

CSRP 477: Architectural Requirements of a System Retrieving Adaptive Image Objects

Malcolm McIlhagga, Ian Wakeman
School of Cognitive and Computing Sciences
University of Sussex
Brighton BN1
9QH

September 29, 1998

Abstract

We provide a number of design choices which should be addressed when designing an adaptive application. These choices are common to all adaptive applications, and are fundamental in determining how the application is adapting to changing resources, while providing user utility. We have illustrated their use in the design of an Image Proxy for the WWW (Wakeman *et al*, 1997).

1 Introduction

There are a number of adaptive applications, ranging from those attempting to reduce congestion and limited bandwidth, such as *vat* (Van Jacobson and McCanne, 1992), *vic* (McCanne and Van Jacobson, 1995), *ivs* (Bolot, Wakeman and Turletti, 1994) and *rat* (Perkins, Hardman, Kouvelas and Sasse, 1997), to User Interface Management Systems such as *Amulet* (Myers, McDaniel, et al., 1997) which attempt to adapt to the display capabilities through battery management schemes within the operating system. The rise of middleware has generated renewed interest in providing generic adaptation policies, but we believe that the interfaces to manipulate resource requirements are phrased in the wrong vocabulary

If we regard each component as supplying an OI interface which can be adjusted independently of the other components, the designer has a choice over each component of an implementation. The number of independent axes which contain implementation choices is the degree of freedom of the application. Adjusting the OI of a component can be viewed as placing the it at some new point in the implementation choice, resource usage, application utility space. As we select a new implementation choices to decrease a component's resource usage, we generally decrease the utility of the overall application.

Thus for each of the possible component implementation, we could in theory measure the utility of the resultant application. The designer must understand this space of possible implementations, for it is from it that they choose a degradation path matching resource variations.

3.2 How Trans arent Should the Degradation be?

Having plotted the implementation space, the designer must next decide how the resource usage is to be monitored and controlled. For networked applications, this has generally been a congestion signal such as packet loss. Processor sensitive applications monitor, by examining the process queue size (Kouvelas and Hardman, 1997). The designer must next assess the availability of resources upon which to switch implementations, and whether to include the user within the loop. Systems which exclude the user and use closed loop feedback need to worry about the stability of the control loop, and the effect that varying utility will have on the user. Systems which include the user must ensure that the user is educated about the need for their collusion in adapting the application. Systems in the latter category include battery monitors asking the user about closing applications, or video conferencing system whose adaptation would reduce the quality to less than the minimum demanded (Bolot et al., 1994).

3.3 Run time versus Design time Behaviour S ecifications.

Having investigated the implementation space and decided upon how degradation should be controlled, the designer must now determine when the possible trajectories through the implementation space are decided - when designing the application or through interpreting some set of instructions later on.

Design-time behaviours include the use of reactive protocols for congestion, such as those in the video conferencing tools vic and ivs. These have fixed trajectories through the implementation space. In general design-time policies are easier to encode.

However, as the applications become more generic and are used for more disparate tasks, the semantics of the application are increasingly determined by the context of their use, and the utility of the application is difficult to pinpoint. An application which simply reports the state of some object has simple semantics and will only be used in a very limited set of contexts can have it's behaviour pre-determined. Conversely, a web browser is used in many different ways (Light, 1998) and the semantics of the retrieved pages depend highly upon context. As designers, we can only make best guesses about the utility of various implementations, and so we can only select trajectories through implementation space that approximate the profile of some imaginary user. In real situations, we must allow users to override a generic profile to determine which aspects of the application semantics are important to them at run-time.


```
media.image.* : true : toMono;
```

When the browsers is always used over low bandwidth links:

```
media.image.jpeg : meta(progressive) == false :  
toProgressive, submit default;  
media.* : true : scale 75, submit author;  
media.* : true : reduce 50, submit author;  
media.* : true : compress 10;
```

Policies can be combined with environmental information to select paths of degradation depending upon the available constraints, so if the user wishes to ensure that all download times are less than five seconds, they first attempt to reduce the quality, then if this fails, they scale the object, using the resolution described in the next section.

```
media.* : SIZE/BANDWIDTH>5 : reduce 50;  
media.* : SIZE/BANDWIDTH>5 : scale 75;  
media.* : true : compress 10;
```

An author's policy to make a JPEG fit in the available space:

```
media.image.jpeg."www.site.org/pics/mypic.jpeg":  
MEDIA-WIDTH > DISPLAY-WIDTH : scale (MEDIA-WIDTH  
/ DISPLAY-WIDTH);
```

4.1 Policy Resolution

Policy scripts are compiled together to create a single executable policy. The policy can act on the multimedia object with which it is associated through the various interfaces discussed earlier.

When the policy is activated (asked to transform its media), it goes through a number of parsing and resolution phases. These determine which rules are relevant to the multimedia object, which rules can be removed or overridden by others and they establish a definitive precedence order between rules from different policy sources.

Initially any rule which applies to media other than that of the attached multimedia object are removed. Then rules which clash are resolved according to the following criteria:

1. user rules have precedence over authorial and default rules, except when 3. or 4. is in place,
2. authorial rules have precedence over default rules, except for 3.
3. user and authorial rules can specify submission to default rules.
4. user

- Reduction involves firing rules that utilise the reduce interface; thus reducing the quality of the attached multimedia object and so improving download time. The reduction phase is a multi-pass operation. Each pass of the rules further reduces the multimedia object. Passes are repeated until no rule is fired, that is, all of the media's size and quality requirements are met. Of course this may never happen! So, the multi-pass mechanism is constrained by a set of heuristics which can identify looping, the exhaustive limits of compression and rules which reduce the multimedia object to the edge of our perception.
- Scaling is the process of altering the dimension of the Media. Images are scaled in the X,Y dimension, as is Video. For some media types scaling is not meaningful. It is not meaningful to scale ASCII and it is only meaningful to scale Audio in terms of amplitude or tone, qualities that can be adjusted to suit the user, but do not effect download time.
- Compression is a simple one-pass non-lossy compression. Media that benefit from this are those which, unlike JPEG and MPEG, do not support their own compression. E.g. any mime:text/*, can benefit from non-destructive compression. This multi-phase multi-pass process of resolution is at the heart of what makes the policy work. It combines the disparate needs of user, author and the system designers.

5 Fixed Trajectories: Determining a Default User Policy

There are three good reasons for the existence of a default policy. First, it is necessary to provide a reasonable behaviour for an object if no policy is provided by user or author. Second, it is sensible to insist on certain behaviours unless the user or author overrides them. It is sensible to use lossless compression on objects that implement no compression of their own, such as text, HTML, postscript, etc. We want to be able to provide default policy decisions that only effect the quality of the media slightly but massively improve download time or display usability, such as a small reduction in the quality of a JPEG or the frame rate of video. Finally, the default policy is the basis from which the user can incrementally develop their own policies.

6 Policy politics: Use and Usability

As system designers, we expected that developing an application that attempted to maximise utility under resource constraint would unconditionally be "a good thing". However, for shared applications in which the parties do not necessarily share the same goals, issues arise about which of the parties should retain power.

The degree of contro-

at least users have redress now.

But policies have been designed to allow either side to claim or relinquish control. An author must defer to user preferences where they exist and can only request certain standards of presentation. Conversely, a user may define a policy that is in part overridden if an author has produced a policy and in part defines their definite display and download needs. Thus, they can receive something close to the intentions of the author (if they so wish) or something that is "now readable" on their palm top. In the final analysis, if both parties have conflicting policies, then the user's must triumph. After all, it is the user that must wait while images download and the user who has such special needs as small displays or disabilities that necessitate differing policies. Our prototype software is a workable alternative to browsing with the images disabled and therefore goes a long way to ensuring that the author's vision is delivered, if only slightly modified.

An interesting philosophical perspective of the conflict between the author and user policies is that we are seeing a concrete representation of the post-modern clash over control of the text. Distribution of media over networks has provided another form of distance between the author and the reader, where the network and display may change the experience wished for by the author into something completely different, even before the reader brings themselves to bear. By making the changes forced by networked delivery explicit and configurable, we enable the author and the user to enter into a dialogue about what the media is intended for, and use policy precedence to resolve the asynchronous dispute harmoniously. But this dialogue can only occur if authors account for various alternate representation to their work.

7 To Conclude

In the development of an adaptive application following some sort of design criteria that incorporates the user is fundamental to the success of that application in optimising its utility in the face of an uncertain environment. The use of open implementation in our component design allowed integration of policy code into the application; the application now has a mutable trajectory through implementation space. This maintains the utility of the application to the user through the optimal use of available resources.

We are currently extending the Image proxy to a general proxy architecture in which any media can become active. We are also investigating authoring tools that encourage the generation of alternate representations of multimedia presentations, and that allow authors to express their policies.

8 Bibliography

L. Clark and M. A. Sasse (1997). Conceptual Design Reconsidered - the Case of the Internet Session Directory Tool. In: *People and Computers XII: Proceedings of HCI'97*, August 1997, pp. 67-85, Bristol.

Armando Fox and Steven D. Gribble and Eric A. Brewer and Elan Amir (1996). Adapting to Network and Client Variability via On-Demand Dynamic Distillation. In: *Proc. Seventh Intl. Conf. on Arch. Support for Programming Languages and Operating Systems (ASPLOS-VII)*", October 1996, Cambridge Ma.

- Ian Wakeman, Malcolm McIlhagga and Andy Ormsby (1998). Signalling in a Component Based World. In: *Proceedings of the First IEEE Open Architectures for Signalling*. April 1998, San Francisco, Ca. Van Jacobson and Steve McCanne (1992). Visual Audio Tool - vat Manual Pages.
- G. Kiczales (1996). Beyond the Black Box: Open Implementation. In: *IEEE Software*, January 1996. Gregor Kiczales and John Lamping and Cristina Videira Lopes and Anurag.
- Mendhekar and Gail Murphy (1997). Open Implementation Design Guidelines. In: *Proceedings of International Conference on Software Engineering*, May 1997, Boston Ma.
- Isidor Kouvelas and Vicky Hardman (1997). Overcoming Workstation Scheduling Problems in a Real-Time Audio Tool. *Proceedings of Usenix Annual Technical Conference*, 1997, Anaheim Ca.
- Ann Light (1998). *Interactivity on the Web*. <http://www.cogs.susx.ac.uk/users/annl/tax.html>.
- Chris Maeda (1996). A Metaobject Protocol for Controlling File Buffer Caches. In: *Proceedings of ISOTAS '96*.
- Steven McCanne and Van Jacobson (1995). vic: A flexible framework for packet video. In: *Proceedings of ACM Multimedia*, November 1995, San Francisco Ca.
- Brad A. Myers and Richard G. McDaniel and Robert C. Miller and Alan S. Ferreny and Andrew Faulring and Bruce D. Kyle and Andrew Mickish and Alex Klimovitski and Patrick Doane (1997). The Amulet Environment: New Models for Effective User Interface Software Development. In: *IEEE Transactions on Software Engineering*, June 1997, 23 6, pp. 347-365.
- Colin Perkins and Vicky Hardman and Isidor Kouvelas and Angela Sasse (1997). Multicast Audio: The Next Generation. In: *Proceedings of INET 97*, June 1997, Kuala Lumpur, Malaysia.
- Edward J. Posnak and R. Greg Lavender and Harrick M. Vin (1997). An Adaptive Framework for Developing Multimedia Software Components. In: *CACM*, October 1997, 40 10, pp. 43-47.
- Nick Sharples and Ian Wakeman (1998). *Netbase: Gaining access to Internet Quality of Service from an Application*. Technical report: CSRP 476. School of Cognitive and Computing Science, University of Sussex.
- Jean Bolot, Ian Wakeman and Thierry Turletti (1994). Multicast Congestion Control in the distribution of Variable Bit Rate Video in the Internet. In: *Proceedings ACM SIGCOMM94*, August 1994.