

Personalised Web Experiences: Seamless Adaptivity Across Web Service Composition and Web Content

Ian O'Keeffe and Vincent Wade

Centre for Next Generation Localisation, Knowledge and Data Engineering Group,
School of Computer Science and Statistics, Trinity College Dublin, Ireland
{Ian.O'Keeffe, Vincent.Wade}@cs.tcd.ie

Abstract. Users have become accustomed to a web that is more than an interactive hypermedia but is a complex mix of rich multimedia services and hypermedia content. Users are now contributors and active participants on the web. However, Personalisation technologies, such as Adaptive Hypermedia, have so far focused almost exclusively on adaptive content delivery resulting in their failure to become a high impact technologies. The absence of rich multimedia services in the current generation of Adaptive Hypermedia Systems means that they do not live up to the expectations of users. By providing personalised web experiences that combine both services and content in a seamless environment such systems could not only live up to the expectations of users but could exceed them. This paper presents a system that supports the adaptive selection and sequencing of both content and services in a unified manner. By applying techniques used in content based Adaptive Hypermedia to services with making use of the state of the art in service composition, this system delivers personalised web experiences that combine adaptively selected and sequenced content and services. The integration of appropriate content with services can improve the experience of the user as well as making the activity more efficient.

1 Introduction

Users have become accustomed to a web that is more than an interactive hypermedia but is a complex mix of rich multimedia services and hypermedia content. Users are now contributors and active participants on the web, communicating with each other using a range of content sharing and collaborative tools e.g. Gmail, MSN Messenger, Flickr, YouTube, wikis and blogs. They are no longer simply consumers of content but are also content creators and publishers.

In this environment, users seamlessly move between interacting with services, for example voting, rating, annotating, communicating, etc, to content interaction, e.g. viewing, navigating, etc. Furthermore, users frequently combine available services and related content in order to carry out complex activities. An example of such an activity, within an educational context, is a peer review activity, which not only requires access to appropriate content but also requires

services to support authoring, submission, annotation and discussion. Not only do these services need to be made available but they must be presented in a specific order according to the requirements of the activity.

Personalisation technologies, such as Adaptive Hypermedia [1], have so far focused almost exclusively on the adaptive delivery of interactive content, e.g. AHA! [2], KnowledgeTree [1] and PersonalReader [3]. As we move to next generation web technologies, there is a need to provide a combination of adaptive selection and sequencing of multimedia content with adaptive selection and sequencing of user centric services. We define the notion of a **personalised web experience** (PWE) as an experience that involves the integration of the personalised selection and presentation of content, personalised service adaptation and personalised service composition. Thus the personalised web experience provides a significant engagement of the user in carrying out activities on the web.

Presented in this paper is a radical rethink of Adaptive Engines (AE), where the AE supports adaptive composition of web services as well as multimedia web content. Such next generation AEs effectively generate adaptive service workflows and adaptively compose content, seamlessly integrating the adaptive selection, composition and presentation of content and services. This work builds upon existing AE technology and integrates portal and semantic web business process and planning techniques to support the unified AE.

The rest of this paper is structured as follows, the current state of the art in the adaptive composition of content and services is discussed as well as an analysis of the key differences between content and service composition. This is followed by the design and implementation of a system that combines these techniques to deliver PWEs. Finally, an example use case that has been built using this system is then presented.

2 State of the Art

Adaptive Hypermedia Systems (AHS) can be characterised as systems that dynamically compose multimedia content based on the needs of the user. Such systems often personalise the delivery of content based on the interests of the user, tasks that they are undertaking or contextual information such as network bandwidth or screen size.

Systems such as APeLS [4] and AHA! [2] use separate models to represent the user, application domain, available content and the adaptation rules. These systems are all content centric with no support for the composition of services. Another AHS that does take a step towards addressing this limitation is KnowledgeTree [1]. It supports the integration of ‘intelligent content’ [5] into the personalised hypermedia that it generates. Pieces of intelligent content are Java applets that are embedded into the content resources used by KnowledgeTree and integrate directly with KnowledgeTree’s user modelling service, Cumulate [1]. Although this approach does allow for a richer level of user interaction, it does not allow for the composition of computationally complete services. This would require, for example, support for the parameterisation of the services (both

inputs and outputs), the management of the information flow between services and support for error handling.

The techniques used to select and sequence services [6] typically fall into two categories, namely workflow and AI Planning. Workflow based techniques allow services to be composed into a flow of execution that explicitly links services together. Such workflow compositions are typically built manually via rule bases or scripting using technologies such as BPEL [7] and YAWL [8]. Such compositions can be considered to be static as there is limited scope for the workflow sequencing to change and references to the services are explicitly embedded in the workflow. The CAWE framework [9] attempts to apply adaptivity to workflow execution through the use of an abstraction mechanism that allows the workflow to refer to abstract services. Information about the user is then used to select appropriate concrete services or pre-existing service compositions on a just in time basis. This approach supports the adaptive selection of services but not the adaptive composition as the workflow remains unchanged.

The second category of composition techniques are those taken from the AI Planning [6] domain. These attempt to dynamically compose services to satisfy a set of goals using techniques such as Hierarchical Task Networks [6] or situation calculus. In such systems, the services are modeled in terms of the input parameters, the conditions under which the services can be invoked and the effects that the invocation would have. Examples of systems that compose Web Services (WS) using AI planning techniques include XPlan [10] and Shop2 [11].

3 Comparing Content and Service Composition

Unlike content, the behaviour of services in a composition can also be influenced by the adaptive selection of appropriate parameters. Parameterisation of content, although possible, is very seldom offered. This is in contrast to services, which nearly always have heterogeneous input and output requirements and data flow imperatives. Through parameterisation, it is also possible to provide a level of adaptation with respect to the behaviour and functionality of a service. This provides an additional layer of adaptivity when selecting services in comparison to the selection of content. In order to take advantage of this added flexibility, it is important that the mechanism used to select services is aware of the parameterisation of the available services and is capable of configuring some of these parameters so as to modify the behaviour of the service to better suit the needs of the user.

Another important difference between content and service composition is that service compositions have an explicit information flow between the services. This information flow is fundamental to the correct operation of a service composition. If it is broken in some way then the composition will not behave as intended. There exists an information flow between resources in a content composition. However, this flow is more conceptual as it relates to the sequencing of the topics covered by the content. This information flow, although important, for example in an educational context, is not critical for the system to support as

it is typically performed by the user. This is especially illustrated in AHS such as APeLS and AHA! in which users have freedom to jump between different content resources.

In addition to allowing the behaviour of a service to be modified, the parameterisation of services also allows more complex services to be created through the composition of many individual services. The careful management of the information flowing into and out of services can be utilised in order to provide functionalities that were not previously available from existing services.

4 Architecture

A high level view of the system architecture is provided in figure 1. As can be seen from the diagram, the Adaptive Engine is the central component of the system. It is responsible for the reconciliation of the available metadata models through the execution of the narrative (adaptive sequencing rules), which strategically guides the adaptation process. The metadata models describe the available content and services as well as providing information about the learner. The narrative's encapsulation of the strategy for composing a PWE allows a designer to describe the sequencing for a PWE and how it can be adapted.

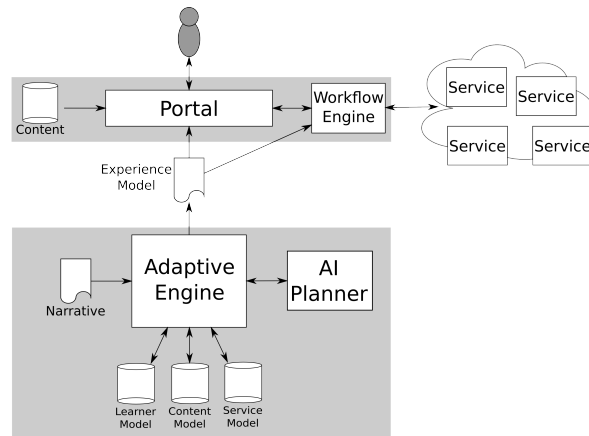


Fig. 1. Diagram representing a high level view of the system architecture

The AE and the narrative that it executes support a basic set of five patterns for describing the sequencing of services. These patterns have been adopted from the work of van der Aalst et al [12] and provide a sufficiently rich range of behaviours to describe the progression between tasks within a PWE. The five patterns supported are: sequence, parallel split, conditional branch, synchronisation and simple merge.

As the AE adaptively sequences the necessary content and services, it makes use of an AI planner to dynamically select appropriate services. The use of the AI planner allows for the adaptive selection of appropriate services as well as the composition of new services when an appropriate service does not already exist. Service selection can be influenced by the AE as it is responsible for setting the selection criteria for the planner.

As a result of the narrative execution, the AE produces an ‘experience model’. This model represents an instantiated PWE in which the appropriate content and services have been identified. In order to make a PWE available to the user, the personalised service composition is instantiated as a BPEL process and deployed to a workflow engine. This brings together the necessary services and controls the users access to those services in order to maintain the sequencing as part of the PWE. The user can then access the PWE through a portal interface that combines the appropriate content and services in a single environment.

5 Case Study

As part of the evaluation of this system, several PWEs were developed. One such PWE was developed based on a personalised eLearning course used to teach undergraduate students. This PWE combines personalised multimedia content covering topics relating to the use of SQL to create, populate and retrieve information from relational databases. The services in this PWE are quiz’s that correspond to the different sections of the course. The learner must get a passing score in a quiz before they can move on to the next quiz. In the case of the Database Retrieval quiz, the learner is also provided with an example database service. This service allows the learner to retrieve information from a database on a topic that they are interested in (the learners are given a choice between two subject domains) in order to answer the questions.

The PWE consists of the delivery of personalised content in parallel with services that have been adaptively sequenced, selected and composed for the learner. The learner is free to browse the content that has been selected and sequenced for them so that they can jump around the hypertext document as the wish. This is not the case with the services. The strategy guiding the sequencing of the services requires that the learner can only access the appropriate services. When answering the Database Retrieval Quiz, the learner has access to both the quiz service and the example database service in parallel.

6 Conclusions

This paper discussed how users on the web are moving from being purely content consumers to being active participants in web based activities. Also discussed was how personalisation technologies such as AHS can enhance the user experience by supporting the adaptive selection and sequencing of services as well as content.

To illustrate how this could be achieved, this paper considered the issues that a AHS would need to address in order to support this functionality and presented

the architecture for a system capable of providing the user with a personalised web experience combining both hypermedia content and multimedia services.

Acknowledgements This research is supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngli.ie) at Trinity College Dublin.

References

1. Brusilovsky, P.: Adaptive Navigation Support. In: Volume 4321 of Lecture Notes in Computer Science. Springer Verlag, Heidelberg, Germany (2007) 263–290
2. De Bra, P., Smits, D., Stash, N.: The design of aha! In: HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia, New York, NY, USA, ACM (2006) 171–195
3. Henze, N., Krause, D.: Personalized access to web services in the semantic web. In: 3rd International Semantic Web User Interaction Workshop, Athens, Georgia, USA (2006)
4. Conlan, O.: The Multi-Model, Metadata driven approach to Personalised eLearning Services. PhD thesis, University of Dublin, Trinity College (2005)
5. Meccawy, M., Brusilovsky, P., Ashman, H., Yudelson, M., Scherbinina, O.: Integrating interactive learning content into an adaptive e-learning system: Lessons learned. In Richards, G., ed.: Proceedings of World Conference on E-Learning, E-Learn 2007, Quebec City, Canada, AACE (2007) 6314–6319
6. Rao, J., Su, X.: A survey of automated web service composition methods. In: Proceedings of the First International Workshop on Semantic Web Services and Web Process Composition SWSWPC, Springer-Verlag (2004)
7. Curbera, F., Andrews, T., Dholakia, H., Golland, Y., Klein, J., Leymann, F., Liu, K., Roller, D., Smith, D., Thatte, S., Trickovic, I., Weerawarana, S.: Business process execution language for web services. (2003)
8. van der Aalst, W.M.P., Hofstede, T.A.H.M.: YAWL: Yet another workflow language. *Information Systems* **30**(4) (2005) 245–275
9. Ardissono, L., Furnari, R., Goy, A., Petrone, G., Segnan, M.: A framework for the management of context-aware workflow systems. In: WEBIST 2007: Proceedings of the 3rd International Conference on Web Information Systems and Technologies. (2007) 80–87
10. Klusch, M., Gerber, A.: Evaluation of service composition planning with owls-xplan. In: WI-IATW '06: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, Washington, DC, USA, IEEE Computer Society (2006) 117–120
11. Wu, D., Sirin, E., Hendler, J., Nau, D.: Automatic web services composition using SHOP2. Workshop on Planning for Web Services (2003)
12. van der Aalst, W.M.P., ter Hofstede, A.H.M., Kiepuszewski, B., Barros, A.P.: Workflow patterns. *Distributed and Parallel Databases* **14**(3) (2003) 5–51