Invited paper

Adaptive Multimedia Documents

Stefan Wirag and Kurt Rothermel

University of Stuttgart, Institute of Parallel and Distributed High-Performance Systems (IPVR), Stuttgart, Germany

Multimedia documents are of importance in a wide range of application areas, such as education, training, advertising and entertainment. Since multimedia documents may comprise continuous media, such as audio and video, the presentation of those documents may require a significant amount of processing and network resources. The amount of resources available during a presentation depends on the system configuration and the current system load. Hence, it can happen that there are not enough resources to render a multimedia document according to the specification, resulting in a reduced presentation quality, if the presentation is possible at all. To cope with those situations, different versions of the same document can be specified, one for each potential configuration or probable load situation. A better approach is to have only one document that can be adapted to different system configurations and load conditions. To enable this approach, multimedia documents have to be specified flexible so that different presentations can be compiled depending on the resource situation. In this paper, we analyse the adaptivity of multimedia documents and investigate to what extend existing document models support this type of adaptivity.

Keywords: multimedia document models, synchronization, adaptivity

1. Introduction

Due to their expressive power, multimedia documents have become attractive for many application areas, such as education, training, advertising or entertainment. Multimedia documents combine continuous media objects (e.g., video, audio or animation) and discrete media objects (e.g., graphic or text). Besides its media objects a document also includes temporal and spatial layout information. In a distributed environment, documents are typically stored on servers, from which they are retrieved for presentation. The actual presentation takes place at a so-called presentation terminal (or terminal for short), which might be a work station, a PC, a Set-Top-Unit, or even a mobile device in future times. When a user initiates the presentation of some document at a terminal, the terminal takes over the responsibility for orchestrating this presentation, i.e. it schedules the access to remote servers, the playout of individual data units, the activation of interaction intervals allowing the user to impact the presentation, and so forth.

In order to present a multimedia document, a certain amount of resources is needed. For example, processing and buffer resources are needed at the terminal to present the document, while network resources are required to transfer the media objects associated with the document from the server to the terminal. The amount of resources available strongly depends on the system configuration and the current system load.

Let us consider the system configuration first. Clearly, a workstation connected to a high-speed network allows for a higher quality presentation than a PDA linked to a radio network. One approach to overcome this problem of heterogeneity is to have different versions of the same "logical" document, one for each potential configuration. Another approach is to have one document that is able to adapt to the capabilities of the underlying system. We prefer the second approach since it avoids redundancy and does not have to predict numerous configurations. For example, with the second approach a digital library would have to store only one version for each document, without taking into account the capabilities of the terminals potentially used to access the library.

Even if we only consider one type of terminal and one type of network, variations in the system load may cause different amounts of resources to be available. Without resource reservation, the resources available for a presentation may change while the presentation is in progress, i.e., resource shortages cannot be avoided. If the underlying system provides for resource reservation, the resources needed to present a (mono-media) object (e.g., a video clip) can be determined and reserved prior to its presentation. However, reservation in general cannot prevent resource shortages to occur during the presentation of interactive multi-media documents (Gecsei (1997)). For interactive documents, it is not feasible or even possible to reserve all resources required to render the entire document in advance. Rather, resources are reserved and released incrementally while the presentation progresses, which again can lead to resource shortages. Consequently, even in the case of resource reservation it may happen that less resources are available than expected.

When a resource shortage occurs, there are basically three ways to react to it. Firstly, the presentation is aborted, which is the only choice if the underlying system is non-adaptive. Secondly, the quality of the presentation is degraded *somehow*, meaning in a system-controlled manner. In this case, the underlying system is adaptive, however the author of the document cannot impact how the quality is degraded. Thirdly, the presentation is adapted in a *user-controlled manner* to the given resource situation. To enable the last alternative, flexible document models are required, which allow to compile different presentations from a given document specification, depending on the resource situation.

Flexible document models provide for the specification of presentation alternatives with different resource demands. Presentation alternatives can be represented by media objects of different type, e.g. subtitles can be defined as a "lower quality" alternative for a speech sequence. Also alternatives may be of the same media type but differ in the quality of presentation, e.g. the desired playout rate of a video may be 20 frames per second, which may be reduced to 10 frames per second due to resource shortages.

The focus of this paper is temporal adaptivity in the context of orchestrated multimedia presentations. In recent literature, various document models and synchronization concepts have been proposed for multimedia presentations. There exist several variations of time-line models: Computer Inc. Apple (1991), Blankowski et al. (1991), Drapeau (1993), Hamakawa et al. (1994), Gibbs (1991), HyTime (1992), MHEG (1995). In these models media objects are aligned on time axes (time-lines). Another popular approach is to express temporal relationships of presentation elements by means timed petri-nets: Coolahan and Roussopoulos (1983), Little and Ghafor (1990), Qazi et al. (1993), Prabhakaran and Raghavan (1993), Senac et al. (1996). Further, there exist approaches where the temporal layout of a presentation is defined as a net of interrelated instants or intervals: Buchanan and Zellweger (1993), Candan et al. (1996), Schnepf et al., (1996), Wahl and Rothermel (1994). An extensive overview of media synchronization concepts and their features is given by Perez-Luque and Little (1996). In this paper, we investigate the temporal adaptivity of multimedia documents and discuss the document models that provide for flexible specifications.

The remainder of the paper is structured as follows: In Section 2 we analyse types of adaptivity of multimedia presentations. In Section 3, existing flexible document models are presented and compared. Finally, we conclude with a summary of our results.

2. Types of Adaptivity

In this section, we will discuss various types of adaptivity of multimedia documents. Before that we have to introduce a general document model.

A document comprises a set of media items. To permit expressive presentations, documents must include a rich variety of media types, such as text, images, audio, video and animations. In the document model assumed below, each media item is represented by a single media object, which is associated with a set of attributes. Which attributes are associated with an object depends on the object's media type. To describe temporal aspects, single media objects representing continuous media items have the two attributes *duration* and *playout speed*. While the first attribute describes how long the media item is presented, the latter specifies how many

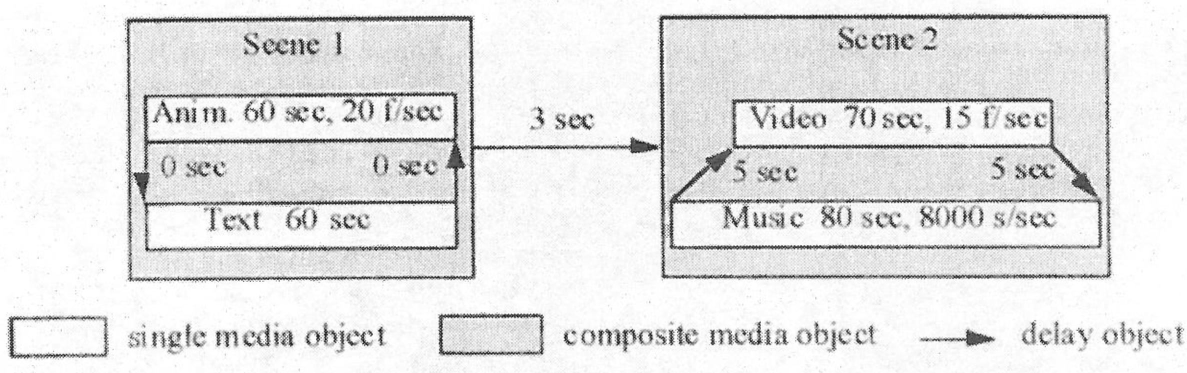

Fig. 1. Example specification

data units of the media item are displayed per second during its presentation. Clearly, a single media object representing a discrete media item has only the attribute duration.

Temporal relationships define how single media objects have to be combined temporally to make up a presentation. In our model, *delay objects* describe temporal relationships between the start and the end of single media objects. In other words, delay objects represent a delay, which is defined by an attribute.

Besides single media objects the model includes *composite media objects*. A composite media object is composed of a number of single media objects, other composite media objects and the delay objects relating to these media objects. A composite media object can represent a document fragment such as a scene or a page. Composite media objects appear to the surrounding composite media objects like a single media object.

Figure 1 shows a document specification according to this model. The example consists of two composite media objects presented after each other with a delay of 3 seconds. The first composite media object consists of two single media objects, an animation presented for 60 seconds simultaneously with a textual explanation. The animation will be presented with 20 frames per second. The second composite media object consists of a video, which is presented for 70 seconds with 15 frames per second, and a piece of music presented for 80 seconds with 8000 samples per second. The music starts 5 seconds before the video and ends 5 seconds after the video.

Using this document model as a basis, adaptivity is possible on the object level and the attribute level.

2.1. Adaptivity on the Object-level

Documents supporting adaptivity on the object-level may contain alternative single media objects, which represent the same or similar information in different form. Depending on the resource conditions, the underlying presentation system can choose the appropriate alternative. For example, information on a technical subject might have a number of alternative representations, including an explanation in form of an audio, a textual description and an explanation in form of a video. As different information forms have various resource requirements, the presentation system can handle different resource situations. But adaptive documents might not only contain alternative single media objects. If the definition of presentation alternatives demands that several single media objects are related by delay objects in a different manner, composite media objects representing alternative presentations might be specified.

Adaptivity on the object-level is not restricted to media objects, documents might also offer completely different temporal relations between the same media objects. For example, sometimes it makes no difference whether an advertising animation is presented before or after a videoclip. Such alternative temporal relations might cause a different overlapping of media objects and therefore may cause different resource requirements. Hence, a document may contain alternative delay objects.

Often, multimedia presentations contain media objects that are "nice to have" but contribute only little to the intended information transfer, such as an animation of a company logo on top of a page. The playout of such information can be omitted when there are not enough resources available. To describe those features, document models have to integrate appropriate object-level flexibility abstractions.

2.2. Adaptivity on the Attribute-level

If adaptivity on the attribute-level is supported, documents allow for alternative presentations even if there is no adaptivity on the object-level. For example, it might be specified that a video normally presented with 20 frames per second can alternatively be presented with 15 frames per second, or that the duration of an animation might be between 5 and 10 seconds. This type of adaptivity enables the presentation system to react in a controlled manner to changing resource situations by adapting the playout speed or the duration of a presentation accordingly.

There are several ways to dynamically shrink or extend the duration of presentations, such as modifying the playback rate, discarding or repeating data, or performing an alternative action when there is no more data to be presented (Steinmetz (1990)). Unfortunately, these techniques often reduce the overall quality of the presentation and hence cannot be used indiscriminately. However, within certain ranges, the physical effects of these modifications may not noticeably reduce quality.

Attribute-level adaptivity is also possible for temporal relations because sometimes the start or end instant of a media object can vary within certain limits. For example, if the presentation system is told that the time between a video sequence ends and a succeeding animation starts is allowed to be between 1 to 5 seconds, it has some flexibility to choose the appropriate schedule. In summary, adaptive temporal relations allow to delay the playout of media objects in certain limits, which gives some flexibility for optimizing the temporal alignment of media objects with regard to resource consumption.

2.3. Controlling of Adaptations

As discussed above, adaptive documents enable a presentation system to choose between alternative presentations to cope with changing resource constraints. In a particular resource situation, a presentation system might have different options to adapt a presentation. For example, in a document that is adaptive on the object-level, it might be possible to present either subtitles or a speech as explanation to an animation. If the document supports attribute-level adaptivity, it might be possible to reduce the speed of either a video or an animation. From a system's perspective it is not always obvious which presentation alternative provides for the better quality, since the notion of quality is highly application-dependent. Consequently, the author of the document should be able to express his or her preferences with regard to the various alternatives, depending on the respective application. Hence, the author should be able to assign *quality measures* (e.g., priorities) to each of the alternatives, which guide the underlying system when making its decisions between alternatives.

With regard to the course of a presentation, it has to be defined when and under which circumstances the presentation system can choose a certain presentation from the set of alternatives. For example, under some circumstances it might be desirable that the presentation system continues with the playout of subtitles when the playout of the audio is no longer possible due to resource shortages, sometimes this might not be desirable at all.

With interactive documents users may jump from one part of a presentation to another, where jumps may cause parts of a document to be rendered more than once. In some cases, it may be desirable that the same alternatives and temporal arrangements are selected in successive presentations of the same document part. For example, let us assume a document contains a scene where a textual description can be presented instead of a video-clip. If in the first presentation of the scene the textual description was presented, the user might be confused about a video presentation when he jumps back to the scene. Here, the presentation system should select the same alternative in each rendition of the scene.

Analogous requirements may hold for the attribute-level adaptivity. For example, it may be required that the presentation speed of a continuous media object never changes while the presentation is in progress, it may be only

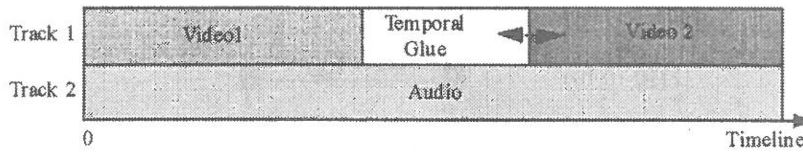

Fig. 2. Specification with temporal glue

changed (in a certain range) before the presentation starts. Moreover, if document parts are rendered more than once, it might be desirable that some object behavior, e.g. the presentation duration of a single media object, is equal in each rendition.

To enable authors to control when the presentation system is allowed to choose a certain presentation alternative or switch to another presentation alternative, so-called *selection policies* have to be specified and associated with the sets of alternatives.

3. Flexible Multimedia Document Models

Various temporal models and synchronization concepts have been developed for multimedia presentations. In this section, we analyze a number of flexible temporal specification techniques that have appeared in recent literature.

3.1. Temporal Glue

The multimedia presentation builder Mbuild (Hamakawa et al. (1992), Hamakawa et al. (1994)) adopts a time-line approach, where media objects are positioned on different tracks (Figure 2). Delays between media objects in a track are represented by temporal glue objects. Temporal glue acts like TEX's (Knuth (1986)) spatial glue and has similar flexibility parameters controlling its stretchability and shrinkability. In the Mbuild system, media objects have a predictable duration, and only temporal glue objects are flexible. Interaction is not supported by this system at all.

The scheduling algorithm of Mbuild exploits the flexibility given by temporal glue objects to compute the optimal layout of a presentation, where a presentation is defined to be optimal if all glue objects in a track have the same extension relative to the specified stretchability and shrinkability. However, the scheduling algorithm does not take into account the underlying system's resource situation. Consequently, this type of adaptivity only simplifies the specification of documents.

3.2. Firefly

The multimedia presentation system Firefly (Buchanan and Zellweger (1992), Buchanan and Zellweger (1993)) is event-based: A media object is described by a start- and an end-event. An event is called an internal event if it lies between a start- and the corresponding end-event. Internal events mark certain instants that can be temporally related to other events. A third type of event is an asynchronous event, which allows to model unpredictable behavior, such as user interaction. Figure 3 shows the graphical representation of this model.

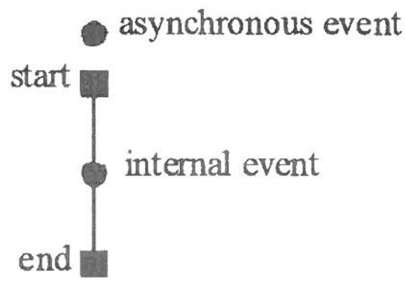

Fig. 3. Modeling of media objects

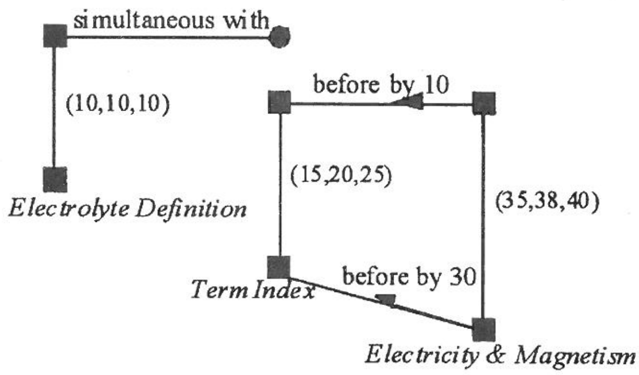

Fig. 4. Firefly example specification

Delays between events in a media object are described by a minimum, optimal and maximum value. Additionally, costs are defined for shrinking the delay between two events to the minimum or extending it to the maximum. In other words, Firefly integrates adaptivity on the attribute level. Moreover, Firefly proposes flexible temporal relations to specify temporal dependencies between events. It can be expressed that two events have to occur simultaneously or that there is a (flexible) delay between two events. It is not possible to assign costs to flexible delays. Firefly provides for simple button-oriented interaction but integrates no abstractions to compose composite media objects to structure a specification. Hence, complex specifications are quite hard to handle.

Figure 4 shows a document specification according to the Firefly model. In the specification, the *Term Index* and the *Electricity & Magnetism* media objects represent the main track. The duration of the *Term Index* presentation is flexible, it can have a minimum of 15 seconds and a maximum of 25 seconds, the optimal duration is 20 seconds. The presentation of the Electricity & Magnetism object will be started 10 seconds after the Term Index and will be finished 30 seconds after it. The duration of the *Electricity & Magnetism* presentation is also specified flexible. The asynchronous event above the Term Index object indicates that it is possible to start a *Electrolyte Definition* object of fixed duration by interaction.

The focus of Firefly was to provide a powerful technique to ease the temporal specification of documents. Therefore, Firefly exploits the flexibility given by a specification to compute the optimal layout of a presentation. The layout of a presentation is defined to be optimal if the accumulated costs for shrinking and extending — if any — are lowest. Let us assume that in the example in Figure 4 the costs to shrink the duration of the *Term Index* are lower than to shrink the duration of the *Electricity & Magnetism* object. Then Firefly would implement a duration of 38 seconds for the *Electricity & Magnetism* object and a duration of 18 seconds for the *Term Index*. However, the resource situation of the underlying system is not considered when the layout of a presentation is computed.

3.3. CHIMP

In CHIMP (Candan et al. (1996)), the temporal layout of a presentation is modeled in a flexible manner by so-called *difference constraints*. Difference constraints are linear constraints of the following form:

$$x_1 - x_2 \leq b$$

By using a variable to denote a presentation event (start time st of an object presentation or end time et of an object presentation), difference constraints may be used to create a flexible temporal specification, describing the possible range of values between two events. For example, let us assume that there exist two media objects o_1 and o_2 with durations 40 and 50 seconds, respectively. We want o_2 to be presented after o_1, and the presentation of o_2 to be started within 10 seconds after o_1 finishes. This can be described using the following constraints:

$$\begin{array}{ll} st(o_1) - et(o_1) \leq -50 & et(o_1) - st(o_1) \leq 50 \\ st(o_2) - et(o_2) \leq -40 & et(o_2) - st(o_2) \leq 40 \\ st(o_2) - et(o_1) \leq -10 & et(o_1) - st(o_2) \leq 0 \end{array}$$

The temporal layout of a presentation is described is this way, because difference constraints allow to represent a specification as directed graph and an efficient shortest path algorithm can be applied to generate a presentation schedule. CHIMP does not support user interaction and composite media objects.

CHIMP allows to specify alternative constraints and to assign priorities to them. Hence, the model allows for adaptivity not only on the attribute-level but also on the object-level. For example, let us assume that two media objects o_1 and o_2 with durations of 30 respectively 20 seconds should not overlap in a presentation. Moreover, we want that the delay between o_1 and o_2 is not greater than 5 seconds and that the presentation of o_1 before o_2 is preferred, but it should also be allowed to present o1 after o_2. This can be described by the following constraints:

$$\begin{array}{ll} st(o_1) - et(o_1) \leq -30 & \\ st(o_2) - et(o_2) \leq -20 & \\ st(o_2) - et(o_1) \leq 5 & \text{Priority: 80} \\ st(o_1) - et(o_2) \leq 5 & \text{Priority: 40} \end{array}$$

$$\begin{array}{ll} et(o_1) - st(o_1) \leq 30 & \\ et(o_2) - st(o_2) \leq 20 & \\ et(o_1) - st(o_2) \leq 0 & \text{Priority: 80} \\ et(o_2) - st(o_1) \leq 0 & \text{Priority: 40} \end{array}$$

To describe that o_1 can be presented before or after o_2, constraints representing both arrangements are introduced. To express that the order where o_1 is presented before o_2 is preferred, higher priorities are assigned to the constraints representing this arrangement.

In CHIMP presentation schedules are compiled considering resource constraints. The scheduler tries to find start- and stop-instants so that the presentation does not require more resources than available. When computing a schedule, the scheduler tries to satisfy as many difference constraints as possible. If not all constraints can be satisfied, constraints with low priorities are omitted first.

3.4. TIEMPO

Tiempo is an interval-based model (Wahl and Rothermel (1994), Wirag (1997)), in which documents are composed of single media objects, such as video or text, and composite media objects, such as scenes. A presentation object is modeled by a temporal space, a presentation interval and a projection. The *temporal space (TS)* represents the content and layout information associated with the media object. The *presentation interval* represents the period in which the media object is presented. The *projection* describes which and how many data units of the TS are presented per second in the presentation interval. The concept of a projection allows to present media objects with other than recording-time properties. Interaction objects, such as buttons or sliders, have additional *interaction intervals*. Each such interval represents the period in which a particular user-interaction (e.g. click with mouse, insertion of text) is accepted. An interaction interval is described by its extent and the state that changes when the user triggers the associated interaction.

The TS of a single media object is defined by the (temporally ordered) data units of the object, e.g., the TS of a video object consists of a temporally ordered set of frames. The TS of a composite media object is composed by a number of single or composite media objects. A top composite object's TS represents the whole multimedia document with the associated presentation interval. To define the temporal layout of a composite TS, presentation intervals of included media objects are arranged by *interval operators* (Wahl and Rothermel (1994)).

Figure 5 shows the interval operators defined in Tiempo. In the simple specification example in Figure 6, the interval operator "before" specifies that the animation should be started 1 second after the video has ended. The "while" operator describes that the presentation of the music starts and ends simultaneously with the video. To define the temporal layout of possible interactions, interaction intervals are related to other intervals by interval operators.

An interaction is described as a so-called *reaction relation* between an interaction interval and the affected elements. A reaction relation consists of a trigger-condition and actions. The actions describe behavior modification of projections, presentation or interaction intervals. Actions can start and stop intervals, pause, freeze and speed up the playout (Wahl et al. (1995), Wirag et al. (1995)). The trigger-condition describes how the state of the interaction interval

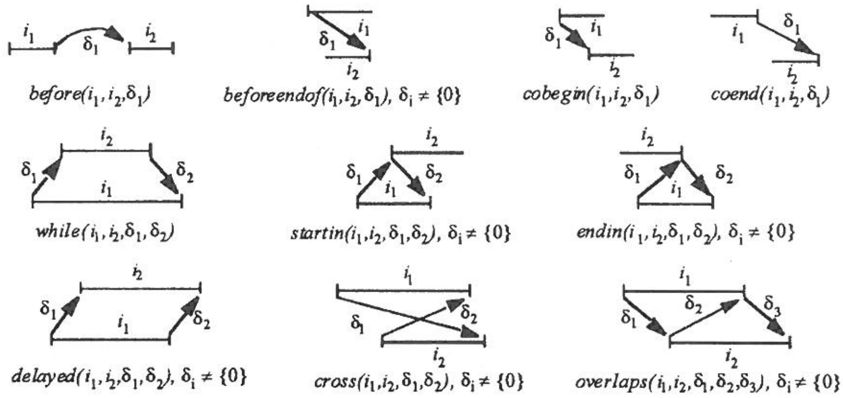

Fig. 5. Interval operators

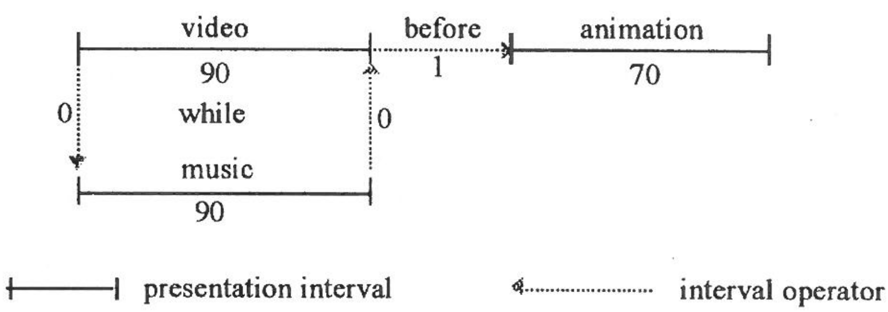

Fig. 6. Tiempo specification example

must change (e.g. a state change of an interaction interval of a button from "not-pressed" to "pressed") before the actions are executed.

The Tiempo model allows for adaptivity on the object- and attribute-level (Wirag (1997)). Adaptivity on the object level can be specified by so-called *selection groups*. A selection group contains a number of presentations that can be selected alternatively. Whenever a selection group is performed, the underlying system selects and presents exactly one of these alternatives. A presentation alternative can be a single or arbitrarily complex composite media object. Selection groups can be nested to specify selections at various levels of abstraction. Moreover, priorities can be assigned to presentation alternatives to indicate which alternative should be preferred when more than one can be implemented. The choosing of presentation alternatives is controlled by the following selection policies:

- *first*: For a selection group associated with this policy only one selection can be made per document presentation. The selection is performed when the selection group is rendered the first time. If it is rendered more than once, always the presentation alternative selected first is to be implemented.
- *static*: With this policy one selection can be made per rendition of the selection group. Once a selection has been made it cannot be changed during the presentation of the selection group is in progress. In contrast to the previous policy, however, different alternatives can be selected for each rendition of the selection group if it is rendered more than once.
- *dynamic*: With this policy the selection may be changed even while the presentation of the selection group is in progress.

Figure 7 shows an abstract representation of a document part with nested selection groups. The outer group includes two presentation alternatives. While the first alternative is a text object, the second one consists of an animation object and another selection group. The inner selection group provides three presentation alternatives, a speech object, a subtitle sequence and an empty object. According to the selection

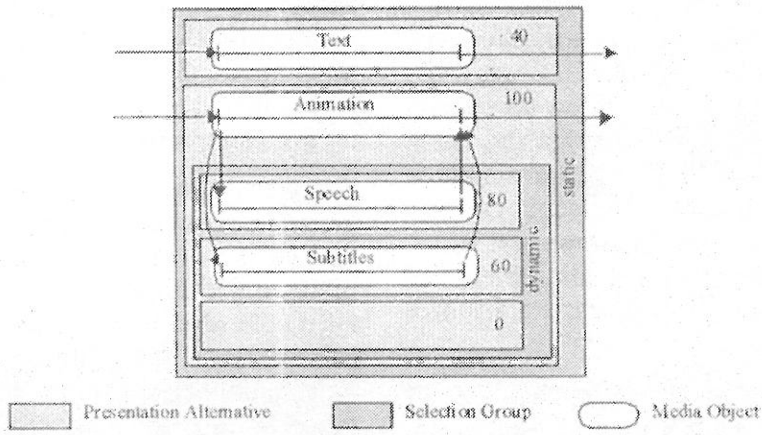

Fig. 7. Document specification with selection groups

group semantics, the presentation system has the option to present the text or the animation with either the speech sequence, the subtitles or no further explanation.

For the outer selection group the selection policy is *static*, which means that it is not allowed to switch between the text presentation and the animation while the presentation of this selection group is in progress. The selection policy for the inner selection group is dynamic, i.e., it is possible to switch between the speech, subtitle and 'empty' presentation whenever it is required. For example, when there is a resource shortage, the presentation can switch from speech to subtitles or may even omit subtitles also. In the outer selection group, the alternative with the animation has the highest priority (100) and hence should be selected provided the required resources are available. In the inner selection group, the speech object is the preferred alternative (priority 80), followed by the subtitles (priority 60) and the empty object (priority 0).

In Tiempo, adaptivity on the attribute-level is modeled by so-called *Quality of Service (QoS) ranges*. QoS range arguments can be used to specify the extent of presentation and interaction intervals, the presentation speed and the delays implied by interval operators. Based on this information, a presentation system can select extent, speed or delay values within the specified QoS ranges according to the current resource situation. To indicate which values of QoS ranges are preferred, priorities are assigned to the contained values. The priority structure of a QoS range is defined by anchor points. The QoS range

$$[10 : 30, 35 : 100, 55 : 70]$$

might define the extent of a presentation interval, which can be between 10 seconds and 55 seconds. In the example, an extent of 10 seconds has the priority 30, an extent of 55 seconds has the priority 70, and an extent of 35 seconds has a priority of 100. With the extents between 10 and 35 linear increasing priorities from 30 to 100 are associated, and with the extents between 35 and 55 linear decreasing priorities from 100 to 70 are associated. Here, the presentation system should implement an extent of 35 seconds for an optimal presentation quality. To define when a value from a QoS range can be chosen, the same selection policies can be assigned to QoS ranges as to selection groups.

Based on the concepts described above, an adaptive scheduling algorithm for environments with the best-effort service model was developed (Wirag (1997)). This algorithm is able to dynamically adapt presentations taking into account the availability of resources as well as the QoS constraint specified by the user. When computing adaptations, the scheduling algorithm tries to find presentation alternatives of selection groups and values of QoS ranges such that the presentation does not need more resources than available, and the sum of priorities of the chosen presentation alternatives and QoS range values is maximized.

3.5. Discussion

Table 1 shows a comparison of the presented document models with regard to their abilities to express adaptivity.

All temporal models discussed in this section provide for adaptivity on the attribute-level. Mbuild offers flexible temporal relations, while Firefly, CHIMP and Tiempo allow to specify both flexible presentation durations and flexible temporal relations. In Tiempo, presentation speeds can also be flexible. Hence, in contrast to the other models, Mbuild allows not to express that the rendition of a media object can be finished premature to compensate a resource shortage.

Firefly offers quality measures for flexible presentation durations, while Tiempo supports quality measures for presentation durations, delays and presentation speeds. The quality measure concepts of Firefly and Tiempo are equivalent. In both models, a certain quality can be assigned to each value of a value range. Tiempo additionally defines selection policies on the attribute level. Whereas in Mbuild and CHIMP it is not possible to control the selection of presentation alternatives on attribute-level explicitly. All possible variations of durations or delays are treated as equal.

Adaptivity on object-level is only provided by CHIMP and Tiempo. In both models it is possible to specify alternative temporal relations. In contrast to CHIMP, Tiempo supports the specification of alternative media objects on different levels of abstraction. Hence, Tiempo enables authors to define presentation alternatives dependent on the selection of other presentation alternatives. Further, the definition of alternative composite media objects is only provided by the Tiempo model.

The possibility to omit less relevant media objects can be specified explicitly in Tiempo only. In CHIMP, difference constraints representing media objects are omitted automatically if not all media objects can be presented simultaneously. Hence, it might happen that even media objects important for the content of the presentation are not rendered. Both models allow to assign priorities to alternatives, but only Tiempo supports the concept of selection policies.

The lack of adaptivity on the object-level in Mbuild and Firefly stems from the fact that both models were developed to ease document specification rather than supporting adaptable presentations.

4. Conclusion

Documents can be adaptive on the object-level and on the attribute-level. Documents that support adaptivity on the object-level can integrate the same information in different form, they can allow different temporal arrangements of media objects or it can even be allowed to omit

			Mbuild	Firefly	CHIMP	Tiempo
Attribute-level	Media Objects	Flexible Durations		■	■	■
		Flexible Speed				■
		Quality Measures		■		■
		Selection Policies				■
	Temporal Relations	Flexible Delays	■	■	■	■
		Quality Measures				■
		Selection Policies				■
Object-level	Media Objects	Pres. Alternatives				■
		Omitting Objects			■	■
		Quality Measures			■	■
		Selection Policies				■
	Temporal Relations	Pres. Alternatives			■	■
		Quality Measures				■
		Selection Policies				■

Table 1. Comparison of document models

less relevant information. Adaptivity on the attribute-level enables the specification of variable durations, playout speeds or delays. Besides abstractions to model adaptivity, flexible document models should provide mechanisms to control the usage of presentation alternatives when adaptations are performed. This means, it must be possible to define the quality of presentation alternatives with regard to the overall presentation of the document. Further, flexible document models should offer means to control the choosing and switching between presentation alternatives.

The Tiempo model offers a rich set of abstractions to specify adaptive multimedia documents. It allows for a high degree of adaptivity on both the object- and attribute-levels. Further, it offers various means to control the adaptation of documents. In sum, Tiempo allows to create multimedia documents that can be adapted to a wide range of system configurations and load conditions.

Currently, an authoring system is being developed for the Tiempo model. The authoring system will allow a mainly graphical specification of adaptive multimedia documents. Further, a presentation system for environments with the best-effort service model is developed that is capable to adapt presentations to changing resource situations.

References

COMPUTER INC. APPLE, *QuickTime Developer's Guide*. Developer technical Publications, 1991.

G. BLAKOWSKI, J. HUEBEL, U. LANGREHR, Tools for Specifying and Executing Synchronised Multimedia Presentations. Presented at the *Proceedings of the International Workshop on Network and Operating System Support for Digital Audio and Video*, (1991) Heidelberg, Germany.

M. C. BUCHANAN, P. T. ZELLWEGER, Scheduling Multimedia Documents Using Temporal Constraints. Presented at the *Proceedings of the International Workshop on Network and Operating System Support for Digital Audio and Video*, (1992) San Diego, CA.

M. C. BUCHANAN, P. T. ZELLWEGER, Automatic Temporal Layout Mechanisms. Presented at the *Proceedings of the ACM International Conference on Multimedia*, (1993) Anaheim, CA.

K. S. CANDAN, B. PRABHAKARAN, V. S. SUBRAHMANIA, CHIMP: A Framework for Supporting Distributed Multimedia Document Authoring and Presentation. Presented at the *Proceedings of the ACM International Conference on Multimedia*, (1996) Boston, Massachusetts.

J. E. COOLAHAN, N. ROUSSOPOULOS, Timing Requirements for Time-Driven Systems Using Augmented Petri Nets. *IEEE Transactions and Software Engineering*, 9 (1983), 603–616.

G. D. DRAPEAU, Synchronization in the MAEstro Multimedia Authoring Environment. Presented at the *Proceedings of the ACM International Conference on Multimedia*, (1993) Anaheim, CA.

J. GECSEI, Adaptation in Distributed Multimedia Systems. *IEEE Multimedia*, April–June 1997, 58–66.

S. GIBBS, Composite Multimedia and Active Objects. Presented at the *Proceedings of the Conference on Object-Oriented Programming Systems, Languages and Applications*, (1991) Phoenix, Arizona.

R. HAMAKAWA, H. SAKAGAMI, J. REKIMOTO, Audio and Video Extensions to Graphical User Interface Toolkits. Presented at the *Proceedings of the International Workshop on Network and Operating System Support for Digital Audio and Video*, (1992) San Diego, CA.

R. HAMAKAWA, H. SAKAGAMI, J. REKIMOTO, Mbuild — Multimedia Data Builder with Box and Glue. Presented at the *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, (1994) Boston, Massachusetts.

HYTIME *Information technology — Hypermedia/Time-based Structuring Language (HyTime)*. ISO/IEC IS 10744, 1992

D. E. KNUTH, *The TEXbook*. Addison–Wesley Publishing Company, 1986.

T. D. C. LITTLE, A. GHAFOR, Synchronization and Storage Models for Multimedia Objects. *IEEE Journal on Selected Areas in Communication*, 8 (1990), 413–427.

MHEG *Information Technology — Coded Representation of Multimedia and Hypermedia Information Objects*. ISO/IEC IS 13522, 1995.

M. J. PEREZ–LUQUE, T. D. C. LITTLE, A Temporal Reference Framework for Multimedia Synchronization. *IEEE Journal on Selected Areas in Communication*, 14 (1996), 36–51.

B. PRABHAKARAN, S. V. RAGHAVAN, Synchronization Models for Multimedia Presentation with User Participation. Presented at the *Proceedings of the ACM International Conference on Multimedia*, (1993) Anaheim, CA.

N. U. QAZI, M. WOO, A. GHAFOOR, A Synchronization and Communication Model for Distributed Multimedia Objects. Presented at the *Proceedings of the ACM International Conference on Multimedia*, (1993) Anaheim, CA.

J. SCHNEPF, J. A. KONSTAN, D. }-C. DU, Doing FLIPS: FLexible Interactive Presentation Synchronization. *IEEE Jounal on Selected Areas in Communications*, 14 (1996), 114–125.

P. SENAC, M. DIAZ, A. LEGER, P. DE SAQUI–SANNES, Modeling Logical and Temporal Synchronization in Hypermedia Systems. *IEEE Journal on Selected Areas in Communications*, 14 (1996), 84–104.

R. STEINMETZ, Synchronization properties in multimedia systems. *IEEE Journal on Selected Areas in Communication*, 8 (1990), 401–412.

T. WAHL, K. ROTHERMEL, Representing Time in Multimedia Systems. Presented at the *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, (1994) Boston, Massachusetts.

T. WAHL, S. WIRAG, K. ROTHERMEL, TIEMPO: Temporal Modeling and Authoring of Interactive Multimedia. Presented at the *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, (1995) Washington DC.

S. WIRAG, K. ROTHERMEL, T. WAHL, Modeling Interaction with HyTime. Presented at the *Proceedings of the GI/ITG-Fachtagung Kommunikation in Verteilten Systemen*, (1995) Chemnitz, Germany.

S. WIRAG, Modeling of Adaptable Multimedia Documents. Presented at the *Proceedings of the International Workshop on Interactive Distributed Multimedia Systems and Telecommunication Services*, (1997) Darmstadt, Germany.

S. WIRAG, Adaptive Scheduling of Multimedia Documents. Technical Report 1997/12, University of Stuttgart, Germany, 1997.

Received: March, 1998
Accepted: May, 1998

Contact address:
Stefan Wirag, Kurt Rothermel
University of Stuttgart
Institute of Parallel and Distributed High-Performance Systems (IPVR)
Breitwiesenstr. 20–22
D-70565 Stuttgart
Germany
e-mail: {wirag, rotherme}@informatik.uni-stuttgart.de

KURT ROTHERMEL received his doctoral degree in Computer Science from Stuttgart University in 1985. After spending a sabbatical at the IBM Almanden Research Center he joined IBM's European Networking Center in 1988. He left IBM in 1990 to become a Professor of Computer Science back at Stuttgart University, where he now leads the Distributed Systems Research Group. His current interests are communication architectures and protocols, distributed multimedia systems, management of distributed systems, and mobile software agents.

STEFAN WIRAG studied Computer Science at the University of Stuttgart. He received his M.Sc. (Diplom-Informatiker) degree in 1994. He works now as a Research Assistent at the Distributed Systems Research Group of the University of Stuttgart, Institute of Parallel and Distributed High-Performance Systems (IPVR). His current research interests are multimedia documents, especially scheduling algorithms and QoS handling.