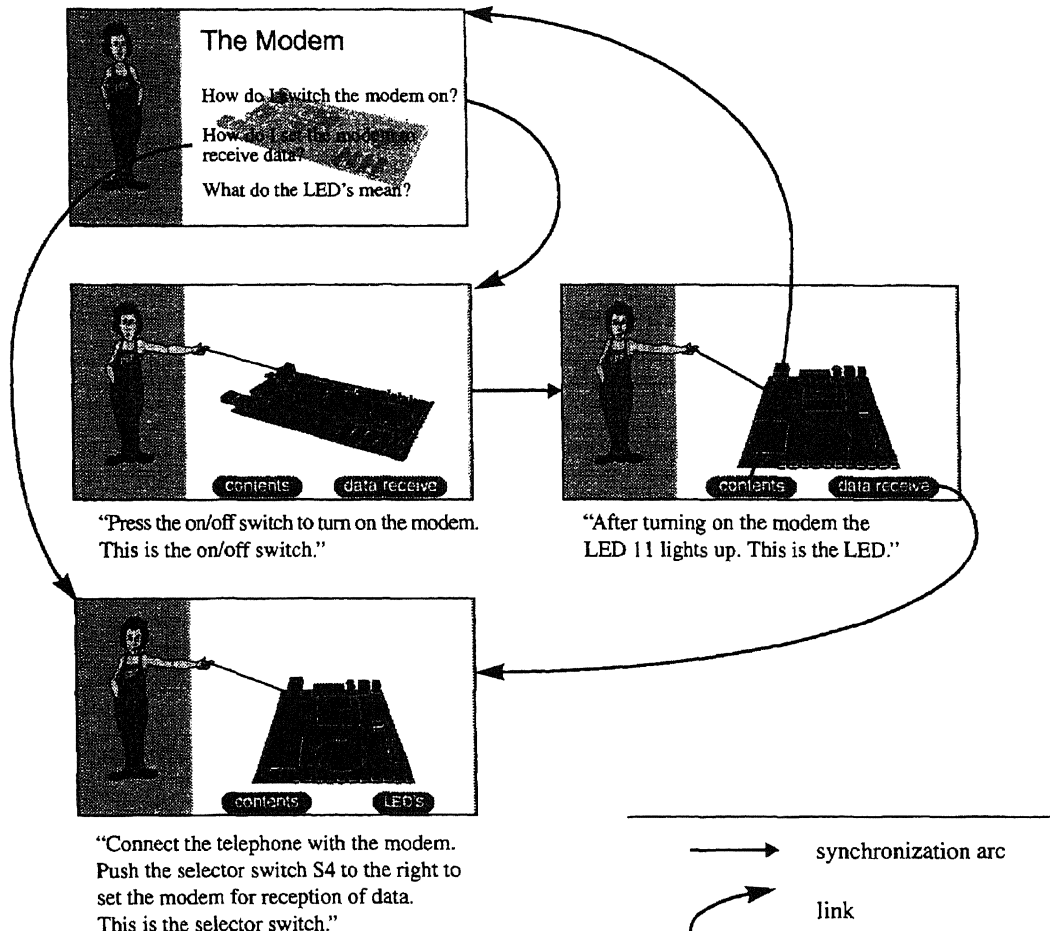


Fig. 1. An example of hypermedia presentation.

leading to a more precise model without losing expressiveness. It provides each of the layers with explicit concepts on which to base their reasoning. Consequently, it also defines for which concepts a layer should communicate with the experts defined in the SRM and when communication with the other layers is required.

The structure of this contribution is as follows. After describing an example hypermedia presentation in Section 2, we present the main features of the AHM in Section 3. In Section 4, we show how the AHM can be integrated with the SRM. In Section 5,

we give a brief discussion of our own work on generating hypermedia presentations.

## 2. Example presentation

We provide an example of a hypermedia presentation<sup>1</sup> (Fig. 1). This consists of a number of multime-

<sup>1</sup> This example is based on a generated PPP presentation explaining how a modem works. <http://www.dfki.uni-sb.de/imedia/java/applets/persona/modem.html>.

dia presentations through which a reader can navigate. The initial screen is a scene composed of four media items: two text objects, an image of the modem, and an animated image of the presenter. These start simultaneously and remain on display, until the reader selects to follow one of the links from each of the three sentences, to an explanation given by the presenter. The first explanation is composed of two subscenes, one following the other. The first subscene shows an animation of the presenter played alongside a spoken commentary. While this is playing, an image of the modem is shown, which is then changed in time with the commentary to highlight the on/off switch. The presentation continues without any interaction from the reader (indicated by the synchronization arc in the figure), and the presenter explains about LED 11. Note that the two buttons allowing the reader to jump back to the initial screen or go on to the next explanation remain on the screen throughout the scene.

### 3. Amsterdam hypermedia model

The purpose of the AHM is to capture multimedia and hypertext aspects in a unified document description. The AHM is a specification of the information needed for storing all aspects of a hypermedia document in a platform-independent way, so that it can be presented on a number of platforms while still reflecting the author's original intentions. The model associates properties, such as duration and size, with a media item and encapsulates these in an object, termed an atomic component. The atomic component, rather than the media item itself, is used as the basis for creating presentations. Other associated properties are semantic labels and presentation specifications. The AHM is described in detail in Ref. [2]. We give a summary here for completeness.

The four main components of the AHM are: (1) the atomic component that hides data dependencies of a media item from the rest of environment and associates presentation information and semantic attributes with the object; (2) the composite component that groups atomic components and smaller presentations into larger presentations; (3) the link that allows the expression of relationships among presentations; and (4) the channel that prespecifies

resources and allows styles to be applied to multiple media items.

In addition to these main components are a number of other objects in the model: (1) media items which are the data objects which form the basis of the presentation; (2) synchronization arcs which allow the expression of temporal constraints among parts of a presentation; (3) anchors that allow links to be attached to parts of media items. We describe these objects within the four main components.

#### 3.1. Atomic component

An atomic component consists of content, presentation specification, attributes and anchors. This is summarised in Table 1.

##### 3.1.1. Content

The content of an atomic component is a reference to a complete or partial media item.

##### 3.1.2. Presentation specification

The presentation specification includes temporal layout, spatial layout and style information. *Temporal information* for an atomic component is its duration. This can be derived from the content for continuous media such as video or audio, can be assigned by the author, or may be derived from the composition structure. *Spatial layout* is assigned through a channel, which has its own size and position, plus the size and position of the content with respect to the channel. The content may have its own intrinsic size, for example, for images or video. *Style information* is of three different types: media item, anchor and transition effect. Media item style is medium-dependent and includes, for example, font for text and background colour for all visual media types. Anchor style determines how the anchors are visualized to the user, e.g., pop-up when the mouse cursor is over the anchor. Transition effect specifies a special effect when the content starts or ends its display, for example, 'fade-in' or 'wipe-left'.

##### 3.1.3. Attributes

Semantic attributes can be associated with the component. These allow selections from an external knowledge representation to be associated with individual parts of a presentation.

Table 1  
AHM atomic component

Content	Media item reference Media-dependent reference	e.g., URL Part of text, image, video, etc.
Presentation specification	Duration Layout Style	Derived from content, specified by author, or determined via surrounding structure Size derived from content, or specified by author, size and position determined via channel Media item style, anchor style, transition effect
Attributes	Semantic description	Related to external knowledge representation
Anchors	Anchor ID Value Attributes Presentation specification	For referring to anchor Media-dependent reference to content Semantic description, as for component Anchor style

Table 2  
AHM composite component

Children	Component references	Atomic or composite components
Presentation specification	Synchronization arcs Duration Styles	Tuning between descendant components Determined by children and synchronization arcs Media item style, anchor style, transition effect
Attributes	Semantic description	Related to external knowledge description
Anchors	Anchor ID Value Attributes Presentation specification	For referring to anchor Component/anchor reference(s) Semantic description, as for component Anchor style

Table 3  
AHM link component

Specifiers	Anchor Direction Context Presentation specification	Component/anchor reference(s) From/To/bidirectional Component reference Anchor style
Presentation specification	Duration Transition	Transition duration Transition effect
Attributes	Semantic description	Related to external knowledge description

### 3.1.4. Anchors

Anchors are identified objects specifying part of a media item and can be used as the basis for creating links and synchronization arcs among components. An anchor consists of an anchor identifier, an anchor value, attributes and a presentation specification. The anchor value is a media-dependent description of part of the content of the component. The attributes have exactly the same function as those for the complete component and should be selectable from the same knowledge representation. The presentation specification is anchor style only.

### 3.2. Composite component

A composite component consists of children, presentation specification, attributes and anchors (Table 2). There are two types of composition: temporal and atemporal. The latter is a collection of multimedia presentations, for example, the three presentations shown in Fig. 1. The former is a grouping of media items belonging to the same presentation, for example, the complete subscene about the on/off switch. We describe the temporal composite component.

#### 3.2.1. Children

The children of a composite component are references to the atomic or composite components contained by the composite.

#### 3.2.2. Presentation specification

*Temporal information* for a composite component is the duration of the children along with their temporal relationships, stored as synchronization arcs. *Spatial layout* is assigned via the channels corresponding to the descendant atomic components. *Style* information consists of media item style, anchor style and transition effect.

#### 3.2.3. Attributes

Semantic attributes can be associated with the component.

#### 3.2.4. Anchors

Anchors are identified objects specifying part of a component. The anchor value in a composite is a reference to an anchor in a descendant component.

### 3.3. Link component

A link allows the expression of a relationship within a collection of multimedia presentations. This is most often expressed to the end-user as a navigation action. A link consists of a list of specifiers, a presentation specification, and attributes (Table 3).

#### 3.3.1. Specifiers

A specifier stores the information for the (possibly multiple) ends of the link, and is itself composed of a number of parts: an anchor, a direction, a context and a presentation specification. The *anchor* is a list of component and anchor identifiers of components that are associated with the end of the link. The *direction* specifies the direction of the semantic relationship among the link-end components. This can be interpreted as a traversal direction when required. The *context*<sup>2</sup> is the collection of objects associated with the end of the link and is given by a component reference. The *presentation specification* for the specifier is anchor style only.

#### 3.3.2. Presentation specification

The presentation specification includes temporal layout and style information. *Temporal information* for a link component is the duration of the transition from the source to the destination context. *Style* information is transition effect.

#### 3.3.3. Attributes

Semantic attributes can be associated with the component. These allow selections from an external knowledge representation to be associated with the relationship expressed using the link. Given that a link may have multiple specifiers, each with its own context, a link is able to record the semantics among a number of source presentations and a number of destination presentations.

<sup>2</sup> This should not be confused with the context expert of the SRM. Link context is a statically specified part of the document structure that defines the scope of the link ends. We will use the full expression link context to avoid confusion.

Table 4  
AHM channel

Media type		Text/Image/Video/Audio etc.
Content	Media item reference	Default
	Media-dependent reference	Default
Presentation Specification	Layout	Size and position defined with respect to window
	Duration	Default
	Style	Default media item style, default anchor style, default transition effect
Attributes	Semantic description	Related to external knowledge description

### 3.4. Channel

A channel is an abstract, rather than a physical, resource. A channel serves two basic purposes: to support multimedia styles (high-level assignment of properties) and to separate out runtime resource issues. A channel (Table 4) requires a media type along with a size and position with respect to some spatial coordinate axes, for example, a window. Other properties can be assigned as defaults for the atomic components that are assigned to the channel.

#### 3.4.1. Multimedia styles

A channel can be used to record default atomic object properties. This use of the channel can be thought of as a generalised multimedia style sheet, i.e., not only specifying presentation styles, but also by stating defaults of the other properties. Every atomic component needs values for each of these properties. For presentations with only a few components, these can be assigned per component. However, for presentations using multiple components, this information needs to be repeated for each one. The use of a multimedia style thus saves repetitive assignment. Values assigned to the channel can be overridden in the atomic component. The more properties that require to be specified per atomic component, the more convenient it becomes to 'hide' them in a channel object.

#### 3.4.2. Runtime resource issues

*3.4.2.1. Playing media item.* In its most basic form, all a channel does at runtime is play a media item—it takes a chunk of data, interprets it according to a data format and displays it on the screen. It also applies the properties assigned to the media item by default from the channel, and considers overrides from the atomic component. A collection of active channels is (almost) equivalent to a document player—the only other information needed is the scheduling, since this requires information from multiple channel streams.

*3.4.2.2. Allocation of resources.* Because the channels provide a description of the resources used by a presentation, the player software can make decisions about the presentation, such as how to distribute the audio streams over the available audio device(s), or whether the presentation can be played at all on the available hardware.

## 4. Relationship with the SRM

In Section 3, we considered how a hypermedia presentation can be modelled using atomic components describing the content, composites and links defining the temporal layout, and channels for specifying the spatial layout. In this section, we integrate the AHM and SRM by specifying for each layer of the SRM which parts of the AHM document description are the basis of reasoning. The results of this section are summarized in Table 5.

### 4.1. Control layer

This layer selects a goal from a set of goals that still have to be met. It influences the message to be expressed by a presentation, but does not make any decisions directly affecting the instantiation of parts of a particular presentation.

### 4.2. Content layer

The content layer takes a single goal passed on by the control layer and refines it to a list of subgoals.

Table 5  
The SRM layers in which the parts of the AHM are generated

SRM	AHM			
	Atomic	Composite	Link	Channel
Control	–	–	–	–
Content	Semantic attributes, media type, anchor IDs, anchor attributes	Children, semantic attributes, anchor attributes	Links among composites anchors, direction, link context; semantic attributes	Channel media type
Design	Style specifications	Temporal constraints within composite, styles, anchor IDs, anchor styles	Link anchor style, transition temporal constraints, transition	Spatial constraints among channels
Realization	Media item, anchor values, size, duration	Synchronization arcs, duration, anchor values	Transition duration	Position, size
Presentation	–	–	–	–

These are, in turn, transformed into a set of communicative acts and relations among these acts. This is carried out by four separate, but communicating processes—goal refinement, content selection, media allocation and ordering. Goal refinement is a task that has to be carried out before any parts of a presentation are created, so we do not discuss it further here.

The content selection process communicates with the application expert, and decides on the semantic content of the presentation. This can be reflected in the document structure using AHM attributes. These can be associated with a component or anchor to express the role of a particular part of the presentation, e.g., recording the knowledge structure used to guide the generation of the presentation in the first instance. The attributes should be taken from the knowledge representation used by the application expert. As an illustration, the picture of the modem at the bottom of Fig. 1 may have the associated attribute 'modem showing switch S4.' The system is then able to reuse the document later for further processing; for example, the user retraces a few steps and then formulates a different goal, and the system is not obliged to regenerate everything. A further advantage is that the description of the object can be used either to find an existing media item in a database, or to generate it.

The media allocation component assigns a particular medium to a communicative act, producing a media communicative act. In terms of the AHM, a

decision has been made as to which media type will be used. This determines, not the selection of a particular channel, but the restriction of potential channels to those of the appropriate media type. Note that while the media items have not been finalised in the content layer, this has no influence on the ability to create structures in the AHM. Since the atomic component and not the media item is the basis of a presentation, compositions and links can be created independently of the final media items.

The result of the ordering process, not necessarily linear, is an ordering of the media communicative acts. Hyperlinks, composition, timing relations and spatial layout are all different ways of expressing ordering relations among media items. When a communicative act is to be expressed using a number of media items, these methods can be traded-off against one another. For example, rather than display many items together, links can be made among smaller groups of the items, or the sequential playing of items may be a suitable alternative for laying out items next to each other at the same time [4]. These different ways of 'ordering' the material have to be captured within the content layer, but may not be finalised until the design or realization layer. We thus restrict the content layer to specifying a structure of composites and links, and leave the decisions on temporal and spatial layout to other layers.

The composition structure defines a hypermedia presentation consisting of a collection of multimedia presentations with relations among them, for exam-



ple, the different explanations of the modem in Fig. 1. The collection of presentations is an atemporal composite, containing a number of children, each of which is a temporal composite. One communicative act can be instantiated as an atomic or composite component, and a relation can be instantiated as a link. Although a link requires the specification of anchors, these are dependent on the media items that have not yet been determined. However, on the basis of semantic attributes, anchors can be assigned to components. These are specified using only anchor IDs that can later be instantiated with the appropriate anchor values within the media items. The direction and link context can already be specified, since these are independent of the final media items.

In summary, in the content layer, the parts of the AHM that need to be generated are: (i) the atemporal structure including the children of each composite; the atomic components; (ii) the links among components including anchor IDs, direction and link context; (iii) labelling of each part with semantic attributes; and (iv) the media type of the atomic components and thus the channel.

#### 4.3. Design layer

The design layer allows the processes of media selection and layout design to carry on in parallel, imposing no ordering requirements for which of these should be finalised first. The design layer is split into two communicating processes: media design and layout design.

The media design component makes decisions on the 'look and feel' of the presentation. It can communicate with the design expert and make decisions on styles for atomic and composite components, anchors, and transitions. These are the media item styles, anchor styles and special effects described in Section 3.

Some style aspects do not, at first sight, appear to fit into the AHM; for example, PPP Persona points to a part of the internal workings of a modem with a stick (Fig. 1). However, this can be described in AHM terms as follows. The pointing stick is an anchor style for highlighting a particular part of the graphic. Another way of visualizing the same information would be to have the graphic part flash when the relevant part of the commentary is spoken. In-

deed, the Persona himself may be considered as 'only' visualizing a pointing convention with which we are familiar from the physical world.

The layout design component is responsible for the spatial and temporal arrangement of the media items in the presentation plus the transition durations. In this layer, the constraints used for determining the spatial and temporal layout need to be generated. These specifications should be internally consistent before being handed on to the realization layer, which may, nonetheless, be unable to satisfy them requiring the generation of a new set of constraints. The generation of the spatial layout design should occur along those lines described in Ref. [5].

In summary, the parts of the AHM which are determined in the design layer are the media item, anchor and transition styles, the temporal and spatial layout constraints for each temporal composite, plus the temporal constraints for transitions.

#### 4.4. Realization layer

In this layer, the final decision is taken on the media items to be played in the presentation, and their corresponding spatial and temporal layout. Again, the layer is split into two communicating processes, paralleling those in the design layer—media realization and layout realization.

The media realization component ensures that appropriate media items are chosen. These may already exist, or may be generated specifically for the task. In either case, in AHM terms, the atomic component content becomes firmly specified, along with the associated channel and position within the channel. The content of an atomic component is specified either as a reference to an existing media item, or as a complete description in an appropriate specification language. Once a media item has been determined, the anchor values corresponding to the anchors can also be determined. For example, in the content layer for generating the presentation in Fig. 1, it is determined that the on/off switch is to be highlighted within the picture of the modem. In the realization layer, the picture of the modem is instantiated, and the switch can be described in terms of the image. The anchor then relates the semantic description of the switch, stored as an attribute, with the media-dependent description of the visualization of the switch in the image, the anchor value.

The layout realization component calculates the final temporal and spatial layout of each temporal composite based on the constraints specified in the design layer. The duration of an atomic component may be derived from the content for continuous media, or from the composition structure where, for example, static and continuous media are played alongside each other. The duration of a composite component can be calculated on the basis of its children, plus any other temporal constraints specified in the design layer. Synchronization arcs in the AHM can be used to specify starting times of components in relation to one another. Where links define possible paths among separate multimedia presentations, there is also a duration perceived by the user when the link is followed. This is the link transition duration, which is also finalised in this layer.

The layout of the presentation is calculated by a layout realization process, such as that described in Ref. [5]. By specifying the spatial layout of the atomic components, the channels are also determined. This results in the selection of existing applicable channels. The exact layout of an object is determined by the channel through which it is displayed, plus the position and scaling of the object relative to the channel.

In summary, in the realization layer, the final media item is specified along with corresponding anchor values; size, position, duration and start time of all components, including link transitions, are calculated; and channel selection, plus the size and position of each channel is finalised.

#### 4.5. *Presentation display layer*

The SRM separates playing of a presentation from its specification in a similar manner to the AHM, and the AHM itself does not go any further with modelling the runtime behaviour of a presentation. We have carried out work on displaying a presentation [6] conforming to the AHM, where the player software interprets the information supplied in a source document, calculates an appropriate runtime schedule, based on the timing relationships finalised in the realization layer, and does the actual display of the constituent media items. A system conforming to the SRM may similarly generate a document that could

be played by a separate player, such as that described in Ref. [6].

## 5. Discussion

We have introduced the AHM as an explicit document model into the SRM, and have given our interpretation of the parts of the model that should be derived in each layer. The three layers that generate parts of the model are the content, design and realization layers. The content layer should establish the overall connections among different presentations; the design layer should specify the constraints among the components making up a presentation; and the realization layer should ensure that these constraints are resolved to a complete platform-independent description of the presentation. The four main components of the AHM are used as follows. Atomic components allow a presentation to be built up without any initial explicit reference to the media items. Composite components allow the grouping of components, along with temporal relationships, into a single presentation. Links allow the creation of relationships among presentations. Channels allow similar styles and layouts to be reused, rather than defining new positions and sizes for each individual atomic component. A similar notion of consistency is embodied through design styles in Ref. [4].

Our own design work on automatically creating hypermedia documents, described in Ref. [7], takes a slightly different approach to that of the SRM. We consider that a large supply of media items already exists, and that these can be searched through to find representations of the concepts to be included in the presentation. The author (who may also be the end-user) can select the topic of interest and have the system make a selection of relevant media items. These can be collected together in appropriate groupings and displayed using predefined templates of channels. This takes a simpler approach to the rhetorics of the presentation by having the author take the steps carried out in the control layer and goal refinement in the content layer. This requires less analysis work on the goals that have to be met. Our approach does not demand the creation of media items, but the selection of appropriate ones from an existing database. This, in turn, requires semantic

annotation of items to be able to retrieve them. We are carrying out work parsing video segments, and objects within the segments, to facilitate the annotation process. Other work is being carried out on a prototype for assembling items into a presentation for display on the World Wide Web.

### Acknowledgements

The authors would like to thank Thomas Rist for his suggestions on improvements to this paper.

### References

- [1] M. Bordegoni, G. Faronti, S. Feiner, M. Maybury, T. Rist, S. Ruggieri, P. Trahanias, M. Wilson, A standard reference model for intelligent multimedia presentation systems, *Computer Standards Interfaces* 18 (6,7) (1997) 477-496.
- [2] L. Hardman, D.C.A. Bulterman, Document model issues for hypermedia, in: W.I. Grosky, R. Jain, R. Mehrotra (Eds.), *The Handbook of Multimedia Information Management*, Prentice-Hall, Englewood Cliffs, NJ, 1997.
- [3] L. Hardman, D.C.A. Bulterman, G. van Rossum, The Amsterdam hypermedia model: Adding time and context to the Dexter model, *Commun. ACM* 37 (2) (1994) 50-62.
- [4] L. Weitzman, K. Wittenburg, Grammar-based articulation for multimedia document design, *Multimedia Systems* 4 (1996) 99-111.
- [5] W. Graf, Intelligent multimedia layout: A reference architecture, *Computer Standards Interfaces* 18 (6,7) (1997) 517-526.
- [6] G. van Rossum, J. Jansen, K.S. Mullender, D.C.A. Bulterman, CMIFed: A presentation environment for portable hypermedia documents, in: *Proceedings of the ACM Multimedia '93*, Anaheim, CA, August 1993, pp. 183-188.
- [7] M. Worring, C. van den Berg, L. Hardman, System design for structured hypermedia generation, *Visual Information Systems '96*, Melbourne, February, 1996, pp. 254-261.



Lynda Hardman graduated in Mathematics and Physics from Glasgow University in 1982. She was the Development Manager for Guide—the first hypertext authoring system for personal computers. She joined the Multimedia and Human Computer Interaction group at CWI in 1992. Her interests include hypermedia reference models, description languages for hypermedia documents (e.g., HyTime, MHEG), hypermedia authoring systems, and the application of multimedia information retrieval to automated authoring of multimedia presentations.



Marcel Worring received his MS degree in Computer Science in 1988 from the Free University Amsterdam (VUA). In 1993, he received his PhD degree from the University of Amsterdam (UvA) on a thesis in computer vision. He currently is an assistant professor in the same department. His focus is on multimedia documents, in particular, the use of computer vision for analyzing multimedia data like video streams and scanned paper documents and the subsequent use of this information in content-based presentation. He is also an active member of the recently established Intelligent Systems Lab Amsterdam (ISLA) in which researchers from computer vision, database research, knowledge engineering, language understanding and AI join forces on various research topics in information science.



Dick Bulterman is a senior researcher at CWI in Amsterdam, where he heads the Multimedia and Human/Computer Interaction project. From 1988-1994, he led CWI's Department of Computer Systems and Telematics. Prior to joining CWI in 1988, he was on the faculty of the Division of Engineering at Brown, where he was part of the Laboratory for Engineering Man/Machine Systems. Other academic appointments include visiting professorships in Computer Science at Brown (1993-1994) and in the information theory group at TU Delft (1985) and a part-time appointment in Computer Science at the University of Utrecht (1989-1991).