3 The Amsterdam Hypermedia Model

We define the Amsterdam Hypermedia Model (AHM) based on the requirements in Chapter 2. The model is first defined per element of the model followed by a discussion of the interdependencies of the presentation specifications—temporal layout, spatial layout, style information and activation state. The model is shown to satisfy the requirements stated in Chapter 2. The document models implicit in a selection of hypertext, multimedia and hypermedia systems are described in terms of the AHM, thus demonstrating that the model is sufficiently rich for expressing a range of presentations.
This chapter is based on work presented in [HaBR93a], [HaBR93b], [HaBR94], [HaBu97].

3.1 Introduction

This chapter defines a hypermedia document model, the Amsterdam hypermedia model (AHM), demonstrates that the model satisfies the requirements given in Chapter 2, and shows that the requirements are sufficient for describing a range of hypermedia presentations, see also Table 3.1.

The model is first defined per element after which the presentation specifications common throughout the model are discussed. This provides insight into how the different elements of the model are interdependent. The presentation aspects discussed are temporal layout, spatial layout, style information and activation state. These aspects are of fundamental importance in a hypermedia document description where presentation aspects have semantic as well as aesthetic consequences.

In order to demonstrate that the AHM is able to express the presentations created by a wide range of authoring systems, we select a number of representative hypertext, multimedia and hypermedia systems and describe their implicit document models in terms of the AHM. In order to demonstrate that the model is sufficiently simple to be implementable we supply a description of the parts of the model which have been implemented in the CMIFed environment.

The structure of this chapter is as follows. Section 3.2 gives the definition of the AHM. Section 3.3 discusses the runtime interdependencies of the model elements. Section 3.4 gives a summary of the model, shows that the model meets the requirements from Chapter 2, and discusses the model as a whole. Section 3.5 describes a number of hypertext, multimedia and hypermedia systems and describes their implicit document models in terms of the AHM. Appendix 1 contains a specification of the parts of the model which have been implemented in CMIFed.

3.2 The Amsterdam hypermedia model

This section defines the Amsterdam hypermedia model. It is ordered in terms of the elements of the model: media item, channel, atomic component, composite components, link component.

We begin with an overview of the requirements from Chapter 2 and their correspondence to the elements in the model, shown in Table 3.1. Each document element description states the definition of the element, states the requirements for that element in a document conforming to the model, presents the requirements in tabular form, and discusses various aspects of that element. The requirements for a document specified in terms of the model are collected together in Table 3.6.

3.2.1 Media item

A *media item* is an amount of data that can be retrieved as a single object from a store of data objects or is generated as the output from an external process. The form of a media item is outside the scope of the AHM. For a media item to be included in a document the requirements laid upon it are that the temporal and spatial dimensionality are known and that the corresponding duration and extent are also calculable. The duration may be specified as indefinite.

3.2.2 Channel

A *channel* defines a spatial position and extent and collects together a number of presentation and semantic attributes that are applicable to a particular media type. A channel consists of an identifier, a presentation specification, attributes and a media type, Fig. 3.1.

The *identifier* is a globally unique identifier

The *presentation specification* stores a channel reference[1], spatial information for visual media types, and style information.

- The channel reference is a reference to another channel or a system-defined window.
- The spatial information specifies the position and extent of the channel. The position and extent are specified with respect to another channel, given by a channel reference, or a window.
- Style information includes media item style, anchor style and transition special effect.

Requirements \ AHM elements		Channel (section 3.2.2): Pres. spec.	Channel: Attribtues	Channel: Media type	Atomic component (section 3.2.3): Pres. spec.	Atomic: Attribtues	Atomic: Anchors: Pres. spec.	Atomic: Anchors: Attributes	Atomic: Anchors: Value	Atomic: Content	Composite components (section 3.2.4): Pres. spec.	Composite: Atributes	Composite: Anchors: Pres. spec	Composite: Anchors: Attributes	Composite: Anchors: Value	Composite: Children	Link component (section 3.2.5): Pres. spec.	Link: Attributes	Link: Anchors: Pres. spec	Link: Anchors: Attributes	Link: Anchors: Value	Link: Specifi ers
Within component layer	Media item			*						*												
Anchoring	Part of media item								*	*												
Storage layer	Properties of instance				*	*	*	*														
	Composition:																					
	temporal										*					*						*
	atemporal										*				*	*						
	spatial	*																				
	Linking																					*
	Semantic attr.		*			*		*				*		*				*		*		
Presentation Specifi cations	Temporal layout				*		*				*						*					
	Spatial layout	*			*												*					
	Style:																					
	media item	*			*						*											*
	anchor	*			*		*				*		*									*
	transition	*			*						*						*					*
	link																*					
	Activation state				*						*											*

TABLE 3.1. Hypermedia document model requirements and the elements of AHM

1. Throughout this chapter we will adhere to the standard terminology as it is evolving in the literature on hypermedia models. There are several exceptions, however, which we note here.
We use the term *component reference* to replace the Dexter term component specifi cation. This is to prevent overloading of the word specifi cation. We introduce *media item reference*, *anchor reference* and *channel reference* in addition to the component reference. These have the equivalent meaning of Dexter indirect addressing, allowing a search which returns a globally unique identity (UID) for the object. The corresponding resolver functions are implicit in the model.
Relationship terminology, such as child, parent, ancestor and descendant, is in terms of the document structure, and is not an object-oriented class structure.

The channel presentation specification may include default properties for the spatial and style information applicable to the media type associated with the channel, such as a scale factor, Z-order, orientation, background colour etc.

The *attributes* allow semantic information to be associated with the channel, for example the natural language used (for text or audio channels). A description of the attributes themselves falls outside the scope of the AHM.

The *media type* is the specification of one or more data formats that can be played on the channel.

Document requirements

The identifier and media type are required for all channels. The channel reference, position and extent are required for channels with a visual media type. The style and attributes are optional. The channel reference, position and extent are meaningless for non-visual media types. The document requirements are summarised in Table 3.2.

Model elements			**Required**	**Optional**
Channel	Presentation Specification	Channel ref.	*[a]	
		Position & extent	*[a]	
		Style		*
	Attributes			*
	Media type		*	

TABLE 3.2. Channel and required/optional document specifications

a. For a visual media type.

Discussion

The AHM channel is based on the CMIF channel, where the components of the channel are more explicitly specified. A spatial hierarchy is defined top down from a system defined window The channel referencing structure is a strict hierarchy.

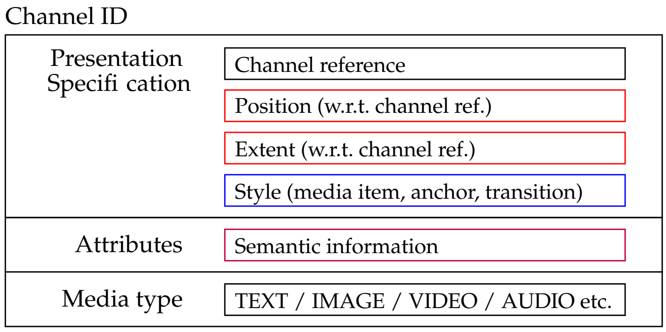

Figure 3.1. AHM channel

Associating semantic attributes with a channel allows complete streams of information to be selected, independently of the composition structure of the document.

We postpone the discussion of the style aspects of channels until Section 3.3.2.

3.2.3 Atomic Component

An *atomic component* specifies information pertinent to a media item, including the data needed for displaying it, Fig. 3.2. It consists of an identifier, a presentation specification, attributes, anchors and content.

The *identifier* is a globally unique identifier

Semantic *attributes* can be associated with the component. They enable the creation of knowledge structures and information retrieval of the components. A description of the attributes themselves falls outside the scope of the AHM.

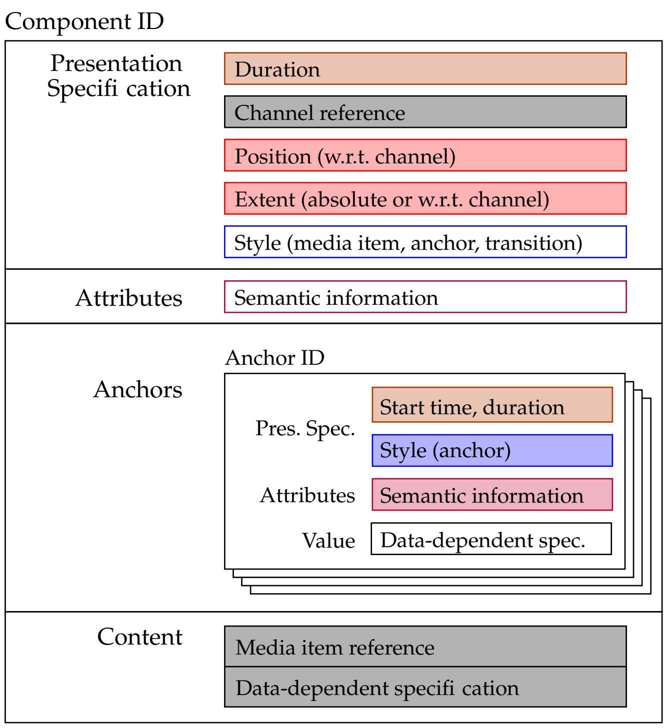

Changes from Dexter are shaded.

Figure 3.2. AHM atomic component

Presentation specification

The *presentation specification* consists of temporal, spatial and style information.

- Temporal information for an atomic component is the duration of the display of the content. The duration can be derived from the content and associated data format of a continuous media type, or can be explicitly specified for any media type. If the derived and specified durations are not the same (for a continuous media type) then the duration that is propagated throughout the model is the specified duration.
- Spatial layout is the position and extent of the display of the content, its orientation, its aspect ratio and its Z-order. The extent of the component can be derived from the data format of the content, can be specified relative to the channel reference or can be assigned an absolute value. The position is specified relative to the channel reference. Other layout information can be assigned through a channel reference or can be assigned to the component itself. The position, or other layout information, may be a function of time.
- Style information applicable to an atomic component is media item style, anchor style and transition. The style information can be assigned through the channel reference or can be assigned to the component itself. The style information can be media item, anchor and transition special effect.
  Transition information can be associated with the beginning and end of the display of the associated content. This requires the specification of a special effect along with a duration. This duration is part of the transition and is independent of the duration or start time of the component.

Anchors

Anchors form a media-independent interface between links and the components. The anchor structure for an atomic component is made up of an anchor identifier, a presentation specification, semantic attributes and a value, Fig. 3.2.

- The anchor *identifier* is unique within the component and is used for referring to the anchor from other components by specifying its containing component identifier along with the anchor identifier.
- The *presentation specification* specifies temporal information and anchor style applicable to the anchor. The former allows the association of a start time and a duration to the anchor. This is applicable only to a non-continuous media type, since otherwise the anchor value determines this implicitly.
- Semantic *attributes* can be assigned to an anchor. Attributes associate concepts directly with the data representation of a real-world object, which allows the creation of knowledge structures and information retrieval. A description of the attributes themselves falls outside the scope of the AHM.
- The anchor *value* specifies a part of the content of an atomic component using a data-dependent specification. If the content is not the complete media item associated with the component then the anchor value is

restricted to be within the bounds of the data-dependent specification of the content.

Content

The *content* specifies the data for the atomic component and consists of a media item reference, which is operating system dependent, and a data-dependent specification of part of the media item, which is data type dependent.

Document requirements

The parts of an atomic component that are required are the channel reference, the duration of the presentation specification and the content. The duration may be derived from the content. The other presentation specifications, attributes and anchors are optional. The content requires a media item reference and the data dependent specification is optional. The content need not be specified if a knowledge structure only is being created, but in this case the attributes are required. An anchor specified within an atomic component requires an anchor identifier and a value. Anchor attributes and presentation specifications are optional. The anchor value need not be specified if a knowledge structure only is being created, but the attributes would be required. The document requirements are summarised in Table 3.3.

Model elements			**Required**	**Optional**
Atomic Component	Presentation Specification	Duration	*[a]	
		Channel ref.	*	
		Position		*[b]
		Extent		*[b]
		Style		*
	Attributes			*
	Anchors			*
		Anchor ID	*	
		Pres. Spec.		*
		Attributes		*
		Value	*[c]	
	Content	Media item ref.	*[d]	
		Data-dep. spec.		*

TABLE 3.3. Atomic component and required/optional document specifications

a. The duration may be indefinite or unpredictable.
b. For a visual media type.
c. Each anchor value specification is restricted to being within the content, i.e. within the data-dependent specification of the media item reference.
d. The intrinsic duration and spatial extent are also known.

Discussion

The AHM atomic component is based on the Dexter atomic component with more detailed specification of the content and of the presentation specification.

Semantic attributes serve the same role as in Dexter. Our intention is to use these for providing cataloguing of the atomic component so that it can be found later in an information retrieval process and for creating semantic knowledge structures. An approach to this is detailed in [WBHT97].

The presentation specification for a component in Dexter is not specified in any detail. We have chosen to explicitly separate out temporal and spatial layout information from other possible presentation specifications including style information. Recording the temporal and spatial layout information is required so that components can be combined together in multimedia presentations, where the technical (as opposed to the artistic) problem is to pack items with temporal and spatial extents into a three dimensional space—or four dimensional, if the presentation consists of three dimensional objects, such as those in virtual reality applications.

The duration in the presentation specification of an atomic component is one of the document requirements. For continuous media, however, the duration can be deduced from the content and its data format, making a specified duration redundant. Without a specified duration, the rest of the environment would need access to the content data and a means of interpreting it in order to calculate the duration, thus we require its specification in the model.

The AHM atomic anchor is based on the Dexter anchor, with the addition of semantic attributes and presentation specification. Neither AHM nor Dexter specifies the form of the identifier or the value. The presentation specification for an atomic anchor allows a start time and a duration to be associated with an anchor referring to content of a non-continuous media type. For example, a music score could have each note as an anchor value and have each highlighted from a particular time for a particular duration. This would be recorded as the start time and duration in the presentation specification for the anchor value referring to a note. The anchor attributes allow knowledge structures to be created and searches to be carried out at a finer level of detail than complete atomic components, [BuKW94].

Allowing anchors to have a duration leads to a number of uses. One of these is to allow different links from the same non-continuous media item whose source anchors have the same anchor value but different start times. A second is to allow the consecutive highlighting of different anchor markers[2], as described in the music example above. A third is to enable synchronization between two streams so that the start times and durations of anchors in a non-continuous medium, e.g. written text, are dependent on the timing of corresponding

2. "Link marker" in Dexter.

anchors in a continuous medium, e.g. a spoken commentary. This would require two synchronization arcs from each audio anchor to each text anchor.

An anchor value is assumed to have a data-dependent defi nition. This is useful for the case that a data type can have anchor values specifi ed within it, such as HTML. An alternative is to allow explicit spatial and temporal extents to be specifi ed, so that a part of, e.g. an image, can be given in terms of its location on screen rather than in terms of the image data. The disadvantage of this is that if the source data is edited there is no record of where the anchor value should still be.

An assumption about an anchor value that should not be made is that it is a single temporal interval or a connected area. A single atomic anchor in an image could be, for example, all the leaves of a tree. In a video, an anchor may be all occurrences of a bouncing ball throughout the video. Neither the spatial nor the temporal extents need be contiguous.

In AHM, the content is not directly included within the atomic component, in contrast with Dexter which assumes that the data associated with the component is contained within the component. Including the content directly is seen as a disadvantage in [GrTr94a], and in the CMIF model it is explicitly stated that data can be contained within the node or stored as a reference from the node. The advantage of allowing references to the data become more apparent with larger media types, such as video, where a database can be tuned to the idiosyncrasies of the data type. Referencing the data also allows media items, or parts of items, to be re-used in other components with different semantic or presentation characteristics.

In AHM, the content is referenced using a media item reference rather than an identifi er This allows media items to be dynamically selected for inclusion in a presentation, e.g. for including the most recent photo of a head of state or for allowing alternative data formats to make appropriate use of available network bandwidth.

3.2.4 Composite components

An AHM *composite component* is a single element referring to a collection of atomic, composite and/or link components. The composition types in AHM are temporal and atemporal. Temporal composition is a grouping of components which are temporally related to one another. Atemporal composition is a grouping of components with no associated temporal relations.

As with the atomic component, an AHM composite component has an identifi er, a presentation specifi cation, attributes and anchors, but instead of content a list of children, Fig. 3.3. We specify the presentation specifi cation of temporal and atemporal composites separately.

The *identifi er* is a globally unique identifi er

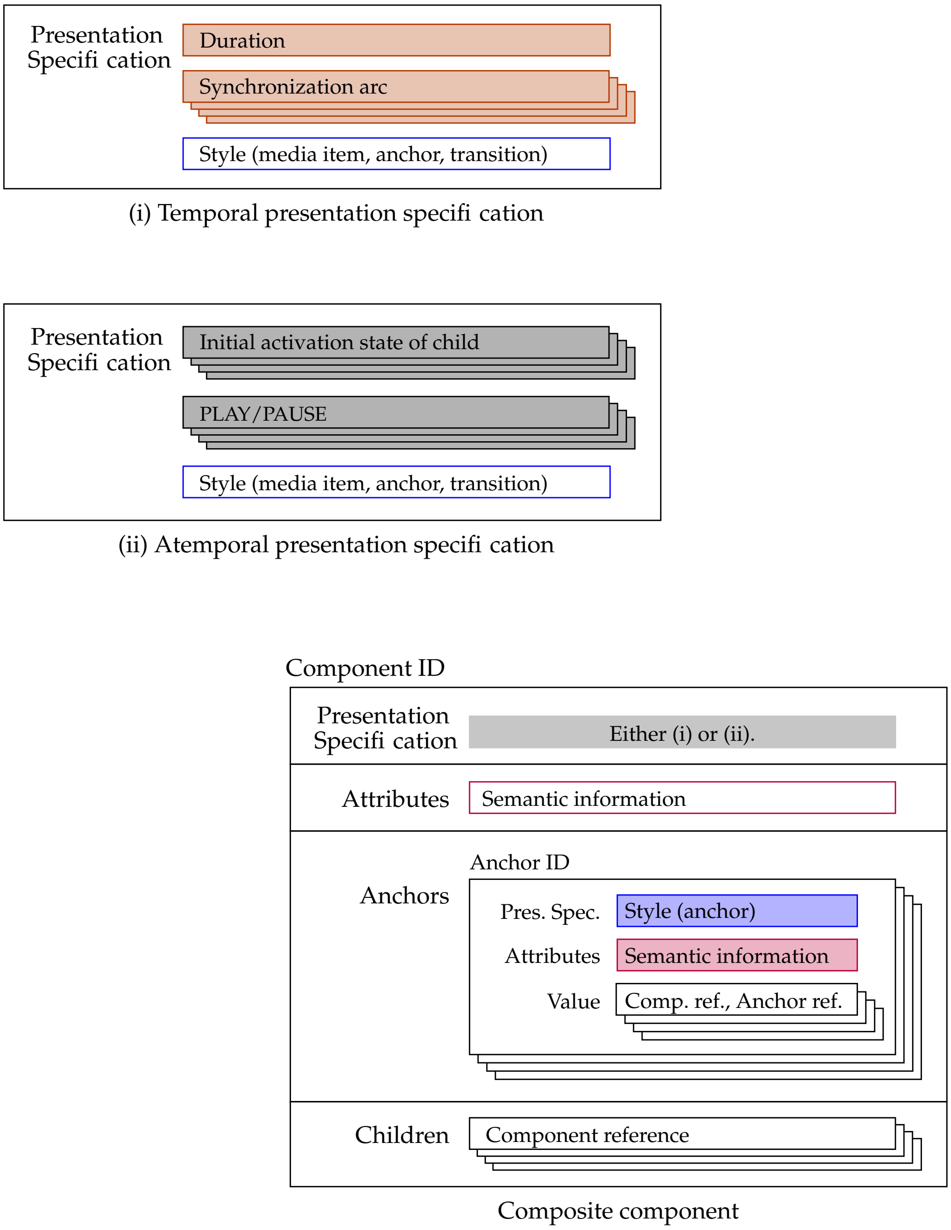

Changes from Dexter are greyed.

Figure 3.3. AHM composite components

The *attributes* allow the attachment of semantic information to the composite as a whole, for the creation of knowledge structures and for retrieval purposes. A description of the attributes themselves falls outside the scope of the AHM.

Temporal composite presentation specification

The *presentation specification* for a temporal composite consists of temporal and style information, Fig. 3.3(i). The temporal information consists of synchronization arcs and a duration.

- A *synchronization arc* specifies the timing constraint between two parts of a presentation and in doing so establishes a single time axis shared by the ends of the arc. It consists of a source and destination, a scheduling interval and a synchronization type, Fig. 3.4. The source and destination specify anchors which are the end-points of the temporal relation. Each is a component reference/anchor reference pair plus a START or END attribute. The anchor may be a point in time, but will more likely be an interval, although not necessarily contiguous. The START/END attribute specifies whether the scheduling interval starts from the beginning or the end of the source anchor reference and extends to the beginning or the end of the destination anchor reference.
  The scheduling interval specifies the temporal relation between the source and destination of the arc. The start time of the destination is relative to the source.
  The synchronization type specifies tolerance and precision properties.
  The source and destination component references are restricted to referring to a component which is a descendant of the temporal composite or the temporal composite itself.
  The children of a temporal composite with associated content require to be located along the same time-axis. This is achieved by the association of synchronization arcs. The intrinsic duration of the composite is the result of combining the durations of the children along with all the specified synchronization arcs. The duration may be indefinite.
  The synchronization arc does not have an identifier, since it is meaningful only within the temporal composite in which it is specified.
- The *duration* allows a scaling factor to be applied to the duration calculated from the duration of the children of the composite and the synchronization arcs.
- *Style* information can be specified to apply to all the descendant atomic components of the composite. It can contain media item style, anchor style and transition special effect. It may also include link style if the composite includes links.

Atemporal composite presentation specification

The *presentation specification* of an atemporal composite consists of an initial activation state and style information, Fig. 3.3(ii).

- The *initial activation state* specifies whether each child is played or not when the composite is activated at runtime. The *play/pause state* is the initial state of the child when it is made active. Each child of the composite requires an initial activation and play/pause state.
- *Style* information can be specified to apply to all the descendant atomic components of the composite. It can contain media item style, anchor style, transition special effect and link style.

Anchors

The *anchor* for a composite component has the same structure as the anchor for an atomic component: identifier, presentation specification, attributes and value.

- The *identifier* and *attributes* are the same as for an atomic component anchor.
- The *presentation specification* specifies an anchor style for the descendants of the anchor.
- The *anchor value* is a list of component reference/anchor reference pairs where the component reference refers to a component which is a descendant of the composite component. The component reference can refer to an atomic or composite component. The anchor reference may be omitted, with the interpretation that the complete component plays the role of the anchor. This removes the requirement for introducing a special anchor value for referring to a complete component. The structure of the composite anchor is a hierarchy, since the composition of components is a directed acyclic graph and because of the descendant restriction on the component reference.

Children

The *children* are the components grouped together to form the composite and are given by a list of component references.

Document requirements

A temporal composite requires at least one child with associated content (directly or indirectly). The structure composed of the children and the synchronization arcs must specify a single temporal extent, in other words the compo-

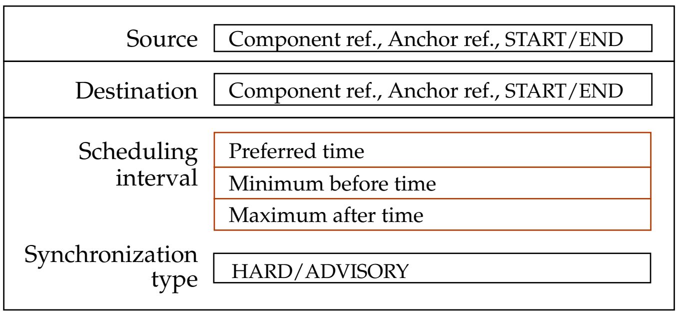

Figure 3.4. AHM synchronization arc.

nents and synchronization arcs must form a connected graph. An explicit duration, extra synchronization arcs, styles, attributes and anchors are optional.Each synchronization arc requires a source and destination component reference/anchor reference and START or END attribute, where both components are descendants of the temporal component from which they are referred to, and a preferred scheduling interval. The source and destination cannot use the same component reference. The synchronization type is optional.

An atemporal composite requires at least one child. Each child is required to have an associated initial activation state and play/pause state. Synchronization arcs cannot be specified among descendants of an atemporal composite since the descendants are temporally independent. Styles, attributes and anchors are optional. The initial activation and play/pause states may be omitted if a knowledge structure only is being created, but in this case the attributes are required.

An anchor specified within a composite component requires an identifier and at least one component reference/anchor reference pair. The anchor reference may be omitted from the component reference/anchor reference pair. Style and attributes are optional. The anchor value may be omitted if a knowledge structure only is being created, but in this case the anchor attributes are required.

The document requirements are summarised in Table 3.4.

Model elements			**Required**	**Optional**
Temporal Composite Component	Presentation Specification	Duration		*
		Sync. arcs	*	
		Style		*
	Attributes			*
	Anchors			*
		Anchor ID	*	
		Pres. Spec.		*
		Attributes		*
		List of anchors	*	
	Children	Comp. ref.	*[a]	
Atemporal Composite Component	Presentation Specification	Initial activ. state	*[b]	
		Play/pause	*[f]	
		Style		*
	Attributes		As for Temporal Composite	
	Anchors		As for Temporal Composite	
	Children		As for Temporal Composite	

TABLE 3.4. Composite components and required/optional document specifications

a. At least one is required.
b. One per child is required.

Discussion
The AHM composite component is based on the Dexter composite component. The AHM, however, references its child components rather than including them, which is the approach taken in Dexter. The AHM gives the advantage of being able to include the same (fully-fledged[3]) component in multiple composite components. The AHM differs slightly from the work in [GrTr94a], where the authors allow both referenced and included components as children of the composite, since this allows any component referenced by a composite also to be referenced by any other composite.

In contrast to Dexter[4], we exclude content from being associated with the composite component, since this introduces ambiguity about whether the other properties of the component (in particular the attributes and presentation specification) apply to the content or to all the descendants of the component. We thus require an extra atomic component to be created around the composite's "content" and have it included in the composition with no special status.

The structure composed of the children and the synchronization arcs must specify a single temporal extent. This requires the specification of a synchronization arc for all but one of the children of the composite.

More than one synchronization arc per child may not be meaningful. Two cases can be distinguished:

(a) the destination component has a finite duration, so that the start and end time of the destination is fully specified with respect to the source;

(b) the destination component has an indefinite duration, so that as well as specifying the start time of the duration with respect to the source, the end time of the destination can also be given in terms of the source. For example, a text item may be specified to start 3 seconds before the end of a video and continue until 2 seconds after the end.

A synchronization arc can be regarded as a specialised link type, since it has two atomic or composite components as end points and other pieces of scheduling information that could be collected together in the link's attributes. We, however, prefer to keep it separate in the same way that Dexter separates attributes of a component from the presentation information. In other words, a link defines a semantic relation among components, whereas a synchronization arc defines a presentation relation (specifically timing). The synchronization arc specifies a constraint between two components sharing the same time axis. A link can connect components which are otherwise unrelated in terms of space and time.

The AHM composite anchor is a new construct. It has been noted that anchoring in composites in Dexter is underspecified[LeSc94], since it does not provide

3. Dexter's composite is a composition of base components plus one set of presentation specification, attributes and anchors.
4. Although the Z specification of the Dexter model does not express this.

semantics for attaching links to embedded atomic components in composite components. Similar problems are also identified in [GrTr94a], where the authors ask specifically whether an anchor in the parent composite can be tied to an anchor in one of its components. By introducing composite anchors we are able not only to re-use anchors in a single atomic component, but also group anchors together, associate new semantics and presentation specifications to the grouping, and attach a link to the group. Note that the children of an anchor of a temporal composite may not be active at the same time, e.g. the same object appearing several times during a sequence of video clips could belong to the same composite anchor. An atemporal composite may also have composite anchors, in which case the activation state of the presentation determines which of the children of the anchor are active.

Dexter uses an identifier reference for referring to an existing anchor. Since we have extended anchors to include attributes, it becomes useful to refer to an anchor on the basis of its attributes as well as via its identifier. This requires the use of an anchor reference rather than an anchor identifier. We thus use the anchor reference in a composite anchor.

An anchor style is often medium dependent (for example underlining a word in text). For a composite anchor, styles are needed which are not necessarily medium independent, but should be functionally equivalent for different media types. For example, for the style "emphasize a destination anchor" text is underlined, an image has a box and sound contains a beep.

An atomic component has of itself no explicit start time. This is captured in the temporal composition structure via the synchronization arcs. For an atemporal composite the children are not related in any temporal way, so that the start time of any child is determined at runtime.

An atemporal composite has no duration. Each of its children may have its own duration if it is a temporal composite or an atomic component.

3.2.5 Link component

A *link component* specifies a relationship among components (atomic, composite or link), Fig. 3.5. It consists of an identifier, a presentation specification, attributes, anchors and a list of specifiers.

The *identifier* is a globally unique identifier.

Semantic *attributes* can be associated with the component, in particular for describing the relationship the link represents.

Presentation specification

The *presentation specification* consists of a duration, a relative position and style information.

- The *duration* is the duration of the transition of the source context to the destination context of the link (the definition of source and destination context is given below). It is specified using a synchronization arc, where the source of the arc is the END of the source context and the destination of the arc is the START of the destination context.

- The *relative position* is the position of the destination context with respect to the source context, where their layout is not prespecifi ed via the channels, e.g. by being in different windows.
- The *style* information consists of a link style which can be used for creating displays of links.

Anchors

The *anchor* for a link component has the same structure as the anchor for the atomic and composite components.

- The *identifi er*and *attributes* are the same as for atomic or composite component anchors.

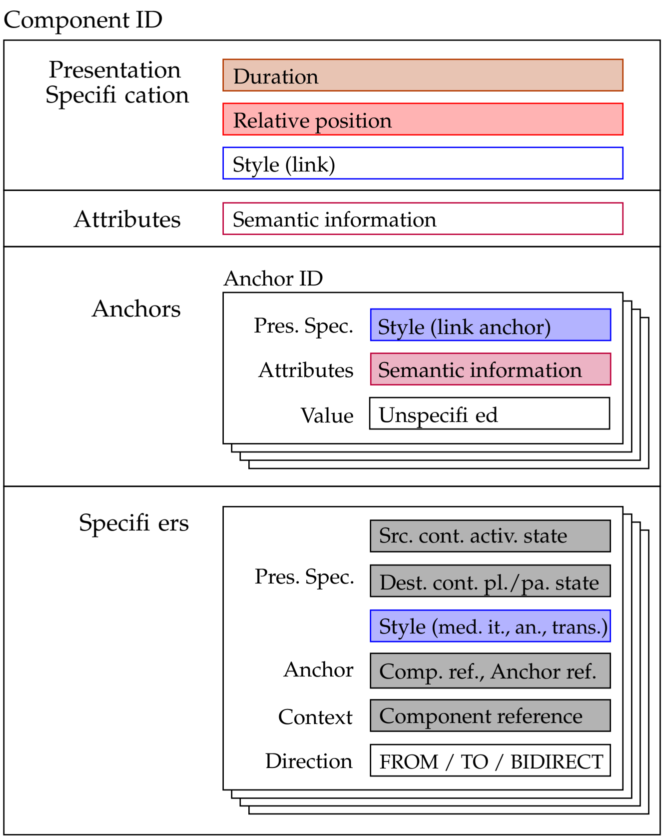

Changes from Dexter are greyed.

Figure 3.5. AHM link component.

- The *presentation specification* is a link anchor style that can be applied to a visual representation of link to link graph structures.
- The *anchor value* is outside the scope of the model.

Specifiers

A *specifier* stores the information for the (possibly multiple) ends of the link and is itself composed of a number of parts: a presentation specification, an anchor, a context and a direction.

- The *presentation specification* consists of a source context activation state, a destination context play/pause state and style information.
  The *source context activation state* determines the activation state for a specifier with direction FROM or BIDIRECT when it is part of the source context. The specifier context can remain active or DEACTIVATE when the link is followed. When it remains active the specifier context can CONTINUE to play or PAUSE.
  The *destination context play/pause state* determines the activation state for a specifier with direction TO or BIDIRECT when it is part of the destination context. The specifier context becomes active and can either PLAY or PAUSE.
  The possible styles are anchor, media item and transition special effect which can be applied when the link is followed. For example, for highlighting a source anchor to show that it has been selected, or for specifying the media item style of the destination context.
- The *anchor* specifies an end-point of the relationship represented by the link. It is given as a component reference/anchor reference pair. The component reference can refer to an atomic, composite or link component. The anchor reference may be omitted, with the interpretation that the complete component plays the role of the anchor.
- The *context* specifies the scope of the relationship at a link end and is given by a component reference. The context component reference is restricted to being an ancestor of the anchor component reference and may be equal to it. The context is thus guaranteed to contain the anchor. The context is further restricted to being an immediate child of an atemporal component, otherwise following the link may violate temporal relationships specified within a temporal composite.
- The *direction* specifies the direction of the relationship represented by the link and can be interpreted as a traversal direction. The values of direction are FROM, TO and BIDIRECT.

The transition information for the link consists of the link duration, the relative position and the special effect stored in the link specifier

Source and destination context

The *source context* of a link is the collection of specifier contexts that are active at the time the reader selects the link. The *destination context* is the collection of specifier contexts that are not part of the source context and that have direction

TO or BIDIRECT, i.e. the collection of specifier contexts that will be made active because of the selection of the link. Note that the source and destination contexts are runtime definitions. Note also that a specifier context may be in neither the source nor destination context, if it is inactive when the link is selected and has direction FROM.

Document requirements

The parts of the structure that require to be specified for a link component are at least one specifier. We do not go into the issues of dangling links, which are discussed in [GrTr94a]. The presentation specification (including transition information), attributes and anchors are optional. For each specifier the source context activation state, the context and direction are required. For each specifier with direction TO or BIDIRECT the destination context play/pause state is required. For each specifier with direction FROM or BIDIRECT the anchor is required. Specifier style is optional. The specifier context component reference is restricted to being the immediate child of an atemporal composite component when the direction is FROM or BIDIRECT and the source context activation is DEACTIVATE. The document requirements are summarised in Table 3.5.

Model elements			**Required**	**Optional**
Link Component	Presentation Specification	Duration		*
		Relative position[a]		*
		Style		*
	Attributes			*
	Anchors			*
		Anchor ID	*	
		Pres. Spec.		*
		Attributes		*
		Value	*	
	Specifiers		*[b]	
		Source cont. activ.	*[c]	
		Dest. cont. pl./pa.	*[d]	
		Style		*
		Anchor	*[c]	
		Context	*	
		Direction	*	

TABLE 3.5. Link component and required/optional document specifications

a. Valid only if destination context is in a different window from the source context.
b. At least one is required.
c. Required only for specifier directions FROM and BIDIRECT.
d. Required only for specifier directions TO and BIDIRECT.

Discussion
The AHM link component is based on the Dexter link component where the presentation specification is further specified to include style and transition information and the specifier is extended to include context information.

A link can be thought of as specifying a temporal composite with source and destination contexts as child components. The difference is that the temporal relations among the child components are brought into play only at runtime when the link is followed.

Link behaviour is discussed in [LeSc94] and [Hala88], where the authors point out that Dexter is in some cases too restrictive (for example by specifying the traversal behaviour) and that the link behaviour should be encapsulated in the storage layer rather than embedded in the particular hypermedia system. It is this type of information that we capture with the combination of context and presentation specification per specifier. Link context is semantic, delimiting the scope of the relationship represented by the link. This is interpreted at runtime as the scope of the presentation which is affected on following the link.

Link specifier context is restricted to being an ancestor of the anchor component reference, so that the anchor is guaranteed to be within (a descendant of) the context. If the anchor were not part of the context then the system would have no way of controlling the activation and deactivation of the part of the presentation that contains the anchor.

A link specifier can act as the source of a link when the specifier has direction FROM or BIDIRECT. When the source context activation state is DEACTIVATE, the context of a source link specifier is restricted to being the immediate child of an atemporal component. We impose this restriction to ensure that prespecified timing relations are not violated. For example, if the specifier context has a temporal composite as a parent then deactivating the source context would leave only part of a temporal composite playing.

Similar issues arise with link destination specifiers. A link specifier can act as the destination of a link when the specifier has direction TO or BIDIRECT. Problems arise when there are synchronization arcs from outside the destination context to within. In such cases, when you play the destination context only the synchronization relationships specified within the destination itself are applicable. Any relations outside the composite are ignored. We thus recommend that the destination context be an immediate child of an atemporal composite, but do not require it in the model.

The link's presentation specification includes the transition duration and relative position applied at runtime when traversing the link. This transition information is stored in the link's presentation specification rather than in each of the specifiers since it pertains to the relationship represented by the link. It means, however, that different durations for the potential multiple link ends cannot be specified. We have had no experience with an implementation of transitions and links with more than two specifiers, so are unable to say whether the model as it stands will be sufficient or not.

The duration of the link transition is specified using a synchronization arc since it is not the duration of a single component but is the temporal relationship between the source context and the destination context. Additional information, such as the synchronization type, is thus desirable.

The link style is applicable to a display of links in a network diagram, for example as implemented in existing hypertext systems [Hala88], [HKRC92].

In Dexter the specifier direction may have the value NONE where, as is pointed out in [GrTr94a], the meaning of NONE is unclear. Our preference for the interpretation of the direction is semantic, that is, it should specify the direction of the semantic relationship among the link-end components rather than a traversal direction. For example, component A "is an example of" component B. The traversal direction is then a presentation property associated with the link type. In the example, the link can be followed in either direction while the relationship is asymmetric. Given that both the semantic and navigational interpretations of direction NONE are not clear, we thus limit the set of directions to FROM, TO and BIDIRECT.

The *anchors* in a link allow other links to refer to it from their specifiers, allowing, e.g., knowledge structures to be created. While we consider link anchors to be outside the scope of the AHM, where linking is interpreted as a relationship that can be presented in terms of navigating among multimedia presentations, others use this construct [GrTr94a], and we see no reason for explicitly excluding it from the document model. The AHM link anchor extends the Dexter link anchor to include attributes and presentation specifications. The link anchor value is outside the scope of both the AHM and Dexter models. Attributes associated with a link anchor can be used to express the semantic relationship being represented by the link to link structure. The presentation specification can be used to specify a style for displaying a representation of the link anchors in a network diagram. We do not discuss this further.

3.3 A runtime perspective of the model

The previous section defines the Amsterdam hypermedia model in terms of the components within the model. This section goes through the presentation aspects expressible within the model and shows how information stored within different components of the model can be combined together at runtime in the final presentation. These aspects are: temporal layout, spatial layout, styles and activation state.

3.3.1 Temporal layout

Temporal information occurs explicitly and implicitly throughout the components in the model. Explicitly in the presentation specifications:

- atomic component duration,
- anchor value duration and start time of an atomic component of a non-continuous media type,
- temporal composite duration,
- synchronization arc scheduling interval,
- link duration,

and implicitly in:

- the content of an atomic component of a continuous media type,
- an anchor value of an atomic component of a continuous media type,
- the children and synchronization arcs of a temporal composite component.

The content of an atomic component, of a continuous media type, has its own intrinsic duration. This can be used as the duration of the atomic component, and stored in the duration specification of the component, or another duration can be specified. Two "contradictory" durations have no consequences for the rest of the model, since the specified duration is the one used, but a system presenting the content would need to ensure that the playing of the content lasted as long as the specified duration, for example by stopping the playing before the end of the content is reached or by playing the content faster.

An anchor of an atomic component has a start time and a duration. When the content is of a continuous media type the start time and duration are calculable from the anchor value specification. For a non-continuous media type these can be specified. The model does not require a start time or duration for the anchor, although this is required for an atomic component. The temporal information for the anchor is not referred to from elsewhere in the model, so does not have to be stated explicitly.

In the AHM temporal composition is "bottom-up" composition, that is, the duration of a temporal composite component is derived from the timing of its (content-containing) children and the related synchronization arcs, e.g. Fig. 3.6. There may also be an explicitly specified duration for the composite, in which case the specified duration should be used in the rest of the model (for example for inclusion in other temporal composites). The player software should ensure

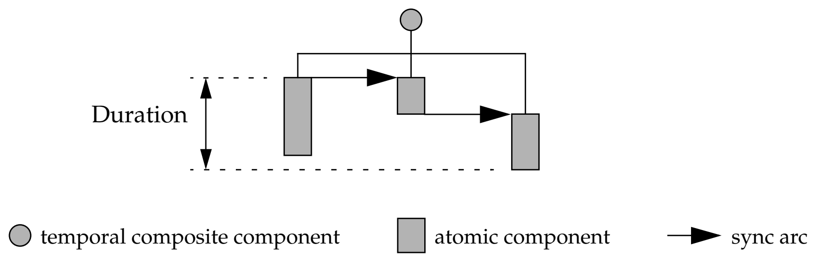

Figure 3.6. Duration of temporal composite component

that the playing of the children of the composite lasts as long as the specified duration, for example by taking appropriate scaling or cropping actions.

It is beyond the scope of the model to resolve the interpretation of inconsistent temporal relations given by the synchronization arcs among the descendants of a temporal composite.

3.3.2 Spatial layout

Spatial information occurs explicitly and implicitly throughout the components in the model. Explicitly in the presentation specifications:

- channel position and extent,
- atomic component position and extent,
- link relative position,

and implicitly in:

- the channel reference in a channel (position and extent),
- the channel reference in an atomic component (position and extent),
- an anchor value in an atomic component of a visual media type (position and extent),
- the content in an atomic component of a visual media type (extent).

The extent of the content of an atomic component can be found in three places in the document model. Firstly, if the content is of a visual media type it may have an intrinsic extent. Secondly, the extent can be specified in the atomic component presentation specification. This may have an absolute value, or may be given in terms of the extent of the content or in terms of the channel extent. The exact terms of specification fall outside the scope of the model. Thirdly, the channel associated with the atomic component has an extent.

The position of the content can be found in two places in the document model. Firstly, in the atomic component as a position with respect to the channel reference. This can be in terms of the position and extent of the channel, e.g. centre left-right and top-bottom with respect to the channel's position and extent. Secondly, in the channel as a position with respect to the channel reference's extent and position. Example position and extents are shown in Fig. 3.7.

The atomic component layout is specified in terms of the channel. The channel can provide default layout specifications, such as "align to top, center left-right, scale to fill", which can be overridden by the specifications in the atomic component, e.g. "center top-bottom". The link specifier layout can be used when the source and destination are in different windows.

Note that the position and extent specifications throughout the model may vary with time.

In contrast to the temporal specification, the spatial composition in AHM is "top-down", that is, the extent of a channel is specified in terms of its channel reference, Fig. 3.7. The top of the hierarchy is a system defined window This allows the layout to be determined independently of the extents of the atomic components, and, more importantly, allows consistent layouts to be designed with no knowledge of the content that is to be played and without imposing restrictions on the extent of the content. Another advantage is that one particular document can be played back on a variety of screen sizes without having to alter the document specifi cation in any way It is the system window which can be scaled and the rest of the layout specifi cations follow

3.3.3 Temporal and spatial layout combined

The layout structure of a presentation is independent of its temporal structure. Both, however, are needed for the complete description of a presentation. The temporal hierarchy and the layout hierarchy meet at the atomic components, illustrated in Fig. 3.8. The durations of the atomic components form the building blocks of the presentation's timing. The channels, defi ned in terms of system windows, form the layout hierarchy. A piece of content is displayed in relation to the channel associated with the atomic component, from the time determined by the temporal hierarchy for the duration specifi ed in the atomic component.

3.3.4 Styles

The styles that occur throughout the model are media item style, anchor style and transition special effect. Link style and link anchor style are also mentioned, but we do not discuss these further. Further specifi cation of the styles themselves falls outside the scope of the model, although examples were provided in Chapter 2.

Media item and transition styles can be found in four places in the presentation specifi cations:

- channel,
- atomic component,
- composite component,

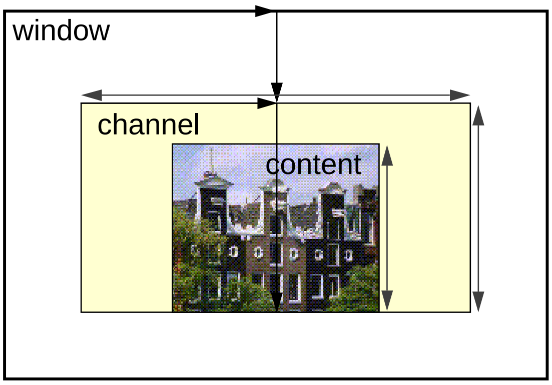

Channel in terms of window:
position centred left-right,
position top at height 25%,
height 60%,
width 70%.

Content in terms of channel:
position centred left-right,
position aligned bottom,
height 80%,
aspect ratio preserved.

Figure 3.7. Spatial layout

- link specifi er

Anchor style can be found in six places in the presentation specifi cations:

- channel,
- atomic component,
- atomic component anchor,
- composite component,
- composite component anchor,
- link specifi er

The application of styles in increasing order of override is the following:

- channel,
- atomic component then atomic component anchor (for anchor style only),
- composite component then composite component anchor (for anchor style only),
- link specifi er

The atomic component style overrides the channel style, since the channel gives only a default to save multiple specifi cation for each atomic component. The composite component style overrides the atomic style, since the same atomic may be included in a number of composites where the context requires a different style, for example background colour. Similarly, the link specifi er style over-

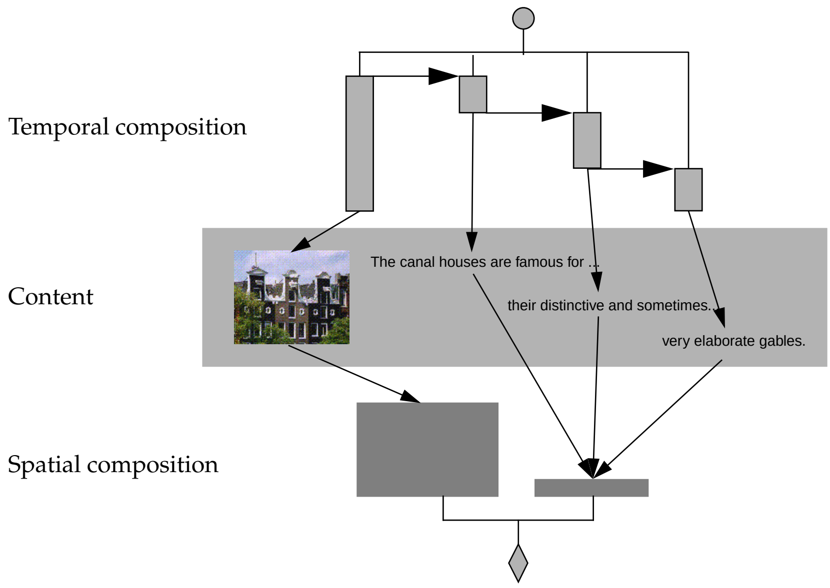

Figure 3.8. Temporal hierarchy meets spatial hierarchy

rides the composite component style since the same composite used in the destination context of a number of links may require different style.

We repeat Figure 12 from Chapter 2 to illustrate where in the model the different styles are stored, Fig. 3.9. Anchor markers are shown as boxed markers for clarity, whereas in an active system they may be indicated more dynamically, e.g. using change of mouse cursor shape.

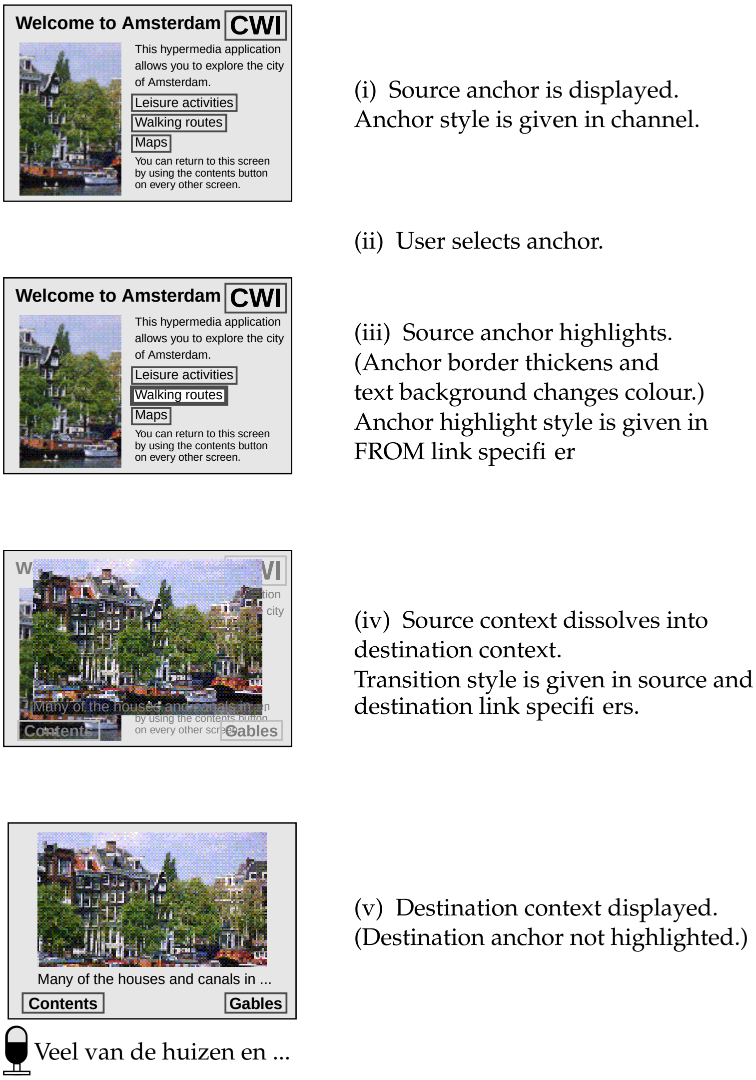

Figure 3.9. Anchor and transition styles on following a link

3.3.5 Activation state

Activation has two aspects— whether an item is active or not, and whether it is traversing its own timeline. Anchor marker selections can be made when the component is active.

Activation state is recorded in two places in the document model:

- with each child in the presentation specification of an atemporal composite component,
- as part of the source context activation state of the link specifier

Play/pause state is recorded in three places in the document model:

- with each child in the presentation specification of an atemporal composite component.
- as part of the source context activation state of the link specifier,
- the destination context play/pause state of the link specifier

Each child of an atemporal composite has an initial activation state. A presentation thus begins with the children of an atemporal component which are marked as active. Each active child starts up as playing or as paused. The activation state is stored in the atemporal composite rather than with each child, since each child may be included in other composites with different activation states, Fig. 3.3. The activation state changes according to links that are followed. The source context behaves as specified in the link specifiers, according to the source context activation state—DEACTIVATE/CONTINUE/PAUSE. The destination context defines the presentations that become active, i.e. those not in the source context and with direction TO or BIDIRECT. The destination context play/pause state specifies whether the specifier context should play or pause.

We first describe the activation state of the simple case of following a link with two specifiers, one with direction FROM and the other TO. The specifier with direction FROM is active and the reader selects the corresponding anchor marker. The source context is continued, paused or deactivated, depending on the source context activation flag. The destination context is activated and is played or paused depending on the destination context play/pause flag.

We now describe the case for a link with multiple specifiers. An example is given in Fig. 3.10. The reader selects one of the active anchor markers, in the example the lower anchor marker on the map. Since the reader was able to select it, the component referenced by the specifier must be active, so the specifier is by definition in the source context. Other specifiers whose context is active are also in the source context, for example the picture of the palace. Each specifier context in the source context with direction FROM or BIDIRECT is deactivated, continued or paused depending on the value of the source context activation flag. In the example the map is continued and the palace is deactivated. If the specifier context in the source context has direction TO then the specifier context is continued. The destination context is the set of non-active specifiers in the link with direction TO or BIDIRECT. In the example the picture of the tower is shown. Each specifier context in the destination context is made active, taking into account the value of the play/pause state flag.

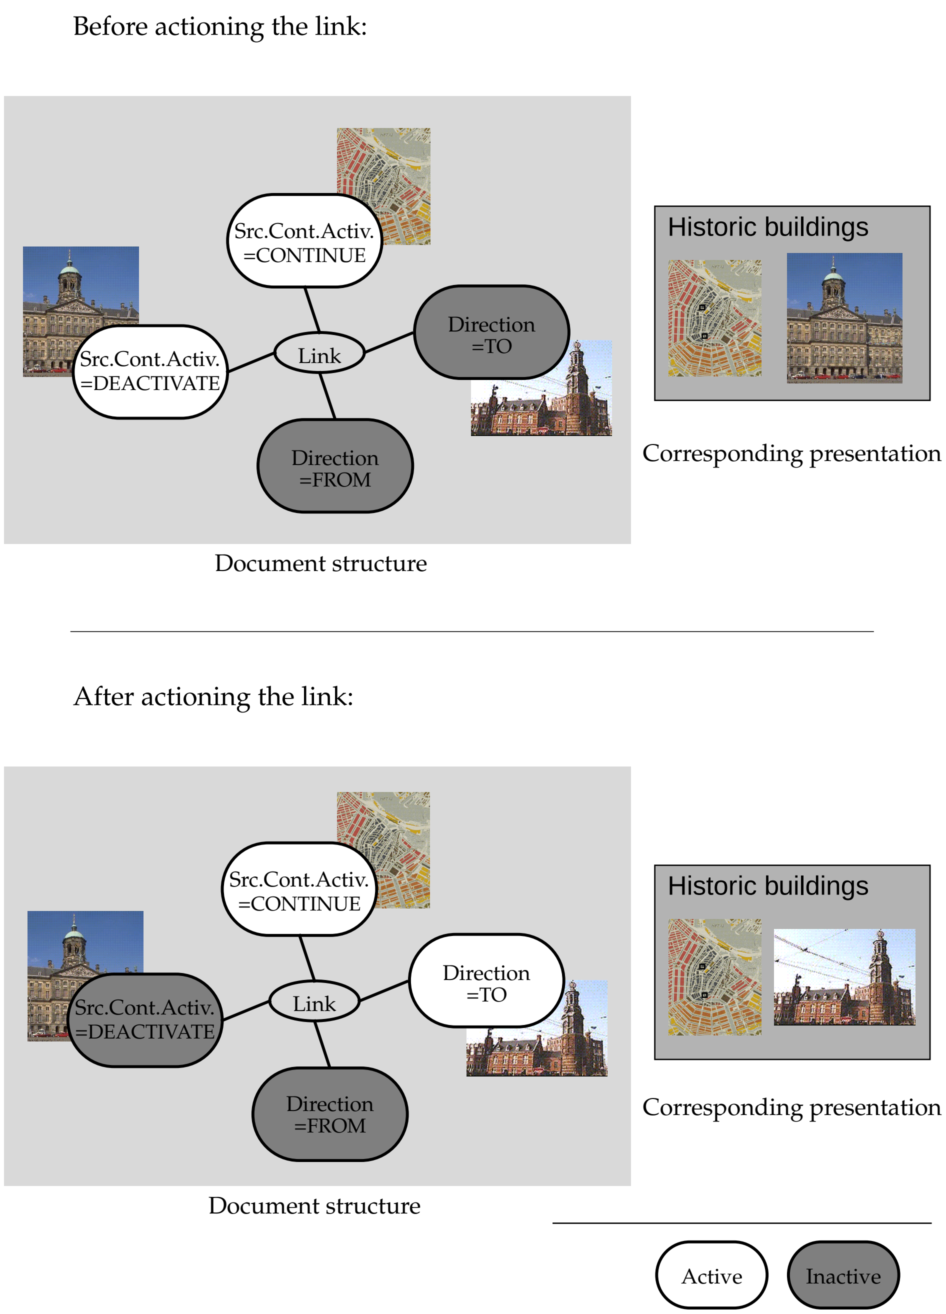

Figure 3.10. Change in activation state on following a link

3.4 Summary and discussion of the model

Having described the model components in Section 3.2 and discussed the presentation specifi cations in Section3.3 we provide summaries and a discussion in this section. A summary of the model and the requirements for a document to conform to the model are given in Section 3.4.1 and in Table 3.6. To demonstrate that the document model meets the requirements stated in Chapter 2, the correspondence between these and the AHM elements is given in Section 3.4.2 and summarised in Table 3.7. A discussion of the model is given in Section 3.4.3.

3.4.1 Summary of AHM

A summary of the AHM is given as a diagrammatic impression in Fig. 3.11. This shows the combination of time and structure in one document model and is based on Figures 13 and 15 in Chapter 2.

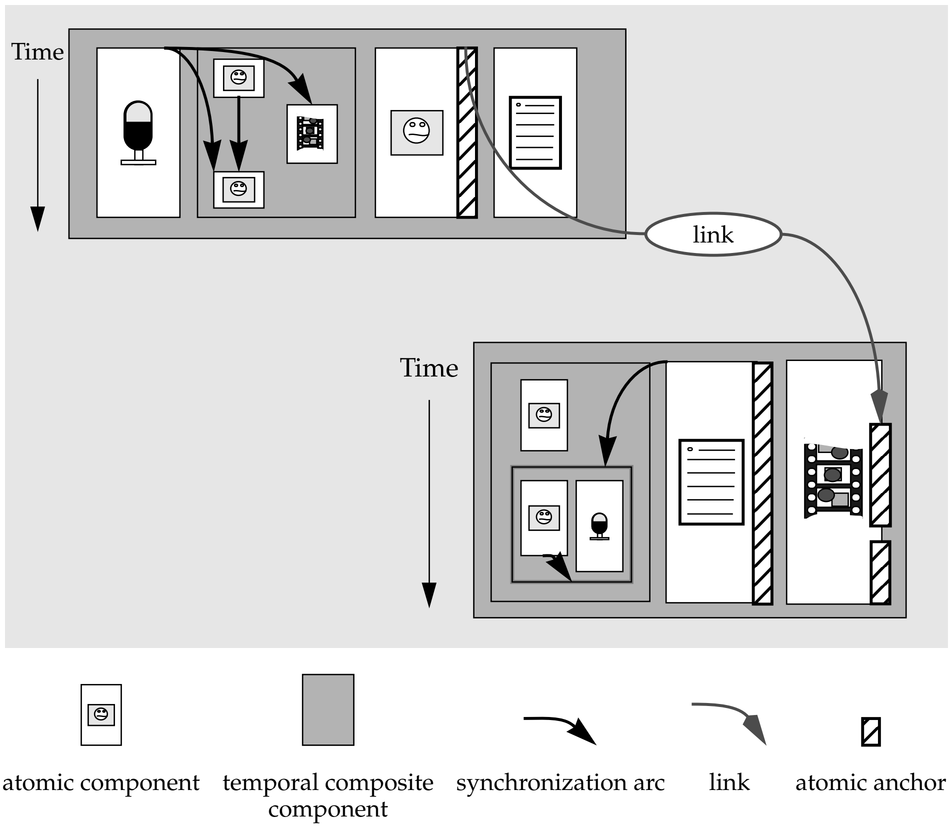

Figure 3.11. Amsterdam hypermedia model overview.

The main components of the model are:
- media item,
- channel (not shown),
- atomic component and atomic anchor,
- temporal composite component including synchronization arcs,
- atemporal composite component (not shown),
- composite anchors and
- link.

Structure elements from the Dexter model have been extended to include explicit temporal and spatial information, and anchors for composite components have been defined explicitly Synchronization arc and channel structures from the CMIF model have been included and are described explicitly.

A more detailed summary of the model is given in Table 3.6, where the requirements for a document conforming to the model are also summarised. References back to the section where each requirement was first stated are also included. Each component requires an identifier, not shown in the table.

A formal description of the model in the Object Z language is given in [OsEl97]. This is included as Appendix 2 of this thesis as a supplement to the definition provided here.

Model elements			**Required**	**Optional**	**Stated in section**
Channel	Presentation	Channel ref.	*[a]		3.2.2 Channel
	Specification	Position & extent	*[a]		
		Style		*	
	Attributes			*	3.2.2 Channel
	Media type		*		3.2.2 Channel
Atomic	Presentation	Duration	*[b]		3.2.3 Atomic Component
Component	Specification	Channel ref.	*		
		Position		*[a]	
		Extent		*[a]	
		Style		*	
	Attributes			*	3.2.3 Atomic Component
	Anchors			*	3.2.3 Atomic Component
		Anchor ID	*		
		Pres. Spec.		*	
		Attributes		*	
		Value	*[c]		
	Content	Media item ref.	*[d]		3.2.1 Media item,
		Data-dep. spec.		*	3.2.3 Atomic Component

TABLE 3.6. AHM elements and required/optional document specifications

Model elements			Required	Optional	Stated in section
Temporal Composite Component	Presentation Specification	Duration		*	Temporal composite presentation specification
		Sync. arcs	*		
		Style		*	
	Attributes			*	3.2.4 Composite components
	Anchors			*	3.2.4 Composite components
		Anchor ID	*		
		Pres. Spec.		*	
		Attributes		*	
		List of anchors	*		
	Children	Comp. ref.	*[e]		3.2.4 Composite components
Atemporal Composite Component	Presentation Specification	Initial activ. state	*[f]		Atemporal composite presentation specification
		Play/pause	*[f]		
		Style		*	
	Attributes		As for Temporal Composite		
	Anchors		As for Temporal Composite		
	Children		As for Temporal Composite		
Link Component	Presentation Specification	Duration		*	3.2.5 Link component
		Relative position[g]		*	
		Style		*	
	Attributes			*	3.2.5 Link component
	Anchors			*	3.2.5 Link component
		Anchor ID	*		
		Pres. Spec.		*	
		Attributes		*	
		Value	*		
	Specifiers		*[f]		3.2.5 Link component
		Source cont. activ.	*[h]		
		Dest. cont. pl./pa.	*[i]		
		Style		*	
		Anchor	*[j]		
		Context	*		
		Direction	*		

TABLE 3.6. AHM elements and required/optional document specifications

a. For a visual media type.
b. The duration may be indefinite or unpredictable.
c. Each anchor value specification is restricted to being within the content, i.e. within the data-dependent specification of the media item reference.
d. The intrinsic duration and spatial extent are also known.
e. At least one is required.
f. One per child is required.
g. Valid only if destination context is in a different window from the source context.
h. Required only for specifier directions FROM and BIDIRECT.
i. Required only for specifier directions TO and BIDIRECT.
j. Optional if the specifier is of direction TO.

3.4.2 Showing the AHM meets the requirements

Table 3.7 gives a summary of the requirements stated in Chapter 2 and shows which parts of the model satisfy the requirement. Most of these are straightforward, in that each requirement is satisfied using one, or part of one, component. Two, however, are satisfied through a combination of different parts of the model: Allen's relations, and the specification of a time axis.

3.4.2.1 Specification of time axis

The model includes no explicit timeline in a temporal composite component. The time axis is instead calculated on the basis of the temporal composition structure, including the duration of the descendant components and the synchronization arcs specifying constraints among the components.

3.4.2.2 Allen's temporal relations

Temporal relations between two components (atomic or temporal composite) are stored as synchronization arcs. Complete relative timing requires the support of all of Allen's temporal relations [Alle83]. We show that synchronization arcs can express these in Fig. 3.12.

We have thus demonstrated that the Amsterdam Hypermedia Model satisfies the requirements for a hypermedia document model as stated in Chapter 2.

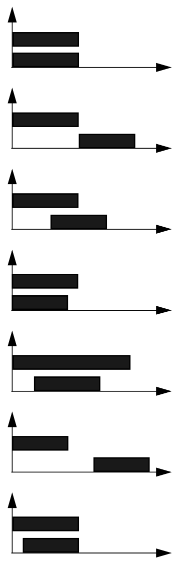

A equals B	from beginning A, to beginning B, delay=0 *and* from end A, to end B, delay=0
A meets B	from end A, to beginning B, delay=0
A overlaps B	from beginning A, to beginning B, delay=t1 *and* from end A, to end B, delay=t2
A starts with B	from beginning A, to beginning B, delay=0
A during B	from beginning A, to beginning B, delay=t1 *and* from end B, to end A, delay=t2
A before B	from end A, to beginning B, delay=t
A ends with B	from end A, to end B, delay=0

Figure 3.12. The Allen time relations expressed as synchronization arcs.

Dexter layer	Document feature	Requirement	Satisfied in
With. comp.	Media items	temporal and spatial dimensions	3.2.1 Media item
Anchoring	Ref. to part of media item	data-dependent specification of part of media item	3.2.3 Atomic Component, Anchors and Content
Storage layer	Instance of media item	reference to (part of) media item, data format, duration, start time, extent, pos., asp. rat., orien., Z-order, style (media item, anchor, transition) start points for links, semantic attributes	3.2.3 Atomic Component, Content 3.2.2 Channel 3.2.3 Atomic Component, Presentation specification 3.2.3 Atomic Component, Anchors 3.2.3 Atomic Component
	Composition	temporal composition, spatial composition, space/time independ. composition anchor composition	3.2.4 Composite components, Temporal 3.2.2 Channel 3.2.4 Composite components, Atemporal 3.2.4 Composite components, Anchors
	Linking	source and destination anchor, source and destination context, transition (duration and special effect)	3.2.5 Link component, Specifiers 3.2.5 Link component, Specifiers 3.2.5 Link component, Presentation specification
	Semantic attributes	associate with: anchors, instances, compositions of anchors, compositions of intances, and links	3.2.3 Atomic Component, Anchors 3.2.3 Atomic Component 3.2.4 Composite components, Anchors 3.2.4 Composite components 3.2.5 Link component
Presentation Specifications	Temporal layout	time axis, start time, duration, scaling, Allen's relations, link transition duration	3.4.2.1 Specification of time axis 3.2.4 Composite components, Temporal 3.2.1 Media item, 3.2.3 Atomic Component, 3.2.4 Composite components, Temporal 3.2.4 Composite components, Temporal 3.4.2.2 Allen's temporal relations 3.2.5 Link component, Presentation spec.
	Spatial layout	space axis, position, possibly changing w.r.t. time, extent, scaling, orientation, aspect ratio, Z-order, link transition spatial relation	3.2.2 Channel 3.2.2 Channel, 3.2.3 Atomic Component (not w.r.t. time) 3.2.2 Channel, 3.2.3 Atomic Component 3.2.2 Channel, 3.2.3 Atomic Component 3.2.2 Channel, 3.2.3 Atomic Component 3.2.2 Channel, 3.2.3 Atomic Component 3.2.2 Channel, 3.2.3 Atomic Component 3.2.5 Link component, Presentation spec.
	Styles	associate with anchors, instances, compositions of anchors, compositions of intances, and links	3.2.3 Atomic Component, Anchors 3.2.3 Atomic Component 3.2.4 Composite components, Anchors 3.2.4 Composite components 3.2.5 Link component
	Activation state	initial activation state change in activation state	3.2.4 Composite components, Atemporal 3.2.5 Link component, Specifiers

TABLE 3.7. Summary of requirements for a hypermedia document model

3.4.3 Discussion of the AHM

We have already discussed the pros and cons of the individual components of the model in the previous discussion subsections. This section discusses the model as a whole, with emphasis on potential improvements and extensions.

The model is based predominantly on the Dexter model, and thus inherits a number of the limitations that apply to the model, e.g., those discussed in [Hala88], [LeSc94]. Some of these are solved by others, in which case we would wish to inherit the solution. We do not intend to resolve them, since the goal of the AHM was not to improve a hypertext model, but to propose a hypermedia model.

A component reference in the AHM is not restricted to any notion of document boundary. The only requirement is that the player has access to the component. For example, if the component reference is given using a URL the player would have to support the HTTP protocol. The main consequence of this is that there is no guaranteed bound to the number of components making up the hypermedia presentation. This is not a problem for activating and deactivating parts of multimedia presentations, but would be a problem if a representation of the link graph was desired.

Channels

Channels are introduced in the model as a means of specifying spatial layout and of allowing the specifi cation of style defaults for atomic components. The introduction of the channel construct in CMIF [BuRL91] was to provide high level control of runtime resource allocation. These two roles remain mixed and somewhat implicit in the AHM. In particular, the channel could be split into a more runtime resource-oriented construct which refers to a layout construct. The layout construct could independently specify layout which could be referred to from multiple channels, particularly useful for items of different media types, e.g. image and video, which are intended to use the same screen space.

While the model does not restrict the movement of a channel with respect to its channel reference with time, this would violate the resource function of the channel.

A limitation of the use of channels for layout is that items cannot be positioned with respect to the content of other items. This is suffi cient for a wide range of presentations, but not, for example, for the overlay of maps of a city taken from different periods to show how the city expanded. Here a part of a map requires to be overlaid accurately onto a part of another map. Scaling in terms of the map dimensions as a whole is insuffi cient. Introducing spatial relations among anchors, i.e. the model construct for referring directly to content, is a potential solution. If layout with respect to content were introduced then the channels would lose their layout independence, thus requiring a more complex implementation.

Anchors
Anchors in the model are based on data-dependent descriptions of the content of atomic components. Since the model makes the dimensions of time and space explicit, any part of the presentation, whether a temporal composite or an individual atomic component, can be specified in terms of the temporal and spatial extents of the presentation. For example, an anchor of a component could be defined which starts at 20% of the duration of the component and continues until 25% of the duration. This type of anchor specification could be used in the same way as the current anchor as the end of a link or a synchronization arc. We have chosen not to include this type of anchor in the model, since the semantics of the anchor value are not clear, and the difficulties of guaranteeing a valid anchor value after editing a media item are even greater than those for data-dependent anchor value specifications.

Link component
While we noted in the link component discussion that we prefer the notion of direction for a link to be based on semantics rather than on traversal, we use link specifier direction in the model as a traversal direction. An extension to the model would be to replace the single direction with a semantic direction and a traversal direction, so that the behaviour of the link could be independent of its semantics. Up until now our own use of the model has been more presentation based than semantically based so we have not yet experienced the need for two distinct directions.

A link has a single duration and spatial relation, whereas it is probably more appropriate to assign these per specifier, so that the timing and placement of the, possibly multiple, destination contexts of the link can be individually specified. The duration would be specified with respect to the time of the user interaction and the spatial relation with respect to the anchor marker selected.

Presentation Specifications
The figures of the model components in this section give the impression that we have specified the presentation specification of all the components completely While we have divided the original Dexter presentation specification into temporal, spatial, style and activation information we do not intend to imply that this is all that a presentation specification can contain. Just as an application is able to define its own semantic attributes and styles, the presentation specification can be extended for application dependent purposes. The AHM states that certain aspects of the presentation specifications occur throughout the model, in particular the temporal and spatial information, which are interdependent. An example of an additional presentation specification is whether an anchor's marker should be displayed or not.

Temporal layout
The model includes no explicit timeline in a temporal composite component. Without an explicit timeline we cannot specify changes of tempo within the

presentation. We are able, in the current model, to specify the tempo of individual components, but *accelerando* and *ritardando* over a temporal composite cannot be specified. In order to incorporate this information, one could include a timeline in the component specific presentation specification for a temporal composite along with specifications of the tempo along the timeline.

Time is treated within the model purely as relative. The initial start time of a document is when the reader starts up a session and all other times are relative to this start time or relative to when the reader follows a link. It would be useful to allow absolute times, such as 18:00:00 MET, to be included in the model, but this would require the development of activation mechanisms other than those already provided by linking.

Spatial layout

While a composite component may refer to a number of atomic components each with their own spatial layout specifications, it might be the case that for a particular composite the author requires a different layout. An example of this is the table top composite in [Grøn94]. The model does not allow the specification of layout information in the composite's own presentation specification. If the model were to include layout information in a composite then an extra level of override would be needed—the channel specifies the default layout, the atomic component overrides it, and the encompassing composite component would take precedence.

Activation state

The current activation states included in the model are whether the presentation is active or not and whether an active presentation is playing or paused. An additional state is possible which distinguishes whether the components are visible or invisible. The assumption is made within the model that if a component is active that it is also visible. This does not, however, allow synchronized streams to be turned on and off while they are playing without losing their current play state. Useful, for example, for playing a video with synchronized commentaries in different languages where the reader is able to select which language to see and/or hear. If a component were to be turned off, in the current model, it is by definition deactivated. Turning it on again would force it to be played from the beginning again. With the addition of a visible/invisible state the component could be turned off, but still continue to play. This would require a mechanism for stating the initial visible state and how this changes via user interaction, akin to that for the play/pause state already defined in the model.

3.5 Implicit document models of existing systems expressed in AHM

Chapter 2 motivated the requirements for a hypermedia document model using a simple presentation. This example provided a lower bound on the features necessary for a document model. This left open the question of whether the requirements would be sufficient for describing presentations created by a broad range of systems. In this section we show that the AHM is able to describe the presentations created by a range of existing systems. We thus demonstrate that the AHM is a valid and useful model of hypermedia.

The systems whose documents we describe with the AHM are selected as being representative of hypertext systems, multimedia systems and, although there are to date fewer of these, hypermedia systems. For each system we give a description of its document model in its own terms, and then a brief description in terms of the AHM. A summary of the systems' documents in terms of the AHM is given in Table 3.8. Note that in this section we can only categorise elements of the implicit document models if they have been reported in the referenced articles. The summaries in each section and the overview in Table 3.8 can thus state only the presence of elements and not the absence of elements. Hence the occurrences of "yes" in the table and the absence of "no".

3.5.1 Hypertext

3.5.1.1 Intermedia

Intermedia [HKRC92] has a database of nodes of media types text, graphics, music and animation. A node of any of the media types can have anchors. A link connects two anchors, and any anchor can be the start or end of multiple links. Sets of links are stored separately as webs. A number of webs can apply to the same set of nodes. Active destinations are introduced using active anchors [PaYS90], i.e. anchors that contain a flag specifying whether the destination node should play or not. This means that while Intermedia does not incorporate time within the document model, it is able to control whether a continuous medium is played on activation or displayed only statically. The original system influenced the development of the Dexter model, and can be described with the model.

In AHM terms

In terms of the AHM, an atomic component has a presentation specification, anchors and content. Spatial information for the extent and position of the window in which the content is displayed is possibly recorded. Each anchor has an identifier, a value and a (visual) marker that is dependent on the application displaying the media item. That an atomic component, of a continuous media type, should start playing on arrival can be recorded as part of the presentation specification of the specifier of the link.

Intermedia does not support composition of atomic components. A Web is a composition of links. A link component is composed of two specifiers each with an anchor and presentation specification. There is no additional link component information. Each specifier has the implicit direction BIDIRECT, since either anchor can play the role of the beginning or end of the link. Any anchor may be referenced by multiple links.

3.5.1.2 Guide

The Guide hypertext system [OWL90] supports the creation of a number of structural objects termed reference buttons, expansion buttons, note buttons, and command buttons. Reference buttons are links to a new position in the same document or in another document, expansion buttons expand information inline, note buttons display additional information in a temporary window and command buttons execute scripts.

Guide has a number of composition types. A type of atemporal composition is the frame which allows one document to be divided into a number of subdocuments. Each subdocument uses the same display screen estate. An expansion button groups atomic components together and initially displays only the content of the expansion button. When the reader selects an anchor marker the content belonging to the corresponding child is displayed in the text-flow of the parent. The AHM has no equivalent composition model, since text flow is not part of a spatial/temporal model, although including text within the flow does require spatial position and extent information. Guide also includes an atemporal/aspatial composition for the inclusion of the destinations of note buttons. These are displayed in separate windows but are stored as part of the document.

The presentation specifications for a number of styles are hard-wired into the system itself, e.g. the destination of a note button always appears in a separate window. Other styles can be specified, e.g. a number of anchor styles are supported by the system. A source anchor marker can be highlighted by changing the cursor when it is over it, the style can change when the reader clicks and the destination anchor marker can be highlighted.

A Guide link has context. A number of expansion buttons can have an enclosing "group" object which specifies that when any child is expanded that any open child is closed.

In AHM terms

An atomic component (a button) consists of anchors and content. Each anchor has an ID and a value. The content can be text, graphics, video or script.

A composite component, an expansion button, is a hierarchical collection of atomic and composite components. This is an atemporal composition where the child is displayed within the text flow of the parent. A composition of expansion buttons, the group, is used to indicate the source context for its children. The frame is an atemporal composition grouping subdocuments which share the same spatial layout. Composition is always by inclusion and not by reference.

A link is not a separate component but is included within the anchor information in an atomic component. It consists of a source and destination specifier A source specifier (direction FROM) consists of a source context activation specification, an anchor style, an anchor and a context. The anchor style is the style for the cursor when it is over the source anchor. The context is deduced from the document structure. A destination specifier (direction TO) consists of an anchor style and an anchor. The anchor style is where the destination anchor will be shown, e.g. at the top of the window, and whether the anchor marker should highlight.

3.5.1.3 Microcosm

The philosophy of the Microcosm system [HaDH96] is to allow the creation of links among documents that are not necessarily part of a hypertext authoring environment. This requires the specification of anchors and links externally to the documents being linked. The Microcosm designers also wanted to provide links without forcing the author to create each one separately by hand. They thus provided the facility for creating links from any occurrence of a word without requiring the author to specify its position in the document. These are termed generic links.

A linkbase is collection of links. Several linkbases can refer to the same sets of documents.

In AHM terms

An atomic component consists of semantic attributes and content. The semantic attributes include the name of the file's author, and any number of keywords or author-defined attributes. The content is a reference to a file and a logical type for the file.

Anchors are not contained within an atomic component but within a link specifier. Each anchor has a value and semantic attributes.

Composition can apply to anchors and is implicit within the generic link component. Composition can also apply to links.

A link has a source and a destination specifier. The source specifier (direction FROM) is a description of one or more occurrences of an anchor within an atomic component. The destination specifier (direction TO) is a single anchor in an atomic component.

3.5.1.4 HTML

A document conforming to an HTML [Ragg97] specification contains a small number of hypertext-specific objects. The basic notion is of a single text flow including presentation information within which anchors can be defined. Anchors can contain the specification of another HTML document or a marker within the same or another HTML document. When the reader selects an anchor marker the link is followed. The browser highlights the anchor markers in different colours allowing the reader to see which link destinations have already been visited.

There is as yet no structural composition in HTML, where the "root atemporal composite" is the complete collection of documents on the World Wide Web. A form of spatial composition is included in the HTML 4.0 proposal [W3C97] called a frame. This divides up the browser window area and allows multiple text-flows to be displayed in the different areas.

In AHM terms
An atomic component (a single file) consists of a presentation specification, anchors and content. The presentation specifications consist of text-flow positioning information and text, anchor and background style information. Anchors consist of an optional identifier, a presentation specification, content and a specification of the link destination.

Spatial composition is possible using the frame construct.

Links are single source, single destination and unidirectional. The direction is from the source anchor in which the destination information is stored.

3.5.2 Multimedia

3.5.2.1 Athena Muse
The Athena Muse system [HoSA89] describes hypermedia documents through the use of a directed graph of packages, where a package can be considered as a small multimedia presentation consisting of text, video and graphics media items. The packages are nodes in a directed graph, with arcs representing possible transitions between the packages. The arcs can be used to represent hypermedia links, where cross references are fired by sending activation signals from one package to another. The network is built on a concept structure with links between high-level abstractions, although how these are translated to the package level is not described.

Each package can be described in an N-dimensional space, where, for the case of multimedia, three dimensions are sufficient for describing temporal and spatial layout. An example is the use of a timeline for attaching subtitles to a video sequence. This is done by introducing a timeline and synchronizing both the video sequence and the subtitles with respect to the timeline.

The structuring of a package does not continue down to the sub-package level, neither does it extend above the package level to group sets of packages together other than as being part of the overall application.

In AHM terms
An atomic component (a package) consists of a presentation specification, anchors, and content. The presentation specification refers to a channel and specifies a duration. It is unclear whether the spatial layout information is specified per channel or per component. The content is a reference to a media item (text, video, graphics). Anchor values can be specified within the content, in particular images.

Composition is restricted to a one-level temporal composition of atomic components and atemporal composition of temporal composite components.

There is no explicit link component, but this can be deduced from the transition arcs between the atomic components. These specify the destination of the link and the temporal and style properties of the transition.

3.5.2.2 Eventor

The Eventor system [ENKY94] has four types of objects: basic, time, composite and input. Each has common attributes such as visual forms and sync data. A visual form has information about its figures and the relations with other data objects. Sync data has information about temporal and spatial synchronizations. Basic objects point to the data and input objects have a set of input ports.

In AHM terms

An atomic component consists of a presentation specification and content. The presentation specification consists of spatial layout (extent and position) and duration. The extent and position of the object can vary with time. The content is a reference to a media item (audio, video, image, text). Anchoring is unsupported in the system.

A composite component is similar to the atomic component where instead of the content it has a list of child objects.

Link components do not exist explicitly and traversals are implemented using a user interaction builder. Jumps can be from an atomic component.

The model also includes a time object, which can be considered an atomic component without any associated content, but with a duration.

3.5.2.3 Integrator

In the Integrator [SFHS91] media items can be sequenced in virtual time, and then mapped for display onto real time through the use of a score. The data items in the multimedia presentation are placed on individual tracks (analogous to staves in a musical score and similar to channels). Timing and synchronization of multimedia items within a single track are determined by their horizontal positions on the track. Timing and synchronization of multimedia items across different tracks are determined by vertical relationships of objects across tracks (similar to synchronization arcs). As well as media item tracks the Integrator allows input or control tracks, and a timing track which allows the timing of the tracks to be altered. Altering the timing is outside the scope of the temporal information specifiable within AHM, although could be added in a system-dependent manner to the temporal information in a temporal composite component.

The duration of a non-continuous item is derived from the track information, so, for example, an image will remain on display until another image occurs on the same track. Composite objects can also be created. A composite object can be placed on the control track of the timeline, and can be opened to view the layout of objects on its own timeline. Time dependencies between objects at different levels of the hierarchy cannot be specified because of the user interface.

Several "flow" operations can be added to the control track of the timeline, including iteration and conditionally branching constructs. Linking information could be extracted from these more general controls.

Transitions are attached to the media items rather than being associated with a link, and can occur at the beginning of an item or join two objects.

It is unclear whether a jump to a different part of the presentation can be specifi ed.

In AHM terms

An atomic component consists of a presentation specifi cation and content. The presentation specifi cation consists of a start time, a deduced duration, a channel reference and a transition style. A channel corresponds to an output device, possibly a window on a video display. It is unclear whether there are properties associated with a channel and whether other properties can be associated with the atomic component. The content is a reference to a media item (still images, video, audio). Anchor values can be specifi ed and can be associated with a link destination specifi ed as part of the information stored with the atomic component.

A temporal composite component is built up from atomic and composite components, in particular serial and parallel synchronization is possible. Synchronization arcs can be defi ned.

3.5.2.4 MET++

In the MET++ application framework [Acke94] a multimedia presentation is a hierarchy of serial and parallel compositions of media items. The building blocks consist of time layout objects and media objects, where each has a start time, a duration and an associated virtual timeline. These are incorporated into a hierarchical structure with the media objects as leaf nodes and the time layout objects as intermediate nodes.

The value of an attribute can vary over time, e.g. the horizontal or vertical position of an object.

In AHM terms

An atomic component consists of a presentation specifi cation and content. The presentation specifi cation consists of a start time, duration and spatial information, in particular position, which can vary over time. The content is a reference to a media item (2D and 3D graphics, images, text, audio, video and user interface components with event handling).

A temporal composite component is a hierarchy of atomic and composite components.

3.5.3 Hypermedia

3.5.3.1 Videobook

The Videobook system, [OgHK90], combines time-based, media composition with linking, allowing the construction of composite multimedia nodes among which a reader can navigate. A scene, consisting of media items and triggers, is played according to the timing and layout presentation parameters associated with its children. Each scene can contain nested sub-scenes. Synchronization of objects is specifi ed by giving the start time of an object with respect to the scene. A trigger object, when selected by the reader, sends a message to its target object which is either displayed if it is a node or executed if it is a process.

In AHM terms
An atomic component consists of a presentation specifi cation and content. The presentation specifi cation consists of a duration, an extent and a position. The content is a media item (e.g. text, image, video, script). An anchor consists of an anchor value and a destination reference. An anchor's value is in terms of its extent, position and duration with respect to the atomic component. The destination of the implicit link is a temporal composite component or an atomic component.

Composition is temporal or atemporal. Temporal composition is a hierarchical collection of atomic and temporal composite components. Atemporal composition is a collection of temporal composite components (scenes, referenced by name).

3.5.3.2 Harmony
Harmony [FSMN91] integrates continuous media items into a hypertext system. Each object is considered a node and there are links between nodes. Links are used for expressing the timing relations between nodes. The notion of an object group is introduced, where, if an object group is the destination of a link, a message is broadcast to all members of the group when the link is traversed.

In AHM terms
Atomic component has content (text, music, graphics, video and animation) and associated procedures (in the object oriented sense). Anchors can be specifi ed in text, video and graphics media types.

A composite component is a temporal (parallel or serial) composition of atomic and composite components.

A link component is a separate component which also describes timing information.

3.5.3.3 HyTime
HyTime [ISO97b], [ISO97b] is a standard for representing the presentation independent structure of hypermedia documents. It embodies its own model of hypermedia, which includes complex hyperlinking, locating of document objects and the scheduling of objects within measured coordinate spaces such as space and time. As a meta-model, HyTime can be used to specify hypermedia models, such as the AHM. HyTime does not, however, provide constructs for presentation specifi c aspects of documents, which form an important part of the AHM. We have created extensions to HyTime architectural forms to express commonly found behaviour in hypermedia systems, in particular for modelling aspects of the runtime layer. These extensions comprise the Berlage[5] architec-

ture: shadow location form, aggregate link form, and duration marker form [ROHB97b].

In AHM terms
An atomic component and all its subparts can be expressed. Temporal and atemporal composite components and all their subparts can be expressed. Linking can be expressed. Channels can be expressed, although HyTime does not address this issue as directly.

The activation/deactivation cannot be expressed, hence the introduction of the Berlage aggregate link, which specifies that all the other children of a composite are deactivated whenever one of them is activated. This is a particular type of atemporal composition implemented in CMIFed, called the choice component.

Synchronization with objects of unknown duration cannot be expressed in HyTime directly, hence the introduction of the duration marker form.

3.5.3.4 MHEG-5
MHEG-5 [ISO97a], [JoRo95] was developed to allow a single representation of a multimedia presentation to be played on a range of end-user platforms over a distributed network. The presentation can be divided up into separate fully encapsulated parts which can be communicated separately, minimizing network traffic and the transfer of unnecessary information. MHEG-5 is defined as a collection of related object-oriented data structures and is a procedural language for which player software can be implemented.

In AHM terms
An atomic component has presentation specifications, anchors and content. The presentation specifications include start time and duration, a layout channel reference and style information. Anchors have an identifier, style information and a value. Content is a reference to a media item.

Temporal composition can be specified including presentation specifications and children. The presentation specifications include a duration, synchronization constraints and style. Atemporal composition is specifiable, in that non-temporally related presentations can be linked to.

A link specifier is implicit within the procedural specification of links. It includes a source context activation state with the choice of CONTINUE or REPLACE.

3.5.3.5 SMIL
SMIL, Synchronized Multimedia Integration Language, provides a declarative way of specifying multimedia documents for the World Wide Web. A specification of the language [BuRL91] and a high-level description [Bult97] are available. The work of the AHM and CMIF played a major role in the development of the

5. Berlage is an important Dutch architect.

requirements for the language. The SMIL requirements for simplicity of development of a player and readability of the syntax, however, have generally lead to simplifications of the model. One area where SMIL is broader than the AHM is the specification of alternative data formats for dealing with the delivery of the same document specification across differing network bandwidths.

In AHM terms
A channel has a position and an extent in terms of a browser window. An atomic component has presentation specifications and content. The presentation specifications include start time, duration and a channel reference and are able to refer to style information (using style sheets). Content is given by a reference to a URL. Atomic anchors can be specified or referred to in the content. Temporal composition is specifiable and is of two types: parallel and sequential. Synchronization arcs can be defined between direct children of a composite. A link has a source and a destination, where each has a component, an anchor and the source has a source context activation state. A source and destination anchor can be specified in terms of a reference to an anchor defined within the content of an atomic component, or in terms of a temporal/spatial anchor value, as described in Section 3.4.3.

3.5.4 Summary of system models expressed in AHM

AHM model elements			Hypertext: Intermedia	Hypertext: Guide	Hypertext: Microcosm	Hypertext: HTML	Multimedia: Athena Muse	Multimedia: Eventor	Multimedia: Integrator	Multimedia: MET++	Hypermedia: Videobook	Hypermedia: Harmony	Hypermedia: HyTime[a]	Hypermedia: MHEG-5	Hypermedia: SMIL
Channel				y[b]		y[b]	y		y				y		y
	Pres. Spec.	Channel ref.											y		
		Position, extent		y		y	y		y				y		y
		Style		y		y									
	Attributes														
	Med. typ.						y						y		
Atomic Comp.			y	y	y	y	y	y	y	y	y	y	y	y	y
	Pres. Spec.	Duration					y	y	y	y	y		y	y	y
		Channel ref.					y						y		y
		Position, extent	?[c]					y[d]	y	y[d]	y		y	y	y
		Style	n[e]			y			y[f]					y	y
	Attributes				y	y			y[g]						
	Anchors		y	y	y[h]	y	y						y	y	
		Anchor ID	y	y		y							y		y
		Pres. Spec.				y								y	
		Attributes			y	y									
		Value	y	y	y	y	y		y		y	y	y	y	y
	Content	Media item ref.	y	y	y	y	y	y	y	y		y	y	y	y
		Data-dep. spec.											y		
Temporal Composite Comp.							y[i]		y	y	y	y	y	y	y
	Pres. Spec.	Duration								y			y	y	y
		Sync. arcs							y			y	y	y	y[j]
		Style												y	
	Attributes														
	Anchors												y		
		Anchor ID											y		
		Pres. Spec.													
		Attributes													
		List of anchors											y		
	Children	Comp. ref.					y		y	y		y	y	y	y

TABLE 3.8. Implicit document models of systems expressed in AHM

AHM model elements			Hypertext				Multimedia				Hypermedia				
			Intermedia	Guide	Microcosm	HTML	Athena Muse	Eventor	Integrator	MET^{++}	Videobook	Harmony	HyTime[a]	MHEG-5	SMIL
Atemporal Composite Comp.			y^k	y^l	y^m		y				y		y	y	
	Pres. Spec.	Init. activ. state		y										y	
		Play/pause													
		Style													
	Attributes														
	Anchors				y^n								y		
		Anchor ID											y		
		Pres. Spec.													
		Attributes													
		List of anchors			y^o								y		
	Children	Comp. ref.	y	y									y	y	
Link Comp.			y		y							y	y		y
	Pres. Spec.	Duration					y					y			
		Rel. position											y		
		Style	y												
	Attributes				y										
	Anchors		Links to links are not considered.												
		Anchor ID													
		Pres. Spec.													
		Attributes													
		Value													
	Specifiers		y^p	y	y	y							y		y
		Src. cont. activ.		y									y^q	y	y
		Dst. cnt. pl./pa.	y												
		Style		y											
		Anchor	y	y	y	y	y	y^r	y		y	y	y		y
		Context		y^s		y^t							y		
		Direction	y^u		y^v		y^w	y^x			y^w		y		y

TABLE 3.8. Implicit document models of systems expressed in AHM

a. The table is filled in for HyTime as " can these elements be expressed directly using HyTime or SGML", and not as " does the HyTime model include these objects".
b. Areas can be defined within the main window
c. It is unclear whether this is recorded in the document or is decided at runtime.

d. The position can vary with time.
e. The anchor style is determined per media type rather than per atomic component.
f. In particular the transition special effect at the beginning or end of the display of the content.
g. In particular a user-specifiable name.
h. But part of link specifier, not of atomic component.
i. Temporal composition is one level only.
j. Only between direct children.
k. Atemporal composition is of links, called a web.
l. An "include in text-flow" composition or a group composite.
m. Atemporal composition is of links, called a linkbase, or of anchors, contained within the link.
n. Anchor composition is specified within a link specifier.
o. Of the form: string specification in file "anywhere" at position "anywhere".
p. A link has two specifiers.
q. By means of the Berlage aggregate link form.
r. Where the anchor reference is whole component reference.
s. Source context is derived from the document structure.
t. The frame structure can be used to derive context.
u. Every specifier is BIDIRECT.
v. A link has one source and one destination specifier.
w. The link component does not exist as such, but the relevant information can be extracted and stored as a link with two specifiers, one with direction FROM and the other TO.
x. It is not clear, but probable, that the implicit link has one FROM and one TO specifier.

3.6 Conclusions

The Amsterdam Hypermedia Model has been designed to capture the elements of structure, timing, layout and interaction that are needed to specify an interactive, time-based on-line presentation. The AHM is not a perfect model of hypermedia, but instead seeks to achieve a balance of expressibility and implementability. In this chapter we have described the choices we have made, and motivated these choices. We have defi ned the model and shown that it can be used to describe the presentations created by a wide range of systems. We have shown that the model is suffi ciently simple to be implementable by providing a description of the parts of the model which have been implemented in the CMIFed environment, Appendix 1.

The main components of the AHM are the channel, atomic component, temporal and atemporal composite components and the link component. Presentation specifi cations that can be associated with these components are selected from temporal, spatial, style and activation information.

Although AHM has its roots in the Dexter and CMIF models it incorporates the following novel extensions:

- The presentation specifi cations within the model have been explicitly stated as temporal, spatial, style and activation information. Each aspect occurs throughout the model and we have shown how the occurrences relate to one another.
- Anchors have been extended to include semantic attributes and presentation specifi cations, including start time and duration for an atomic anchor of a non-continuous media type.
- Content is specifi ed explicitly as a media item reference along with a corresponding data-dependent specifi cation.
- Anchor reference and channel reference in addition to a component reference are used throughout the model.
- Composition of anchors has been introduced.
- Composition of components is of two types: temporal and atemporal. Dexter expressed only atemporal and CMIF expressed temporal. Including both types of composition within one model requires the inclusion of activation state information.
- Activation state information has been incorporated throughout the model. This includes: initial activation state, play/pause state and change in activation state on following a link.
- Link components have been extended to include context in the link specifi er
- Transition information, including transition duration and special effect, has been incorporated in the model.

The model could be extended in the following directions.

- Separate out the spatial layout hierarchy from the channel element and make a reference to it from a channel.
- Introduce spatial layout with respect to content.
- Include a way of selecting between synchronized streams of information, e.g. by introducing a visible/invisible state.
- Include a timeline more explicitly.
- Separate out link semantic direction from link traversal direction.
- Allow the specification of anchors in terms of time and space for atomic and temporal composite components.
- Allow the inclusion of absolute time within the model. This could be associated with children of atemporal composites but would require an extension to the current link activation mechanism.
- Include auto-firing of links.

A formal description of the model in the Object Z language is given in [OsEl97], included as Appendix 2 of this thesis.

The following chapters investigate the authoring aspects of hypermedia presentations.