

Towards Measuring the Cost of Changing Adaptive Hypermedia Systems

Nathalie Colineau, Cécile Paris, and Ross Wilkinson

CSIRO – ICT Centre
Locked Bag 17, North Ryde NSW 1670, Australia
{Nathalie.colineau, Cecile.Paris, ross.wilkinson}@csiro.au

Abstract. As applications evolve over time, it becomes increasingly desirable to be able to adapt a system, enabling it to handle situations in different ways and to handle new situations. We refer to this as the flexibility and maintainability of a system. These features come at a cost. We argue here that they are an important aspect of evaluation, and that we need to measure these costs. To start getting a handle on how one might evaluate these aspects of a system (or of an approach), we turned to our own approach to building AH applications and designed a specific study to allow us to look at these issues.

1 Introduction

The last decade has seen the development of approaches and reference models (e.g., [8] and [11]) for Adaptive Hypermedia (AH) systems, and the development of various methods to provide a range of adaptations to users. There are many AH systems in a various domains (e.g., [2], [6], [7], [5], [12], [4]), and the issue of evaluation has become crucial. Evaluation is a complex (and costly) problem, with multiple dimensions. Most evaluations to-date are empirical studies of specific applications under particular conditions, focussing on one dimension: the end-user. We argue that we must look at evaluation from a number of perspectives, including but not limited to the end-user. In particular, we are interested in a system-perspective.

In this paper, we propose to examine the flexibility and the ease of maintenance afforded by the design of a system. As a start towards this goal, we present a specific attempt to evaluate the flexibility of our system through a scenario in which we performed a series of changes requiring an existing application to produce new or different presentations and adaptations.

2 Current Evaluations of AH Systems

A great deal of work has focused on evaluating AH systems with a focus on usability (e.g., [1], [9] and [12]). In most cases, an experiment is set up to evaluate an adaptive system by comparing it to a standard system. While recent work started to look at the systems' performance and the validity of their design (e.g., [3] and [10]), we propose to look at yet another facet of a system-oriented evaluation, in particular the design of the underlying AH architecture with its impact on the maintenance of an application.

This would provide insight into an AH architecture's customisability, its maintainability, its flexibility to be extended or modified, and the degree of reusability it allows. We believe this would help us understand the benefits and costs associated in developing AH applications and start comparing various architectures.

3 Evaluating the Flexibility of Adaptive Systems

Implementing adaptive systems able to handle specific situations in specific ways is important, but not enough. Most adaptive systems are built around a specific application. Using them in the context of a different application is difficult. As applications evolve over time, it is desirable to be able to change an application, enabling it to respond appropriately to new situations: this requires the ability to handle new situations or respond differently than originally envisaged to known situations. This is important for at least two reasons:

- (1) We are designers not domain experts; it is thus difficult for us to know what is the most appropriate in each situation. The design of an application should allow the experts to take control and configure the application appropriately;
- (2) Situations are dynamic – what is satisfactory today may be inadequate tomorrow. We must be prepared to take on board new requirements as they come in.

These requirements come at a cost. We thus believe that there is another side to evaluation: the ease or cost of developing applications that can be easily configured and maintained to meet changing requirements. Note that we are not talking here of computational learning models, but rather we refer to the manual cost of adapting and maintaining an application.

3.1 What Should Be Evaluated?

What should we take into account to evaluate the cost of maintaining adaptive applications? To get a handle on this issue, we started with our own architecture for AH applications. Our aim was to understand the ease with which an application built using our architecture could be maintained and extended. In particular, we looked at:

- (1) *What changes are needed?* – When there are new requirements, do the modifications require the development of new resources, the implementation of additional functionality to the underlying architecture, or both?
- (2) *Who can do it and what is the expertise required?* – Adaptive systems are now quite complex and require a lot of expertise that may be shared among several individuals (e.g., software engineering, HCI, domain expertise, etc.).
- (3) *How hard it is?* – How much effort and time is required to modify the system to the new requirements?

The objective of this work was not to come up with an absolute figure nor a metric, but to start building an understanding of how flexible our system is, of the cost of extending an application, and of how one might evaluate approaches along these dimensions.

In the following sections, we outline our specific AH architecture, the Virtual Document Planner (VDP). We have already used this architecture to build various applications (e.g., [12] and [4]), and we felt that it was flexible and portable across domains and applications. Yet, we wanted to address our questions above more precisely, especially with respect to maintenance and cost.

3.2 The Virtual Document Planner Overview

The Virtual Document Planner (VDP) is a multimedia generation system that dynamically produces hypermedia documents. It is the main module of Myriad, our platform for Contextualised Information Retrieval and Delivery [11]. The VDP belongs to the class of AH systems working with an open set of (potentially heterogeneous) resources by opposition to AH systems based on a fixed and closed world (cf. [2] and [5]). Instead of providing navigation support and guidance through a (manually authored) hypermedia space, the VDP *creates* a tailored hypermedia space. The VDP generates hypermedia documents using discourse rules to select and organise the content, presentation rules to determine an appropriate way to realise the content and the structure of the presentation, and a set of templates to control the layout of the final document. In the two classes of systems, the adaptation methods used are different. For example, while AH systems with a fixed set of resources tailor the content by removing, inserting, altering or stretching fragments of text, open sources AH systems decide what needs to be included, retrieve the appropriate fragments from various sources and assemble them into a coherent whole. This process is done in our case by the VDP planning engine. In addition, while the former class of systems focuses on the development of adaptive navigation support techniques, the manipulation of links is not an issue in the latter class as systems have control over the generated hypermedia space. However, besides those differences, the two approaches present similarities in the way the adaptation is performed. In particular, they base the adaptation on the information recorded in the user profile, task model and other contextual models.

3.3 The Evaluation Study

To assess the flexibility of our system, we defined a scenario where our application was to generate three multimedia documents, taking into account a series of changes regarding both the content (varying the type and the amount of information provided) and the presentation (varying the way information is laid out). We wanted to evaluate: what needed to be done to enable the application to allow for the changes? Who could do it? And how easy/expensive it would be? Our study allowed us to put to the test our approach and architecture, and to assess its current flexibility to fit new requirements. We found that:

- (1) By decoupling the engine/mechanisms from the resources, the changes were applied to the resources, at distinct layers of our architecture (content, presentation, or adaptation mechanisms). Indeed, while the planning engine used by the VDP is generic and can be reused in different applications, the resources (i.e., discourse and presentation rules) are domain and task dependent, and need to be specified and re-written for each application;

- (2) It was hard to anticipate exactly the amount of work, especially for changes involving content. We realised that the cost of providing new content was not related to the amount of new material to be added. It depends both on the role in the discourse structure of the new content and how it fits with existing content, and on the granularity at which the new content is represented, as this influences the number of discourse rules to be written; and,
- (3) The expertise needed to perform these changes was complex. Although the VDP's resources are highly modular and declarative, all expressed in XML, creating and modifying them requires some expertise. The person authoring a target document needs to understand the syntax in which the discourse rules are written, and, more importantly, how they can be combined to generate a particular document.

We realised that the latter was potentially a significant cost of our approach. As a result, we are now undertaking work to provide domain experts with authoring tools that generate automatically the appropriate discourse rules from the specification of a document structure, thus reducing that aspect of cost. Having the authoring tool coupled with our AH architecture, we can now revisit the cost of changing an application, as illustrated in Table 1.

Table 1. Comparative cost of introducing an authoring tool enabling the system to automate content changes from the author's specifications

	VDP without Authoring Tool	VDP with Authoring Tool
What needs to be changed?	Discourse Rules	Specifications of the document structure
Who is doing the changes?	System designers with discourse expertise	Authors for content
How much effort is required?	Several hours	15 to 30 min

Referring back to our original questions (What? Who? How hard?), we have now significantly reduce the complexity of the expertise and the amount of manual effort required (the new or modified discourse rules are derived automatically from the authoring tool). We thus believe that our analysis, albeit only a starting point, has already allowed us to get a handle on the cost of using our approach and has led us to improvements.

4 Conclusion

Situations and requirements evolve; nobody wants to rebuild a system each time a new need comes in. Thus, maintenance of adaptive systems for new situations is an important issue. While the flexibility of tailoring the provision of information to a wide range of situations is key, we must also understand the cost of adapting an application to new situations. Thus, we must understand the cost of extending an application to enable it to generate new tailored hypermedia presentations and

applications. We have attempted to understand these costs for our approach. This is only a first step towards our goal: articulate a framework for evaluating various approaches to AH applications by addressing issues of flexibility and maintainability.

References

- [1] Boyle, C. and Encarnacion, A.O. MetaDoc: an adaptive hypertext reading system. *User Modeling and User-Adapted Interaction* 4(1), 1-19, 1994.
- [2] Brusilovsky, P. Methods and Techniques of Adaptive Hypermedia. *User Modeling and User Adapted Interaction*, 6(2-3): 87-129, 1996.
- [3] Brusilovsky, P., Karagiannidis, C., and Sampson, D. Layered evaluation of adaptive learning systems. *International Journal of Continuing Engineering Education and Lifelong Learning*, 14 (4/5), 402-421, 2004.
- [4] Colineau, N., Paris, C. and Wu, M. Actionable Information Delivery. In *Revue d'Intelligence Artificielle (RSTI – RIA)*, 18(4), 549-576, Sept. 2004.
- [5] De Bra, P., Aerts, A., Berden, B., de Lange, B. and Rousseau, B. (2003). AHA! The Adaptive Hypermedia Architecture, In *Proc. of the 14th ACM Conference on Hypertext and Hypermedia*, 81-84, Nottingham, 2003.
- [6] De Carolis, B., De Rosis, F., Andreoli, C., Cavallo, V. and De Cicco, M.L. The Dynamic Generation of Hypertext Presentations of Medical Guidelines. *The New Review of Hypermedia and Multimedia* 4, 67-88, 1998.
- [7] Fink, J., Kobsa, A. and Nill, A. Adaptable and Adaptive Information Provision for All Users, Including Disabled and Elderly People. *The New Review of Hypermedia and Multimedia* 4, 163-188, 1998.
- [8] Halasz, F. G. and Schwartz, M. The Dexter Hypertext Reference Model, *Communications of the ACM*, 37(2), 30-39, 1994.
- [9] Hothi, J. and Hall, W. An Evaluation of Adapted Hypermedia Techniques Using Static User Modelling. In *Proc. of the 2nd Adaptive Hypertext and Hypermedia Workshop at the 9th ACM International Hypertext Conference (Hypertext'98)*, Pittsburgh, PA. 1998.
- [10] Paramythis, A., Totter, A. and Stephanidis, C. A Modular Approach to the Evaluation of Adaptive User Interfaces. In *Proc. of the workshop on Empirical Evaluation of Adaptive Systems at UM'2001*, 9-24, Sonthofen, Germany, 2001.
- [11] Paris, C., Wu, M., Vander Linden, K., Post, M. and Lu, S. Myriad: An Architecture for Contextualized Information Retrieval and Delivery. In *Proc. of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems (AH2004)*. 205-214, The Netherlands, 2004.
- [12] Paris, C., Wu, M., Vercoustre, A., Wan, S., Wilkins, P. and Wilkinson, W. An Empirical Study of the Effect of Coherent and Tailored Document Delivery as an Interface to Organizational Websites. In *Proc. of the Adaptive Hypermedia Workshop at UM'03*, Pittsburgh, USA, 2003.