

Adaptive Hypermedia for Personalized TV

Judith Masthoff and Lyn Pemberton
Judith.Masthoff{Lyn.Pemberton}@brighton.ac.uk

University of Brighton, School of CMIS
Brighton, UK
Tel: +44 1273 642490, Fax: +44 1273 642405

Keywords: interactive television, personalized television, adaptive hypermedia, streaming media

Abstract

Adaptive hypermedia techniques have until now been mainly applied to web pages with text and images. In this chapter, we explore how these techniques can be extended to time-based media, particularly in the domain of interactive television.

1 Introduction

The advent of digital television has brought new opportunities for viewers to interact with their televisions by "allowing viewers to access additional information about the programme they are watching, select different camera views of a sports event, play games, access enhanced teletext services, or access the internet and send e-mails" (BBC, 2003). Interactivity is set to expand, as programme makers take advantage of its creative capabilities, broadcasters recognize it as a mechanism for creating new funding streams and governments seek to expand their channels of communication with citizens. More households are adopting digital television: in the UK alone, digital television is expected to be in use by 15 million subscriber homes by 2005 (UK Government, 2001). Interactive television can be seen as the extension of interactivity with computers into the broadcast television world. Digital interactivity, of course, opens the way for adaptivity, which has the potential to benefit a range of viewer types. However, it seems likely that simply applying the adaptivity techniques familiar from the desktop domain to iDTV will not work, for a variety of reasons. These include the nature of the iDTV audience, their relationship with the television medium, technical issues of speed and bandwidth and more pragmatic concerns about cost. In this chapter we explore the potential for adaptivity offered by iDTV, showing how it might be operationalized in and around the broadcast television stream in a range of genres.

Unlike the personal computer, which was a new piece of technology bringing little "baggage" in terms of expected behaviours, the television is already rich in cultural resonances, which we need to understand if we are successfully to make use of the medium in new ways. The television audience, their attitudes to the television and their patterns of television use all differ from their PC equivalents and offer challenges for adaptivity. Televisions are found in almost all homes in the Western world, ensuring a reach beyond that of the personal computer. The potential user population for adaptive

iDTV is therefore rather different from that for PC-based applications. While this difference in user population opens up the possibility of delivering the benefits of personalization to a wider audience, it does suggest that new approaches may be needed.

For many households, the television set is the literal and metaphorical focus for family life, with a rich set of accompanying behaviours. A number of studies (Ling & Thrane, 2002; Lull, 1988; Moores, 1996) have described the television-watching behaviour of families, and a general picture has emerged of people watching programmes in company, watching while engaged in other tasks, using television as ambient atmosphere provision, channel hopping and generally seeing television as oriented to leisure rather than task performance. This is often characterized as a "leaning back" attitude to media consumption, as opposed to the active, engaged model of "leaning forward" to interact with a PC. Ling and Thrane, in their interview study of Norwegian viewers, suggest that television has "a limited potential to engage the individual", finding its niche as a means of relaxation or even as background noise, used for generating ambience rather than for paying attention to specific content. One of their participants goes so far as to say: "I don't watch TV to like, learn anything". This is also reflected in research into attitudes to digital television amongst UK viewers (Ofcom, 2001), suggesting it may be dangerous to assume a willing mass audience for formal "work-like" provision via television.

In contrast to the use of PCs, television viewing is largely a family or social activity (Barwise & Ehrenberg, 1988; Kasari & Nurmi, 1992). According to a large research study in the UK (Livingstone & Bovill, 1999), television is the medium most often shared with family. This tendency towards viewing television in company suggests that individual-based adaptation familiar from PC-based research might not apply.

A final major difference between use contexts for PC and television arises from the broadcast nature of television as opposed to the time-independent nature of most computer-based material. In countries with a limited number of major television channels, this gives a sense of sharing an experience with others beyond the viewer's own living room, which is an attractive feature for many. Although many viewers may be willing to move to a time switching method of interaction, e.g. via a digital video recorder or recordable DVD, others may not want to lose the sociability of the national broadcast. Attempting to integrate adaptivity into the broadcast stream will be a challenge.

This chapter will not discuss in any detail how the television can decide what is best for a certain viewer, but will focus instead on what adaptation possibilities there might be and how the adaptation could be executed using adaptive hypermedia techniques. The main research questions addressed are:

- Where and how can the adaptive hypermedia techniques discussed by Brusilovsky (2001) be applied in the iTV domain?
- How can these techniques be extended in order to deal with time-based media? What are the main problems to overcome?

In Section 2, we present a three part model of interactivity in the iTV domain, and show briefly how adaptive hypermedia techniques can be applied to so-called distribution and extra programme interactivity. In Section 3, we focus on so-called intra-programme interactivity, and show how adaptive hypermedia techniques can be interpreted for time-based media, what the main problems are and how to overcome them. Technical feasibility and evaluation issues are explored in the conclusion.

2 A Three Part Model of Adaptivity for iTV

Definitions of interactive television are still in flux. We will adopt here a modified version of the Henley Centre framework described in (Gawlinski, 2003, p. 6). The model distinguishes three types of interactivity, according to whether the interactivity applies to the entire programme, a non-broadcast activity or the programme content. These are:

- Distribution interactivity
- Extra-programme interactivity
- Intra-programme interactivity

2.1 Distribution interactivity

Distribution interactivity has as its object the entire programme rather than the programme content. For instance, the Electronic Programme Guide (EPG) allows viewers to scroll through an interactive schedule, search for programmes in particular genres and possibly set reminders. Some EPGs also include cast listings, film reviews and multimedia content. Certain EPGs already include a touch of personalisation as an interface feature: the channels that are being broadcast are greyed out if the household has not signed up for the appropriate subscription. In this way the viewer won't be misled into making a wrong choice, but can still see what they are missing. A more lightweight version of distribution interactivity is the "now and next" box widely used to allow viewers a brief description of the current and following programme on a specific channel. The electronic programme guide is considered a "must have" for multi-channel viewing where hundreds of channels may be available (Gawlinski, 2003). Viewers interviewed as part of the GO Digital project considered the EPG to be the most appreciated interactive service (Go Digital Group, 2003).

EPG's have been the favoured testing ground for much of the research that has been carried out to date into adaptation for interactive TV (e.g. Ardissono & Buczak, 2002). Some adaptive EPGs have been developed to market, for instance Personal TV, originally developed at University College Dublin (Cotter & Smyth, 2002; Bonnici, 2003). Adaptive EPGs filter the enormous number of programmes on each day, to recommend viewing material. They base their recommendation on their prediction of the viewer's likes and dislikes, using two approaches. The first of these, content-based filtering, selects material based on how similar it is to material appreciated by the viewer in the past, as recorded in a model of past viewing choices. The second technique, social filtering, takes into account the opinions of like-minded users, making selections based on how much the material has been appreciated by other viewers who resemble our viewer. The advent of Personal Video Recorders and Video on Demand makes adaptive

EPGs even more vital: instead of just choosing from programmes currently being shown, the viewer may be able to view broadcasts covering months or years.

Currently, EPGs focus on recommending programmes the viewer is likely to enjoy, typically based on genre preferences (Uchyigit & Clark, 2003). It would be possible to operate with a variety of different criteria. For instance, EPGs could draw attention to programmes that contain content of local interest, that are broadcast with accessibility support of a particular kind or whose content is relevant to a particular stage in a curriculum. Such a system would need a rich model of the viewer, indicating, for these examples, their location, their abilities and their learning interests, together with metadata about TV programmes, indicating the topics they address, the media used and so on.

Current adaptive EPGs supply recommendations to individual viewers. However, given that TV is often watched in company, it would be natural to extend these recommendations to *groups* of viewers, taking into account what is appropriate for the individuals in the group. Viewing groups can be quite heterogeneous, and age, gender, life experience, intelligence, and personality influence what types of TV programmes people enjoy (Livingstone & Bovill, 1999). In our previous work (Masthoff, 2004), we have discussed how one can adapt to a group of viewers, in such a way that each individual enjoys or benefits from the broadcast.

2.2 Extra-programme interactivity

In extra-programme interactivity the focus is neither the programme, nor the content but some other activity available via the television set. A major category here is made up of Internet-style hypertext browsing facilities such as Internet on television, teletext-like pages and the "Walled Gardens" of information provided by iDTV platform operators. Typically these would include local events, news, sports, cinema listings and so on, as well as specifically television-related information. Another major category is games: these are currently the most popular of all services provide via iDTV. Chat and email, which would also fall into this category, are already available on interactive television. For example, the NTN network enables US customers in bars and restaurants to play interactive games via iTV and chat with each other as they do so. Quico (2003) reports large dinner parties being organised in Portugal for groups of people who meet via interactive TV messaging facilities.

Adaptivity could be used to select appropriate other viewers to communicate with. Existing chat and email applications could be adapted fairly straightforwardly for this, in a range of domains. For instance, viewers of learning programmes might be put in contact with other learners, who may be more proficient and able to help them, less proficient and in need of the kind of help that would challenge the learner to verbalise their knowledge, or are struggling with the same problems. The I-Help system (Greer et al, 2001) is an example of how learners may adaptively be brought in contact with each other. In the sorts of informal learning that seem most natural to the television context, formal exercises will not be appropriate and bringing learners into contact with each other (and teachers) seems therefore particularly important. Consumers could share their opinions about products advertised or sold via television, allowing others to make a well-

informed choice between alternative products. In domains such as sport, entertainment and news, viewers could share opinions and could chat to others who share (or dispute) their tastes.

2.3 Intra-programme interactivity

Services in this category allow the viewer to interact with content of the broadcast stream, to create what is known as "enhanced television". This has the great advantage of not requiring the viewer completely to abandon watching the broadcast stream while interacting with the programme. There are two main subcategories.

Information Interactivity

The viewer can access information relating directly to the programme (or advertisement) content. Documentaries are often enhanced in this way. The standard in the UK for this type of service was set by the BBC's natural history edutainment programmes, such as *Walking with Beasts* (Gawlinski, 2003), where viewers could elect to watch additional textual information and video footage, and could choose between commentaries with different levels of scientific detail. Another high-profile example was the ITV coverage of the 2001 General Election in the UK, where viewers could navigate via an interactive map to watch real time coverage of local events in their own constituency, while still keeping up with the overall picture. Typically this information is set out in pages of hypertext, with the programme displayed in reduced size in a corner of the screen.

Participation Interactivity

These services allow the viewer to choose between options during a programme or advert, for instance in order to play along with contestants in a game show, to vote for a contestant in a talent show, to buy a product or to select which player to concentrate on during a football match. For instance, programmes for toddlers are enhanced by simple games, practicing colours and shape recognition. These interactive opportunities can be closely synchronised with the broadcast stream. For instance, one service asks viewers to predict the outcome of penalty shots in football matches, giving the viewer only a few seconds to make the choice.

In the remainder of this chapter, we concentrate on the potential for intra-programme adaptivity, within the broadcast television stream itself.

3 Adaptivity in the broadcast television stream

3.1 Selecting streams and items

One adaptive hypermedia technique discussed by Brusilovsky (2001) is the selective showing and hiding of content elements and links. In traditional hypermedia, content elements are typically static text and images.

3.1.1 Interpretation of this technique in the iTV domain

Information in the television context is presented in a complex, multimedia structure, and the entities or objects potentially subject to adaptivity are multiple. A television

programme can be conceptualised as consisting of multiple “streams” presented simultaneously to the viewer. For instance, a programme might include overlaid video streams, multiple audio streams (such as background noise, music, speech, and the audio description for visually impaired viewers), and textual and graphical overlay streams (such as closed captioning, broadcaster logo, text commentaries, some of which may be animated). For an example see Figure 1. The main difference from the content elements in traditional hypermedia is that streams are transient and time-based: what the viewer is seeing and hearing has a beginning and an end, changing over time.

Each stream may be either presented or suppressed. For instance, an audio description stream could be presented and a closed caption stream suppressed for a viewer who is visually impaired. An overlay stream with the latest stock market figures could be presented or suppressed depending on the viewer’s suspected interests and so on. These choices could be made on the basis of an appropriate user model.

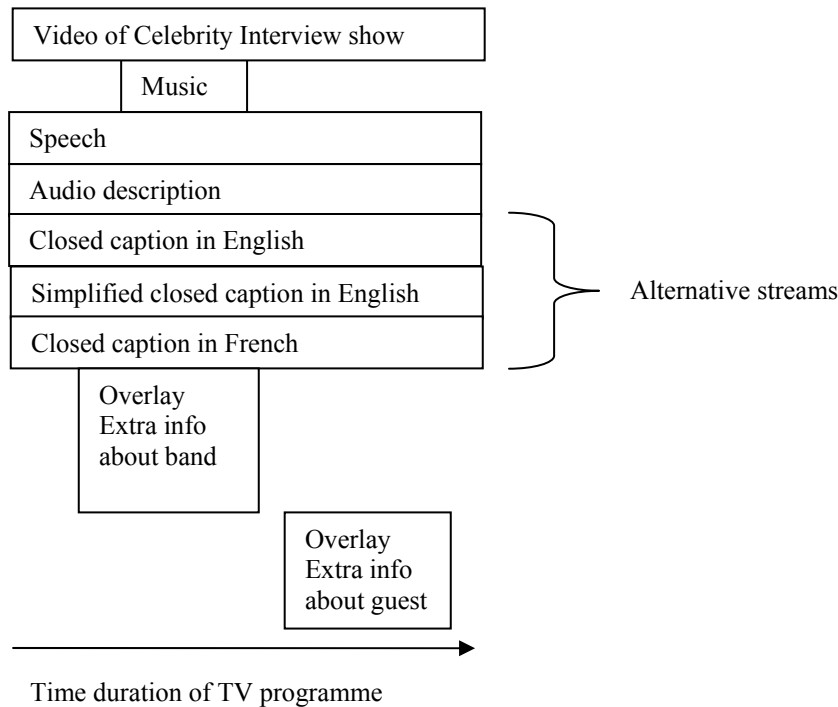


Figure 1. Example of a television programme consisting of multiple streams.

Streams may be alternatives for each other. For instance, two alternative video streams might be designed to fill the same space on the screen, or two audio tracks might give alternative commentaries for a programme (as in “Walking with Beasts” discussed above). In this case it will not be possible to present them simultaneously, and the television might choose which one to present. For instance, an overlay stream with the option to buy a Manchester United shirt could be presented as part of a football match broadcast if the viewer is a Manchester United fan, in preference to a similar overlay for an Arsenal shirt.

A television programme can also be conceptualised as consisting of multiple sequential so-called “items”. In some genres, such as dramas, these items are tightly interwoven via skilled editing. In other genres though, the distinctions between items are clear-cut. News programmes, chat shows, holiday programmes, music shows, and advertisement breaks all typically consist of multiple items. An example of a TV programme with multiple items is shown in Figure 2. Note that each item may be composed of a number of different kinds of media stream and that items can have different durations. The duration of an item will tend to be determined by the duration of the video stream.




Figure 2. Example of a programme consisting of a number of items, with each item consisting of a number of streams.

In a personalised system, each item could be presented or suppressed. For instance, instead of showing everybody the same news broadcast, news items could be included or not depending on their relevance or interest to the viewer. Choice of news items tends to incorporate viewer interests, quite often partly derived from where the viewer lives (Dimitrova et al, in press; Maybury et al, 2004).

3.1.2 Problem: how to deal with the fixed time duration of programmes

Given the time-based nature of television, suppressing an item will result in a television programme becoming shorter. Hence, suppressing items is only easily feasible if there are more items than needed to fill the time slot of the programme. One way to avoid the time duration problem would be to allow the set-up box to choose only between alternative items with the same duration (such as news items C and D in the example above). This would however make it impossible to choose between, say, a short superficial treatment of a news item and an in-depth analysis which is likely to take more time. One solution would be to define the in-depth analysis as a separate news item that could follow the superficial news item and be shown in preference to another news item of the same duration. For instance, after having seen news items A and B, viewers could be presented with either a continuation of B (namely C) or a news item on a different topic (D).

3.1.3 Problem: How to deal with commercial and social pressures

There are commercial and social issues that limit the freedom with which items may be suppressed. Companies paying for advertisements to be broadcast will be likely to have strong views about their suppression. Suppressing certain news items may influence people’s political and social views, and may therefore be deemed unacceptable. To avoid

this problem, items for selection might be alternatives to each other, perhaps presenting the same information but from a different angle. For instance, there might be a children's version of a news item as an alternative to an adult one. Alternative advertisements for different models of the same brand might show either the sporty hatchback or family model of the same car brand, according to a viewer model. Another solution would be to change the costing model of advertisements to reflect the percentage of viewers to which it is broadcast.

3.2 "Stretch media": extending streams and items

A second adaptive hypermedia technique is so-called "stretch text," i.e. presenting selected text fragments in condensed form with the option for the user to access more details (Brusilovsky, 2001).

3.2.1 Interpretation of this technique in the iTV domain

Providing viewers with access to additional information is currently a popular way of making television programmes interactive. This often takes the form of a "call to action" icon on the screen which is activated via a key on the remote control to display a pop-up box containing supplementary information. In iTV programmes, information is available in a variety of media, and therefore we will talk about "stretch media" instead of "stretch text".

Audience research has shown that textbox pop-ups are appreciated and acceptable 10 to 15 times per show. By contrast, some episodes of *Walking with Beasts* had as many as 80 pop-ups (Nolan, 2003). This illustrates two points: firstly, that TV programme makers may be willing to provide large amounts of additional material for their programmes, and secondly, that there is an opportunity and need for adaptive selection of the additional material to present. Currently, it is up to the viewers to choose what they want to see and hear, so, in a sense, everything is presented as stretch media. However, people tend to watch TV in leisure- rather than work- mode, suggesting that the amount of interaction required should be kept to a minimum. In these circumstances, an effective approach to adaptivity would be to select default settings appropriate to the viewer in order to present some material automatically and other material as stretch media, while completely suppressing inappropriate material. Content-based and social filtering techniques can be used here, with the viewer model initiated by using stereotypes based on market research, by viewer tracking or by manual parameter-setting by the viewer themselves. Note that additional material can be incorporated into a TV programme as additional streams (such as the textual overlays and alternative commentaries in *Walking with Beasts*), as additional items (such as additional video footage about the making of a movie), or as a time extension of items.




Figure 3. Example of stretch media for a textual overlay stream.

Figure 3 shows an example of stretch media applied to a textual overlay stream. Pressing number 1 on the remote could present a textual overlay as shown in the middle screen. If the user model indicates that the user has a strong interest in painting, the overlay could have been presented automatically. A combination of stretch media and filtering could also be used: when the viewer presses 1, the system could decide, based on the user-model, which overlay stream to present from a set of possible overlays. For instance, an alternative could be shown to somebody more interested in gardening, discussing when to plant this tree, and how to care for it (see right-hand screen). Note that this information does not necessarily have to be related to the original programme topic: for instance, any programme could be changed into a programme that teaches children about colours, by asking them to tell us the colour of an indicated area in the video stream.

Stretch media can also be used for items. For instance, a news item could be introduced in concise form, and the interested viewer could ask for a more in-depth presentation. The condensed version of an item could be a short video summary with an indication to press a certain button to see more. Instead of a video summary, a (scrolling) text overlay at the bottom of the screen could be used, or a short auditory cue. The problem of over-running identified in 3.1.2 would also need to be tackled.

3.2.2 Problem: How to deal with the time-based nature of streams

Presenting a textual overlay stream (as in the example above) when the viewer presses a button could result in the text not being visible long enough to allow the viewer to read it (see ITC, 1999). Similarly, the viewer might miss the start of an auditory overlay. Additionally, some indication for the viewer is needed about what media can be stretched and how to do it, i.e. an anchor of some sort. To solve these problems, for each stream to be used as stretch media, a partner stream is needed that represents its condensed version, indicating *when* the viewer can perform *which action* to get more information about *what*. The “when” is indicated by the duration of the partner stream. The “which action” is indicated by the symbol or sound presented. In the example above, a “1” indicates that the action is pressing the 1 on the remote control. The “what” can be indicated by the location of a symbol, the timing of audio, or the words used. In the example, the location of the “1” indicates that more information is available about the plant.

3.2.3 Problem: How to deal with multiple simultaneous stretching opportunities

Another problem that arises from the time-based nature of streams is the difficulty of dealing with multiple stretching options simultaneously. For instance, one stream might provide additional information about a talk show guest, while a second stream might tell the viewer about the clothes she is wearing. The viewer might be interested in both, and the system could decide to present both as stretching opportunities (for instance, showing “1-guest 2-clothes” at the bottom of the screen). However, when the viewer has seen the information about the guest, time might have run out, preventing her from seeing the information about the clothes. A number of possible solutions could be explored. Firstly, when using a PVR, it might be possible to freeze the programme while viewing additional information, so that the programme continues where the viewer left off. This would work if the duration of the programme could be adjusted accordingly by removing or compressing a later item. Secondly, bookmarks to additional information could be stored for viewing in the non-TV area after the programme has finished. Thirdly, the additional information could be presented on another device, for instance via a text message to a mobile phone. Lastly, multiple stretching opportunities might be presented only in those situations where there was sufficient time to view all of them, or where it was made very explicit that the viewer would be able to select only one to be presented.

3.2.4 Problem: How to obtain stretch material

Producing broadcast material is costly and time intensive, so the question arises as to who would produce the material and what business model could be used. The need for business models was one of the main discussion points at the latest Personalized Television workshop (at User Modelling conference, Pittsburgh, 2003). A standardised architecture would allow material to be provided not only by programme makers but also by specialised companies (perhaps in return for advertisement space: “these annotations were brought to you by”) or dedicated individuals approved by the broadcaster. For instance, in the example above, a language-teaching company could provide annotations on French vocabulary, e.g. “green = vert”. A garden centre could add information on how expensive the tree is and where to buy it. Viewers could sign up for annotations from specific sources (and pay for the service), or the sources could be filtered automatically, using content-based and social filtering. Thomas (2002) discusses the possibility of the general public adding annotations (suitably moderated) and the usability issues involved in doing this effectively.

Research is ongoing on how to automate use of the vast resource of existing broadcast material. For instance, a question in a quiz show about a novel might result in a fragment being shown of a TV dramatisation of that novel. Content-based filtering could be used for this, but it does require marking up of content. Fortunately many programmes will have closed captions, which can be used for text-based filtering. In addition, there has been work on video summarisation, including personalised video summarisation, where summarised videos are produced based on key words the viewer is interested in and the length of time available (Tseng & Lin, 2002). Dimitrova et al (in press) present algorithms for automatically splitting up a news programme into individual news items, and producing a textual summary and image to describe each news item. Automatic summarization can also be very useful for generating partner streams for stretch media.

In addition to using existing broadcast material to augment a programme, information from the Internet can be used. For instance, InfoSip automatically augments film broadcasts with the latest information about actors from the internet, and MyInfo incorporates the latest on-line local weather and traffic information into a news broadcast (Dimitrova et al, 2004).

3.3 Dimming or highlighting streams and items

A third adaptive hypermedia technique is the dimming or highlighting of selected content elements and links (Brusilovksy, 2001). In text-based adaptive hypermedia, this is done by modifying font size, font face, font colour, and contrast between background and font colour.

3.3.1 Interpretation of this technique in the iTV domain

Instead of suppressing streams and items, or allowing the viewer to access them at will, streams and items could be presented, but be made less or more obvious. Dimming or highlighting an item can be achieved in the following ways:

- Reducing or increasing the volume of all auditory streams. For instance, the volume could reflect the relevance of a particular advertisement for the viewer. It may also be possible to dim an item by reducing the brightness and contrast of all visual streams. In traditional adaptive hypermedia dimming is used to allow users to concentrate on key paragraphs and images while still being able to peripherally see the dimmed ones (Brusilovsky, 2001). In the kind of dimming discussed here, the viewer would not have anything else on the TV screen to concentrate on (unless multiple items are shown on the same screen side by side, see below). However, as we discussed above, viewers tend to multitask while watching TV. Reducing the volume would allow them to concentrate on the other things they are doing, such like reading or talking.
- Increasing or reducing the speed of all streams. Increasing the speed will save time, so is one way to fit more content into the fixed duration of a television programme. Reducing the speed could, for instance, allow viewers more time to understand the language used and the points made.
- Presenting the item simultaneously with another item, while using a smaller screen space. Examples are shown in Figure 4. In a sense both items are dimmed, but particularly the item allotted the smallest space. The left-hand example in Figure 4 shows how item B (for instance a football match) can be made proportionally smaller and overlaid on item A (for instance an episode of a soap). The viewer could be watching the soap while keeping an eye on the football. If a goal were scored, the size of the football item could automatically become larger. The middle example shows how item C (for instance a news flash) can be overlaid in its original size on item A (episode of a soap), when for instance only the closed caption stream of item C is shown and no closed captions are shown for item A. Note that while combining visual streams from multiple items in these ways might work quite well, it will be

harder to combine auditory streams. In case of conflicting auditory streams, the volume of the most dimmed item should be put to zero (or very low).




Figure 4. Examples of dimming an item by showing it in a smaller screen portion simultaneously with another item.

Similar ways can be used to dim or highlight individual streams: for instance, modifying the volume of an auditory stream, its speed, or the relative spatial size of visual streams. Textual streams can also be dimmed or highlighted using the font modifications discussed by Brusilovsky (2001).

3.3.2 Problem: How to deal with limited space

Only a limited amount of text can appear on a television screen due to its resolution. This poses problems when highlighting textual streams. Manual scrolling as used on web pages is not an option for television. Automatic scrolling, in which text moves (as in the textual new flash overlays on CNN) impedes readability (Miller, 2003). Another approach, as used in Teletext on analogue UK television, is to split textual overlays into multiple pages, automatically cycling through these pages. It seems most likely that alternative streams will be produced by the broadcaster team for a textual overlay with different font sizes.

3.3.3 Problem: How to deal with synchronisation

Changing the speed of an auditory stream might lead to a lack of synchronisation with the visual stream. A similar problem holds for changing the speed of textual streams. One, rather restrictive, solution would be to only allow changing the speed of streams that do not need to be synchronised with other streams. A more sophisticated solution would require metadata indicating to what extent the speed of streams can be modified without losing synchronisation. The easiest way to deal with synchronisation would be to have alternative, carefully edited, streams available for the various speeds.

3.3.4 Problem: How to avoid interference

Highlighting a stream can interfere with other streams, not just it attracts the viewer's attention, but because, for instance, a larger text overlay could obscure more of the video image or louder speech could obscure the music. Either metadata is needed that indicates the boundary of parameter settings (like the maximum size of a text overlay), or, as before, carefully edited alternative streams can be provided with for instance text overlays in different sizes.

3.4 Sorting items

A fourth adaptive hypermedia technique is the sorting of content elements and links: changing the order in which paragraphs, images, or links are shown on a page, so changing *spatial locations* (Brusilovsky, 2001).

3.4.1 Interpretation of this technique in the iTV domain

As our items are time-based, sorting mainly means changing the *timing* of the presentation of items. For instance, in a news broadcast, the order in which news items are presented could be modified. Viewers might be presented first with the items that are most important or appealing on the basis of their user model. One commercially motivated reason for doing this is that viewers have a tendency to switch channels if they become bored with a programme. Dimitrova et al (in press) discuss the ordering of news items based on the position of the items in the original news broadcasts, the timing of events depicted in the items, and the match between the items and the user profile. Masthoff (2004) discusses how an ideal ordering of items also depends on the mood induced by items, the topical relatedness of items, and the importance of items.

Sorting in the spatial sense can be applied when multiple items are shown simultaneously on the same screen, or multiple stretch indicators are located at the bottom of the stream. Note that the latter requires the locations to be changeable by the television rather than being fixed in the stretch partner streams.

3.4.2 Problem: How to maintain continuity

Many of the techniques suggested so far, involving automatic removal, addition, modification and reordering of elements will be very problematic in that such approaches challenge current practices in television production, which have been honed over approximately sixty years and are underpinned by a wealth of specialist skills, particularly editing for continuity, an important principle for time based media. Continuity is established for instance by overlapping background sounds or music from one scene to the next or by using visual transitions between shots. The presenter in a news programme (and their text) also provides continuity. When suppressing items, choosing between alternative items, or sorting items, it is important to generate continuity. The easiest way to do this would be by careful editing of all alternatives, considering all possible orders of presentation. As this is expensive, it would be likely to restrict the opportunities for adaptation. Alternatively, connecting sentences could be generated for a textual overlay (or audio) that overlaps between two consecutive items, or embodied animated agents could be used in a presenter role. Magnenat Thalmann and Thalmann (1995) discuss automatically generated digital actors for interactive TV. Simple avatars have already been used in live experiments with so-called inhabited television, to represent participating viewers on the screen (Benford et al, 2002). More sophisticated animated agents with generated text, speech and gestures have been developed, for instance in the European NECA project (Krenn et al, 2002).

4 Conclusions

We have shown how techniques from the Adaptive Hypermedia community can be used and will need to be extended to deal with interactive television, in particular with its time-based nature. Though having been inspired by the interactive television domain, we expect that many techniques presented here are more broadly applicable to other uses of time-based media, such as video and radio on-line and on mobile phones.

The techniques described in this chapter seem technologically feasible in the UK. Already, many so-called “video switching” iTV applications exist that allow viewers to select from multiple items, such as Sky News Active (choosing between main news, weather, entertainment, etc), BBC’s Wimbledon (choosing which court to watch), Sky Sports Active (choosing between camera angles) and Channel Four’s Big Brother (choosing between rooms) (Gawlinsky, 2003). According to Gawlinsky (*ibid*), ACTV in the US have also pioneered a variant of video switching where viewers make a number of choices during a programme, resulting in seamless switches of video, giving the viewers the impression of continuing to watch the same programme. Existing applications tend to use multiple channels: choosing between video items is basically implemented as switching between channels. This would make it expensive to select between a large number of items, and current applications typically allow selections between at most four items. Already, some iTV applications allow viewers to select between streams, such as Walking with Beasts (choosing between alternative commentaries). Already, many iTV applications make extensive use of stretch media. The technology is therefore available for the setup-box to make selection decisions on the user’s behalf, if it can build up an appropriate user model. Much research has already been done on building up user models of iTV viewers both for news and advertisement selection (see e.g., Dimitrova et al, *in press*; Maybury et al, 2004; Lekakos & Giaglis, 2004). As in the PC domain, social, and content-based filtering can be used, combined with stereotypes (Ardissono & Buczak, 2002). It has also been shown how this can be extended to cope with multiple viewers (Masthoff, 2004). There have been several proposals to solve the problem of how the TV will know who is watching it (needed if adapting to individuals rather than a household), including infrared cards (Lieberman et al, 1999), personal remote controls, a conversational login procedure and a probabilistic mechanism using the time of day combined with a known probability of a viewer watching at that time (Goren-Bar and Glinansky, 2002). We have not limited our discussion to techniques that are already straightforwardly implementable. For instance, it is unlikely that in the near the future animated agents will be used as news presenters or that existing broadcast material will be reused automatically. However, the technology is developing rapidly, and discussions such as ours might even contribute to giving ideas in what directions to develop it.

The evaluation of adaptive interactive television requires special attention. The evaluation of computer-based adaptive systems is already non-trivial (Masthoff, 2002; Weibelzahl & Paramythis, 2003). For adaptive interactive television, both empirical and heuristic evaluation methods will need to be extended. They need to take account of the different nature of television viewing, such as viewing for relaxation, multi tasking, viewing as background activity, and viewing in groups (Pemberton & Griffiths, 2003).

5 References

- Ardissono, L. & Buczak, A. (Eds.) (2002). Personalisation in future TV. Proceedings of TV'02, the 2nd workshop on personalization in future TV. Malaga, May 2002.
- Barwise, P., & Ehrenberg, A. (1988). *Television and its audience*. London: Sage.
- BBC. (2003). New Media Digital Developments.
http://www.itc.org.uk/the_digital_age/DEVELOPments/interactive_tv.asp
Accessed 26 September 2003.
- Benford, S., Greenhalgh, C., Craven, M., Walker, G., Regan, T., Morphett, J. & Wyver, J. (2000). Inhabited television: Broadcasting interaction from within collaborative virtual environments. *ACM Transactions on Computer-Human Interaction*, 7, pp510-547.
- Bonnici, S. (2003). Which channel is that on? A design model for interactive programme guides. In J. Masthoff, R. Griffiths and L. Pemberton (Eds.) *Proceedings of the first European Interactive TV conference*, Brighton, UK, April2-4, 2003.
- Brusilovsky, P. (2001) Adaptive hypermedia. *Journal of User Modeling and User Adapted Interaction*, Ten Year Anniversary Issue (Alfred Kobsa, ed.) 11 (1/2), 87-110
- Cotter, B. & B. Smyth. (2002). Personalised Electronic TV Guides. *Artificial Intelligence Magazine*, 21 (2). AAAI Press, California.
- Dimitrova, N., Zimmerman, J., Janevski, A., Agnihotri, L., Haas, N., Bolle, R. & Velipasalar, S. (in press). Content augmentation aspects of personalized entertainment experience. To appear in L. Ardissono, M. Maybury and A. Kobsa (Eds.) *Personalised Television*, Kluwer Publishers.
- Gawlinski, M. (2003). *Interactive Television Production*. Focal Press: London.
- Go Digital Group. (2003). GO Digital Key Findings. Available at:
http://www.itc.org.uk/uploads/GO_DIGITAL_KEY_FINDINGS.pdf
- Goren-Bar, D. & O. Glinansky. (2002). Family stereotyping: A model to filter TV programs for multiple viewers. In L. Ardissono and A. Buczak (Eds.) Proceedings of the 2nd workshop on Personalization in Future TV, Malaga.
- Greer, McCalla, Vassileva, Deters, Bull, S. & Kettel, L. (2001). Lessons learned in deploying a multi-agent learning support system: The I-Help experience. In J.D. Moore, C. Luckhardt Redfield, and W.L. Johnson (Eds.) *Artificial Intelligence in Education: AI-ED in the wired and wireless future*. IOS Press: Amsterdam.
- Independent Television Commission (1999) General Guidance on Subtitle Display.
http://www.itc.org.uk/itc_publications/codes_guidance/standards_for_subtitling/s_ubtitling_1.asp
- Kasari, H., and Nurmi, S. (1992). TV audience segments based on viewing behaviour. In Advertising Research Foundation (ARF) and European Society for Opinion and Marketing Research (ESOMAR), Worldwide Broadcast Audience Research symposium. Toronto.

- Krenn B., Pirker H., Grice M., Baumann S., Piwek P., Deemter K. van, Schroeder M., Klesen M., Gstrein E. (2002). Generation of multimodal dialogue for net environments. In S. Busemann (ed.), *KONVENS 2002*, Deutsches Forschungszentrum fuer Kuenstliche Intelligenz (DFKI), Saarbruecken, Germany, pp. 91-98.
- Lekakos, G., & Giaglis, G.M. (2004). A lifestyle-based approach for delivering personalized advertisements in digital interactive television. *Journal of Computer Mediated Communication*, 9.
<http://www.ascusc.org/jcmc/vol9/issue2/giaglis.html>
- Lieberman, H., van Dyke, N., & Vivacqua, A. (1999) Let's browse: A collaborative web browsing agent. Proceedings of the 1999 International Conference on Intelligent User Interfaces.
- Ling, R. & Thrane, K. (2002). "I don't watch television to like learn anything": the Leisure Use of TV and the Internet. *First Monday*, 7, 1 (January 2002).
http://firstmonday.org/issues/issue7_1/ling/
- Livingstone, S. & Bovill, M. (1999). *Young people, new media. Summary report of the research project: Children, young people and the changing media environment*. As accessed on
http://www.lse.ac.uk/Depts/Media/people/slivingstone/young_people_report.pdf .
- Lull, J. (1988). *World Families Watch Television*. Sage: Newbury Park, Cal.
- Magenat Thalmann, N. & Thalmann, D. (1995). Digital actors for interactive television. Proc. IEEE, Special Issue on Digital Television, Part 2, July 1995, pp.1022-1031.
- Masthoff, J. (2002). The evaluation of adaptive systems. In N. V. Patel (Ed.), *Adaptive evolutionary information systems*. Idea Group publishing.
- Masthoff, J. (2004). Group modeling: Selecting a sequence of television items to suit a group of viewers. *Journal of User Modeling and User Adapted Interaction, Special issue on personalized TV*.
- Maybury, M., Greiff, W., Boykin, S., Ponte, J., McHenry, C. & L. Ferro. (2004) Personalcasting: Tailored broadcast news. *Journal of User Modeling and User Adapted Interaction, Special issue on personalized TV*.
- Miller, C. (2003). Interactive television services: Content guidelines. Royal National Institute of the Blind report, UK. Available on
<http://www.theforum.gov.uk/events/presentations/interactive-content-guidelines-rnib.pdf>. Accessed 3 October 2003.
- Moore, S. (1996). *Satellite Television and Everyday Life: Articulating Technology*. Acamedia Research Monographs 18. University of Luton Press.
- Nolan, S. (2003). The interrupting cow. Invited talk, First European Conference on interactive TV, Brighton, UK, April 2-4, 2003.
- Oftel (2001). Digital Television - Consumers' use and perceptions, Aug 2001 (published Sept 2001). Available at:
<http://www.oftel.gov.uk/publications/research/2001/digtv0901.pdf>

- Pemberton L. & R. N. Griffiths. (2003). Usability Evaluation Techniques for Interactive Television. *Proceedings of HCI International*, Crete
- Quico, C. (2003). Are communication services the killer application for interactive TV? Or: "I left my wife because I am in love with the TV set". In J. Masthoff, R. Griffiths and L. Pemberton (Eds.) *Proceedings of the first European Interactive TV conference*, Brighton, UK, April2-4, 2003.
- Thomas, N. (2002). A conceptual model for peer-to-peer interactive television. In J. Masthoff and R. Luckin (Eds.), *Proceedings of the workshop Future TV: Adaptive instruction in your living room*, pp. 10-14.
- Tseng, B.L. & Lin, C-Y. (2002). Personalized video summary using visual semantic annotations and automatic speech transcriptions. *IBM research report RC22610*, October 29.
- Uchyigit, G. & K. Clark. 2003. Personalised multi-modal electronic program guide. In J. Masthoff, R. Griffiths and L. Pemberton (Eds.) *Proceedings of the first European Interactive TV conference*, Brighton, UK, April2-4, 2003.
- UK Government. (2001). Short Survey of Published Material on Key UK Trends: 2001-2011.
- Weibelzahl, S., & Paramythis, A. (Eds.). (2003). Proceedings of the workshop Empirical Evaluation of Adaptive Systems, associated with the User Modeling conference, Pittsburgh, 2003.