Computer Science and Information Systems 11(1):343–367 DOI:10.2298/CSIS121227012S

AMASE: A framework for supporting personalised activity-based learning on the web

Athanasios Staikopoulos[1], Ian O'Keeffe[1], Rachael Rafter[1], Eddie Walsh[1], Bilal Yousuf[1], Owen Conlan[1], and Vincent Wade[1]

[1] Knowledge and Data Engineering Group,
School of Computer Science and Statistics,
Trinity College Dublin, Ireland
{Athanasios.Staikopoulos, Ian.OKeeffe, Rachael.Rafter, Eddie.Walsh, yousufbi, Owen.Conlan, Vincent.Wade}@scss.tcd.ie

Abstract. Personalised web information systems have in recent years been evolving to provide richer and more tailored experiences for users than ever before. In order to provide even more interactive experiences as well as to address new opportunities, the next generation of Personalised web information systems needs to be capable of dynamically personalising not just web media but web services as well. In particular, eLearning provides an example of an application domain where learning activities and personalisation are of significant importance in order to provide learners with more engaging and effective learning experiences. This paper presents a novel approach and technical framework called AMASE to support the dynamic generation and enactment of Personalised Learning Activities, which uniquely entails the personalisation of media content and the personalisation of services in a unified manner. In doing so, AMASE follows a narrative approach to personalisation that combines state of the art techniques from both adaptive web and adaptive workflow systems.

Keywords: Adaptive Framework, Personalised Learning, Learning Activities, Adaptive Services and Workflow.

1. Introduction

The Internet is increasingly being seen as a replacement for the desktop environment providing an integrated platform in which users interact with rich media content and services to carry out complex tasks. In order to further enhance the user’s experience of the web, personalisation techniques need to be applied to web content, services and the workflow coordination of those services and content. Personalisation aims to ensure that media content and services are selected and tailored according to the user’s personal needs, preferences, goals and context [1].

To enhance and improve the interactivity of the user experience, the next generation of Personalised web information systems needs to be capable of dynamically personalising web media, services and workflow in a unified manner [2]. Typically most existing approaches to personalisation on the web have focused on tailoring only multimedia content, which is restricted to adaptive content, selection and navigation but have as yet omitted to consider adaptive workflow and adaptive services. In addition the majority of such personalisation systems use bespoke or "tailor-made" content and services. In the AMAS project[1] we aim to develop innovative techniques and technologies to address these challenges and to support the dynamic and integrated personalisation of web media and services in the domain of eLearning.

Activity is an important part of learning. For example, Active Learning [3] is a model of instruction that places the responsibility of learning on learners through practical activity. In order to assimilate the learning material, the learners must actively engage with it. Examples of active learning include class discussion and 'think-pair-share'. Situated Learning [4] states that learning is a function of activity, context and the culture in which it takes place. It occurs as a result of social interaction, e.g. through workshops, role-playing, field trips, etc.

Learners learn best when the learning experience offered to them meets their needs and stimulates their preferred modes of learning [5]. When designing an adaptive learning environment, we should consider not only adaptive content and adaptive navigation, but also 'adaptive learning activities' [6]. We need to shift from learning objects that are 'retrieved' or 'accessed' to learning activities that are 'experienced'. Learning activities can be considered as specialised workflows, coordinating learning/educational content and tasks. In this case typical participants of the workflow are the learner and the educator. Many research projects such as LADiE [7] and successful learning activity environments such as LAMS [8] have investigated the pedagogic benefits of learning activities, such as for example, a peer review[2]. Personalised Learning Activities provide all the opportunities of learning activities but with the significant advantages of content, services (tools) and workflow being dynamically adapted to benefit individual learners. This customisation can be based on different "dimensions" of the learning occurrence such as the learner's preferences, prior knowledge, competences and context [9]. In AMAS we define the notion of a Personalised Learning Activity as a learning experience that involves the integration and personalisation of the selection, sequencing and presentation of both content and services.

This paper presents AMASE (AMAS Engine), a core technical framework of the AMAS project, which implements a narrative approach to personalisation for the dynamic generation and enactment of Personalised Learning Activities on the web. The narrative approach to personalisation facilitates a powerful mechanism for offering adaptation that can alter and compose content, services or storylines across a variety of granularities to meet the learning objectives of individual learners [10][11]. AMASE combines the complimentary power of state-of-the-art techniques from the domains of both adaptive web and adaptive workflow systems. The actual adaptation process is a hybrid approach, utilising the capabilities of abstracted workflow and rule-based systems. In general, AMASE has been specifically designed to capture many different learning scenarios for individual and collaborating students with highly adaptive and personalised requirements. The learning experience can be personalised at different stages, for example during the generation of the learning activity (workflow), the re-adaptation of the learning activity due to changing requirements, on enactment (when

[1] AMAS (Adaptive Media and Services for Dynamic Personalisation and Contextualisation) SFI project, please refer to http://kdeg.scss.tcd.ie/amas.

[2] In a peer review activity the learner reviews the work of one or more of his/her peers.

selecting the appropriate learning resources), on composition and customisation of the selected learning resources, as well as a result of a monitoring process where specific progress and engagement criteria are monitored for an individual. AMASE has also a number of built-in features for the generation of domain specific relationships and adaptation rules from primitive constructs and model-based rules to support extensibility and reuse. In order to evaluate our approach and technical framework we have implemented a case study involving undergraduate students of an SQL course.

The remainder of the paper is structured as follows: Section 2 provides an overview of the narrative approach to personalisation. Section 3 presents the AMASE approach and framework to provide personalised activity based learning. Section 4 describes an educational activity developed to teach SQL as part of a third level course. Section 5 outlines and discusses the evaluation results obtained from the use of this educational activity in a real world application. Section 6 presents a state-of-the-art review of related work. Finally, Section 7 summarises the main contributions of our approach and framework.

2. Narrative Approach to Personalisation

The narrative approach to personalisation facilitates a powerful mechanism for offering adaptation that can alter and compose content and services across a variety of granularities. Moreover the narrative approach promotes a separation of concerns, e.g. concepts, models and logic, that also empowers the authoring process, thus making adaptive experiences easier to create and recompose [12]. The key constituents of the narrative approach to personalisation are described below.

A narrative encapsulates strategies through which relationships between concepts in a domain are created and selected, in order to fulfill objectives within that domain. At design-time this strategy (or strategies) is authored in order to represent the variety of conceptual paths that comprise all potential experiences, and how models influence these paths. During execution, this strategy is reconciled with the appropriate contextual models (e.g. user model) to produce a conceptual pathway, tailored toward the specific instantiations of those models. The generation of each individual user's experience involves the runtime binding of specific content and services to concepts in this pathway, or the further refinement of a concept through sub-strategy.

The concepts that make up a domain, including their description, hierarchy and relationships, are of particular importance to the narrative approach. The domain is considered to be the conceptual space in which the objectives of the experience being created are defined and has sufficient coverage, through concept relationships, to specify the likely start points from which a strategy may initiate an adaptive experience. Concepts may be expressed as a hierarchy, with high-level concepts potentially being described by many sub-concepts layered below them. The hierarchy is one form of common relationship that may be used to define the domain. Other forms of relationships may be more focused to the adaptive experience desired. For example, in an educational setting the pre-requisite relationships between learning concepts may be defined in the domain [5].

The activity objectives are often expressed with reference to the domain and usually include reference to one or more concepts. This requires that either the strategy has been authored with specific reference to those concepts or it has runtime access to the domain model. The objectives also typically have a descriptor to indicate how to determine whether an objective has been successfully achieved.

Strategy may be considered at two discrete levels: the means by which measured decisions are made as to which concepts should be included in an adaptive experience; and the means through which the binding of concepts and content/services is achieved. A strategy is an approach (e.g. a set of logic or group of policies) that uses contextual models and the desired objectives to identify an appropriate conceptual pathway and set of guidelines in order to achieve those objectives. The strategy remains agnostic to the content/ services that will be used to realize the constituent concepts and it may be, though not always, created independently to the domain.

Contextual models provide evidence for the strategy to make decisions. They may also be used to guide the binding process. Contextual models may be both dynamic, updated in parallel to the execution of a strategy, and static, existing before the strategy is executed and not altered during its execution. For personalisation the most common dynamic model is that of the user. The actions of the user directly or indirectly alter their model, which is used by the strategy as the basis of decisions. These actions may offer the user very direct control over how the strategy executes, thus leading to a more adaptable and scrutable experience [13].

The output of the strategy execution is a specific concept pathway through the domain towards one or more objectives, that has been tailored to the contextual models and which may also contain guidelines to support binding. The execution of the strategy across a set of contextual models is referred to as reconciliation as the data in the individual models needs to be reconciled with the strategy to decide a concrete pathway.

The final step in the realization of an adaptive experience is the binding of concepts to specific assets, e.g. content and/or services, to fulfill them. This may also be considered as a strategic step as there may be many candidate assets that could suitably realize the concept. Contextual models and the metadata describing the assets may also need to be consulted to achieve this step. For example, the capabilities of a particular device may need to be considered when attempting to deliver specific content.

Binding may also lead to no assets being identified as the concept could be too high-level to find an appropriate candidate. In this instance a solution may be to return to the strategy execution phase to break the concept down into sub-concepts from the domain. After binding has been completed it is possible to deliver the adaptive experience to the user as the concepts have been realized with real content and services.

3. AMASE Technical Framework

The AMASE framework supports the generation of personalised learning activities, combining media content and (user centric) services in a unified manner. The framework is based on a number of interconnected components that generate, execute and support the user's interaction with a personalised learning activity. Fig. 1 presents a high level overview of the AMASE architecture and its components.

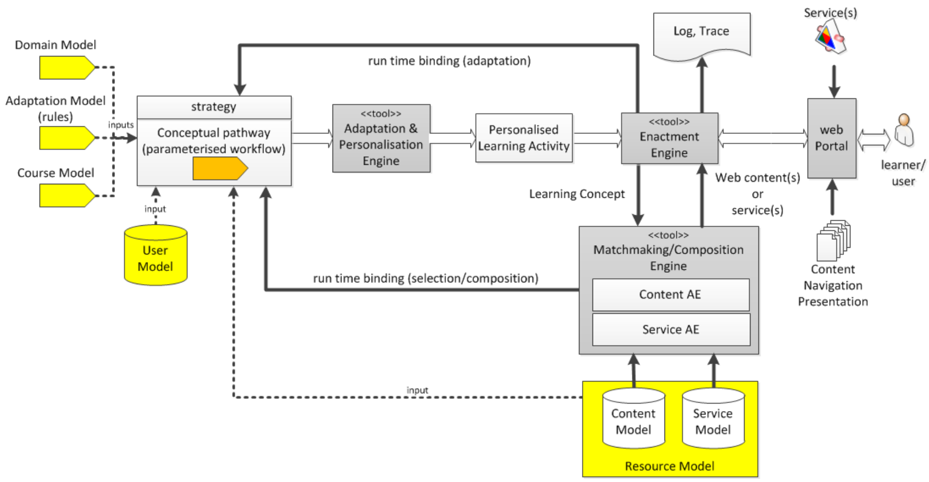

Fig. 1. AMASE Architecture Overview

Initially, an instructional designer by using an authoring tool will specify at an abstract level of detail the strategy – a generalised and parameterised conceptual pathway of learning concepts (content and tasks) that need to be followed for an activity. Next, the Adaptation and Personalisation engine will interpret the strategy, reconcile the appropriate contextual models, apply the adaptation rules and generate the executable personalised learning activity. The executable personalised learning activity will then be initialised and enacted by the Enactment engine. At this stage the personalised learning activity remains abstract as the actual content and services have not yet been selected. The learner's interact with their assigned personalised learning activities via a web based Learning Portal. The framework will resolve (bind) the appropriate content and services to a given learning step (concept) of their personalised learning activity on the fly, via the Matchmaking and Composition component. Once new adaptation requirements are specified and triggered, the adaptation process is repeated at run-time and which will further refine the learning activity of a learner.

The AMASE framework implements a narrative approach to personalisation resulting to a number of key benefits and adaptivity features to support activity based learning. For example it allows AMASE to effectively manage and support the highly dynamic and adaptive nature of real world applications that are driven by pedagogical strategies and learning objectives. It also allows AMASE to address the individual's needs, preferences and context across different requirements and scenarios. The narrative approach to personalisation also allows AMASE to specify and support different modes of learning, including self-directed and group-directed activities within formal, non-formal and self-regulated learning context. It also allows us to provide an agnostic binding, where the actual content and services are selected on the fly to realise a particular learning concept of the generated pathway and re-purpose existing web media and services. From an engineering perspective the narrative approach of personalisation facilitates the organisation and separation of the contextual models from the adaptation logic (rules). Finally, it supports the authoring of new courses and pedagogical strategies, through the reuse of contextual models (domains) and predefined rules.

The following sections explain in detail the components of the AMASE framework and the process followed in order to provide personalised learning activities to learners.

3.1. AMASE Models: Adaptation Capabilities

In AMASE, a number of contextual models are used to support the highly dynamic and adaptive nature of the framework and to provide personalised learning activities to learners according to their individual preferences, needs and context.

Domain Model: Capturing the Conceptual entities

In AMASE the domain model specifies a conceptual namespace where the concepts and their relationships are defined and identified uniquely.

In particular, concepts can be defined across multiple domains, meaning that concepts can be organised together within specific namespaces or topics. In that way, conceptual entities for Newtonian Physics and Astronomy can be organised into different domains, however they can be combined together in a learning activity to teach specific aspects of a course.

In AMASE the adaptation framework has been specifically designed to support extensibility and to combine different types of content and services in a unified manner. As a result, based on a common and abstract definition of a concept, different specialisations are specified. In general, the task-related concepts refer to actions that need to be performed and the topic-related concepts to content that is to be studied as part of a course.

In particular, “LearningTasks” and “LearningTopics” are associated with particular learning requirements that are to be satisfied implicitly (automatically) from the selection of appropriate learning objects. Similarly, there are concept specialisations for “SimpleTasks” and “SimpleTopics” that can refer to any general type of web media source or service. In its core the concept definition specifies a name, a description, a set of keywords and a potential reference to an ontological description. The ontological description of a concept allows AMASE to provide more elaborate decision making based on description logic and related reasoners.

Similarly, the framework has been specifically designed to operate upon various different types of relationships and hierarchies that can be specified among concepts. In particular, we have built concept relationships based on the following primitive types (see Fig. 2). “DirectedRelationships” specify one-to-one directional relationships between a source concept and a target concept. The relationship is only imposed from the source end. “UndirectedRelationships” specify one-to-one relationships between two equivalent concepts. In this case the relationship is imposed by both concept ends. “ComplexRelationships” specify many-to-many relationships among concepts that can participate in directed and undirected relationships. Finally, “Directives” specify relationships between a concept and a predefined term (command).

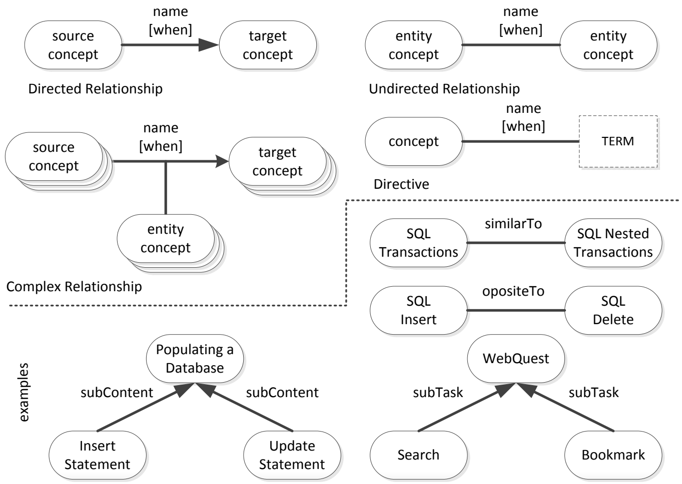

Fig. 2. Conceptual relationships & hierarchies

Upon these general relationship types we have defined specialisations to define “SimilarTo” and “OpossiteTo” relationships (see Fig. 2). A “SimilarTo” is an undirected relationship that is applied equivalently among two concepts to indicate their similarity. For example, when a “SimilarTo” relationship is defined among the “SQL Transactions” and the “SQL Nested Transactions” concepts, then for the selected “SQL Transactions” concept additional similar link(s) are provided, in this case referring to the “SQL Nested Transactions” concept. Quite similarly, an “OppositeTo” is an undirected relationship that is applied among two equivalent concepts to indicate that they are contrary.

In addition, the AMASE framework allows hierarchical relationships to be defined among concepts. For example, the “SubContent” and the “SubTask” relationships are used to define composed-of relationships between content-related and task-related concepts respectively. In that way, we can specify that the “Populating a Database” topic has as “SubTopics” the “Insert Statement” and the “Update Statement” topics, whereas the “WebQuest” task has as “SubTasks” the “Bookmark” and the “Search" tasks. These relationships play an important role on the generation and customisation of the navigation model across concepts as well as to handle the higher level concepts as composite entities. After the concepts and relationships of the domain have been specified, it is then necessary to relate them to the objectives of an activity (see Fig. 3). In that way, we can ascertain if a learning activity that realises a specific course has successfully satisfied its requirements. In AMASE these objectives are mainly categorised as learning objectives and are fulfilled by the interaction or completion of specific learning content and tasks that are part of the personalised learning activity. Based on the learner’s preferences, needs and context these specific learning requirements can be realised differently for each learner. Various in-place monitoring and logging mechanisms trace the learner’s progress and the completion of their objectives, throughout the enactment of their personalised learning activities.

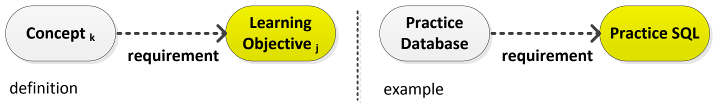

Fig. 3. Relating concepts to objectives

Adaptation Rules

Three different types of adaptation rules are used on the generation of a personalised learning activity. Model-based rules that depict graphically the conceptual relationships and the patterns to match and replace on a strategy and rules that are specified programmatically at low level of detail with a rule language. More specifically, "ConceteRules" are implemented with a specific rule language such as Drools. In this case the rule definition uses a URI to point to the actual implementation file of the rule. It also specifies the parameters that are necessary to be passed to the run-time adaptation environment in order to evaluate the rule. Upon the successful activation of the rule, the body of the rule injects or retracts other specific relationships and rules into the knowledge session that is used to build and personalise the personalised learning activity. All rules at some point will be mapped to primitive workflow operations – that are provided by the adaptation API of the framework, and which will build and personalise the learning activity. The actual rules are stored within a repository, in this case the Drools Guvnor. Fig. 4, provides an example of a concrete rule where if a user has obtained a low score on a test (competency less than 50), a prerequisite relationship between the "Project Phase" and the "Suggesting Reading" is applied. As a result the user will be assigned the "Suggesting Reading" learning task.

```
rule "Recommend"
    when
        um: Usermodel( $cs: competencies)
        Competency (name == 'testSQL', value < 50 ) from $cs
    then
        PreRequisiteRelationship relation = RuleFactory.getInstance().
            buildPreRequisiteRelationship("Project Phase", "Suggested Reading");
        insert( relation );
end
```

Fig. 4. AMASE Concrete Rules

"GraphTransformation" rules provide a more general and flexible mechanism in which more complex and advanced adaptation rules can be defined as search and replace patterns upon the activity graph. Their graphical nature allows learning designers to easy capture their own rules with an authoring tool. In this case two activity graphs have to be provided. A *before_pattern* that indicates the pattern of sequencing concepts that is to be matched in an activity graph and an *after_pattern* that indicates how the activity graph is modified as the result of the match. In addition, a *when* clause is used to define the condition upon which the rule will be triggered. These rules are implemented similarly to the Graph Transformations rules defining the *pre* and *post* conditions on graphs [14]. Fig. 5 provides an example of a graph transformation rule that applies to users with a Constructivism learning objective. In this case the *before* compartment specifies that the A and B in parallel tasks in an activity graph will be replaced by the B and C in sequence tasks that are specified in the after compartment.

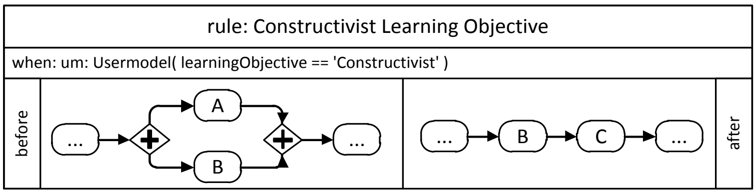

Fig. 5. AMASE Transformation Rules

Finally, there are Relationship rules that specify relationships among Concepts and which trigger the adaptation of learning activities. In this case, upon the fundamental relationships of “Directed”, “UnDirected” and “Complex” relationships, we have predefined a number of rules to express for example “PreRequisite”, “PostRequisite”, “CoRequisite”, “AlternativeTo” and “ReplaceWith” adaptation relationships or requirements (see Fig. 6). These predefined rules are available to a learning designer via the authoring tool (see section 3.2).

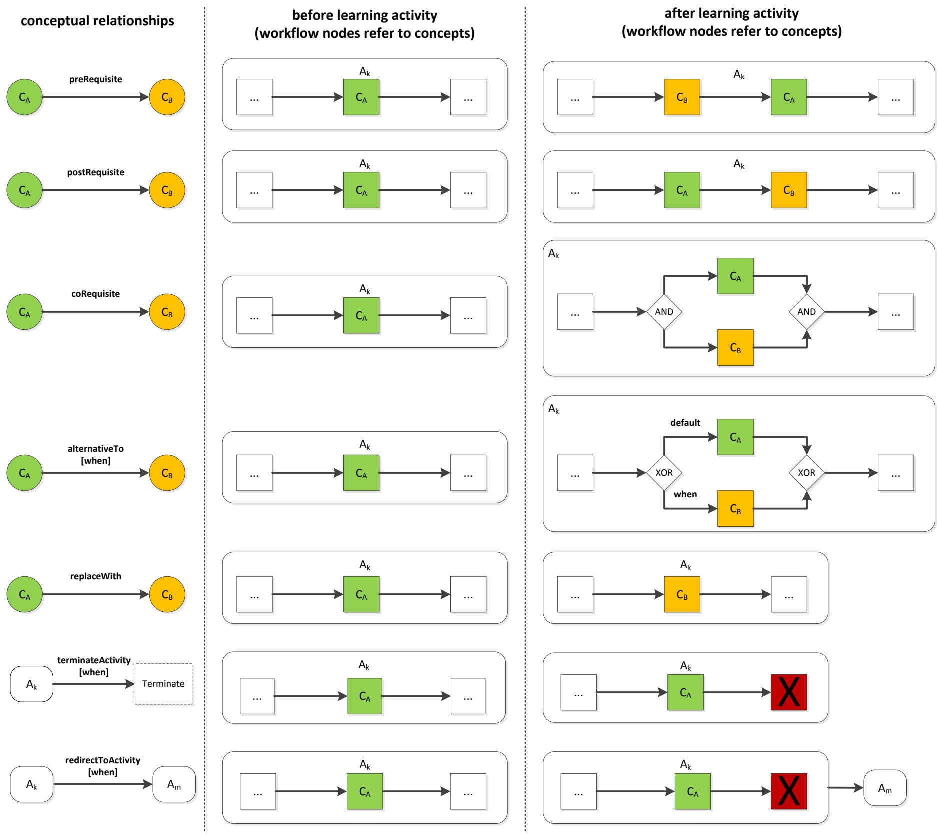

Fig. 6. Predefined Adaptation Rules

More specifically, a “PreRequisite” is a directed rule that applies that the target (B) should precede that of a source (A). For example when a “PreRequisite” relationship is defined among the “Design Database” and the “Implement Database” concepts, the AMASE framework would generate a personalised learning activity where the “Design Database” precedes the “Implement Database” task, even if that task was not part of the initial strategy description (workflow). Quite similarly, a “PostRequisite” is a directed rule that imposes that the source (A) should precede that of a target (B). A “CoRequisite” rule implies that the source and the target should be available in parallel. The “AlternativeTo” rule imposes that can be an alternative selection to target (A) that of source (B), which is available under the successful evaluation of a condition. The condition can be evaluated to true by either the user’s explicit selection or by the run time evaluation of the rule. Finally, a “ReplaceWith” rule replaces the source concept with the target concept.

Strategy: Activity Templates

The strategy is specified as an abstracted conceptual pathway. In effect, a strategy combines the domain model to specify sequences of concepts (content and tasks), a set of rules to apply specific adaptation and personalisation rules, and the user model to parameterise the adaptation and personalisation rules according to the user’s preferences, context and needs. From an implementation perspective, the strategy is captured as a generalised and parameterised workflow (template) and a set of related adaptation rules. Parameters act as placeholders to the workflow template, which specify the variability points upon the common behaviour. Parameterisation allows a strategy to be reused and instantiated differently to a particular context such as for course, domain, group or learner. A template can be parameterised upon the following dimensions:

- The participants or stakeholders – allows binding a workflow node to a specific participant that is a learner or a group of learners.
- Concept Node – allows binding a specific concept (content or task) to an existing workflow node that acts as a placeholder.
- Any Node – similarly but more flexible, allows binding a specific node implementation to a predefined workflow position.
- Sub-Process – allows binding an activity flow to a predefined workflow position.
- Condition – allows modifying the condition on a conditional node.

In AMAS, a course can be associated with many different parameterised learning activities, so under the evaluation of certain conditions, a different starting point can be selected for a specific learner or a group of learners.

Next, the strategy will be interpreted by the Adaptation and Personalisation engine, which will reconcile accordingly the contextual models and generate the personalised learning activity or experience. A strategy can also be used during the binding phase to resolve a learning step (concept) to actual learning content and services by providing the set of selection rules to be used. Finally, new strategies can be specified and applied at run time, for example due to changing of (learning) requirements from an educator.

User Modelling

A User Model captures the preferences, competencies, goals and needs (learning objectives) for each learner. In addition it encapsulates elements that play an important role on the personalisation of the course. For example it captures scores from tests, the level of expertise, preferences on subjects and tools, their level of interaction with the system, their prior knowledge (e.g. MySQL over ORACLE) and learning goals. During the adaptation process such elements play an important role in the parameterisation of the adaptation rules (influence) and the generation of the personalised learning activity.

User Models themselves are stored in a repository, such as an eXist database. In order to extract and combine information about a user from different sources such as their personalised learning environment (e.g. Sakai, Moodle) or social media, we use FUMES [15], a Federated User Model Exchange Service that implements a mapping-based approach to handle heterogeneity across different user models.

3.2. Authoring Phase

An instructional designer by using an authoring tool (e.g., GRAPPLE [12]) will specify at an abstract level of detail the strategy that needs to be followed and the adaptation and personalisation rules that are to be applied for a course.

The aim of the authoring tool is to simplify as possible the authoring process and abstract (hide) many of the implementation details by applying a model-driven approach. As a result an educational designer is agnostic to the underlying mechanisms of the framework. Instead an educational designer relies on graphical abstractions (models) to capture the learning activities and the adaptation strategies. These model abstractions are then automatically transformed to appropriate implementation artefacts. In AMASE the framework has been specifically designed to address the complexity of the authoring phase via the reuse of the domain models and the adaptation rules (see Fig. 6). In particular, the AMASE framework allows new specialised domain rules and relationships to be created from basic constructs (e.g. “DirectedRelationships”) and arbitrary adaptation rules to be easily specified with model-based techniques (see graph transformations in Fig. 5). The authoring tool also provides predefined adaptation rules and pre-configured adaptation tags that can be easily attached (drag and drop) by an instructional designer to learning elements when specifying a strategy. As a result, the educational designer does not need to be an expert on adaptive authoring in order to use the authoring tool and model an adaptive activity such as for an “SQL Database”. Currently, we are developing and testing such an authoring tool that can be used by non-experts for the authoring of personalised learning activities. In general the authoring process is outlined as follows:

The conceptual learning content and tasks that are used in a strategy are defined in a domain model. An educational designer can either reuse a predefined domain model that has been already defined with the authoring tool - for example on “relational databases”, or create his own domain model from scratch. In the latter case, as an educational designer creates the new conceptual entities and relates them with the available predefined relationships, a new domain model is created by the system. Next, the educational designer would have to capture the conceptual pathway of learning content and tasks by using a number of available workflow constructs to indicate for example sequential, parallel or optional execution paths of the learning steps. Finally, upon the strategy the educational designer will have to specify the various adaptation points by selecting and tagging the appropriate workflow entities. The authoring tool provides to an educational designer a number of predefined adaptation rules and tags that an educational designer can use and apply to different courses, groups, learners or learning contexts.

3.3. Adaptation Process: Generating Personalised Learning Activities

Next, in AMAS the adaptation and personalisation engine will interpret the strategy, reconcile the appropriate contextual models, apply the adaptation rules and generate an executable personalised learning activity. The actual adaptation process is based on a hybrid approach combining the advantages and capabilities of workflow and rule-based systems. Rules are used to specify the adaptation effects, evaluate the adaptation conditions as well as to trigger other adaptations. Workflows are used to capture the strategies, as well as support the composition and coordinated execution of learning tasks.

The adaptation process is depicted in Fig. 7. As inputs the Adaptation and Personalisation Engine receives: (i) an optional abstracted learning activity (template) defining partial and parameterised workflows, (ii) a set of adaptation definitions, (iii) adaptation instances to generate the facts that will be inserted in the engine, (iv) a User Model to parameterise the adaptation with the preferences, competencies and objectives of a learner, and (v) a domain model defining the learning concepts and relations.

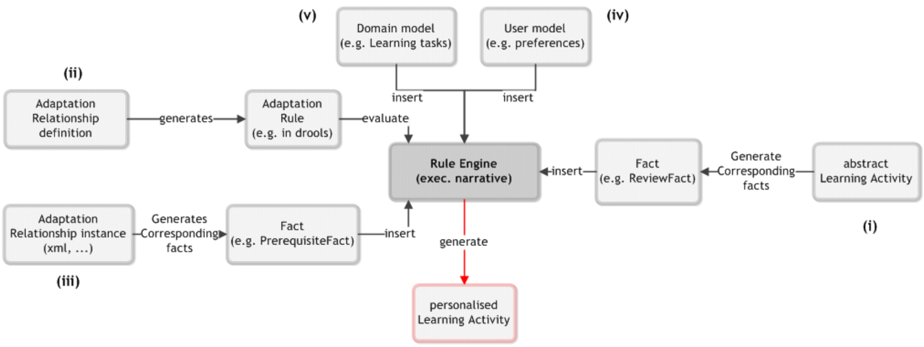

Fig. 7. Adaptation Process

In a case where an abstracted learning activity (template) is not provided, the Personalised Learning Activity is constructed from scratch, exclusively by the adaptation rules. Instead, if a User Model is not provided then the adaptation process continues and generates a learning activity that applies to all users. Finally, in AMASE the adaptation process can be repeated and applied at run time upon an already deployed and executing learning activity. For example, if new learning requirements are identified the Enactment Engine as before would interpret the adaptation strategy, evaluate the adaptation rules and accordingly will trigger an adaptation to further refine the executing learning activity.

3.4. Personalised Learning Activity as Workflows

As a result of the adaptation process, a Personalised Learning Activity is generated as a Business Process Model and Notation (BPMN) workflow specification, ready for execution. At this stage the activity has been personalised but remains abstract, as the appropriate content and services have not yet been selected in order to instantiate the tasks. The next step is to deploy the personalised activity to the Enactment Engine, so that it can be executed and made available to the learner.

3.5. Enactment of Personalised Learning Activity

The Enactment Engine is a jBPM based workflow engine [16] that supports the concurrent execution of multiple learning activities (BPMN workflows) assigned to individual as well as collaborating users. For the enactment of services we consider their functional description, consisting of input, output, precondition, and effects (IOPEs). Learning activities are also stored in a repository, so they can be reused and further customised for different domains and contexts. The current state of an executing learning activity is persistent and stored in a database. This enables us to support long-lived activities as they are dynamically loaded from a database. As the Enactment Engine provides the execution environment for learning activities, it interacts with the strategy, to get adaptation rules that are to be evaluated at run time and accordingly trigger the dynamic adaptation of executing learning activities. The monitoring mechanism monitors the instantiation and initiation of activities.

A learner interacts with the personalised activity through the web-based learning Portal (see Fig. 1). The Portal provides the learner with an environment in which both the content and services that make up the activity are available in an integrated manner. As the learner interacts with the Portal, requests are sent to the Enactment Engine in order to retrieve the appropriate content and services for the learner.

3.6. Binding Learning Activity steps to specific resources

As previously explained the personalised learning activity remains abstract, as for a given learning step the appropriate content and services have not yet been selected. As learners interact with their personalised learning activity via the web-based learning Portal, the framework will resolve (bind) on the fly the learning steps to specific content and services.

More specifically the requests to resolve a particular learning step to learning content and services are initially sent to the Enactment Engine. In turn the Enactment Engine will forward them to the matchmaking/composition component, which will select and compose if necessary, the appropriate content and services on a “just in time” basis. During the selection process a number of contextual models are considered, for example a learner’s model, selection rules, the meta-data descriptions of tasks and the available resources.

The strategy provides the selection rules that are to be applied and which will influence the personalised selection and sequencing of both content and services. The

User Model is used on the matchmaking process to parameterise and influence the selection rules according to the preferences and needs of a learner. The metadata descriptions of tasks and the available resources are also used during the matchmaking process to determine the resources that match the selection requirements of a particular learning step. The metadata descriptions of available resources are stored in the AMASE resource repository.

If the request returns a set of potential matches, it then depends upon the selection directives to select the one or ones to be used. For example the selection can be based on the best match, an arbitrary selection (any), a user based selection, or all possible. Next, the composition service will use a specific UI template to compose or link the selected resources together.

4. Case Study: A Personalised SQL Course

In order to evaluate our approach and technical framework we have implemented an authentic case study, where undergraduate students access an “SQL Database” course for a period of eight weeks and from which they can interact with the Personalised Learning Activity. Next, we describe how the Personalised Learning Activity is generated for two different types of student in Section 4.1, and how they interact with it in Section 4.2.

4.1. Personalisation of the Learning Activity

In general, as part of the course users have to initially practice their SQL skills with a database sandbox environment and then to perform a Web Quest in which they must find and bookmark relevant material from the open web. This is followed by three parallel tasks: that of getting an assignment, designing a database with a design tool and implementing a database. Once users have completed all three activities they continue with the submission of their project. While users perform these tasks, they can participate in an online discussion forum and study specific assigned SQL topics.

In order to illustrate the personalisation of the course we consider two different learners: an expert (user01) and a novice (user02) who has individual learning preferences and prior knowledge. Fig. 8 provides screen captures of the Personalised Learning Activities that are produced for user01 and user02. Due to their different competencies the adaptation process will assign a Practice Database for user02 but not for user01. In particular for user02 the matchmaker will associate that task with an Oracle database tool, due to his/her SQL preference for Oracle (over MySQL). Next, for user02 the general Web Quest task will be replaced by a Questionnaire task, due to his/her learning preference to learn through examples (Constructivism). For user01 the matchmaker will resolve the Web Quest task to a combination of two services, the Search and Bookmarking. Both learners are assigned the Suggested Reading task. For user01 the Study SQL task is parameterised with a few SQL topics that are related to advanced and expert users, whereas for user02 the task is parameterised with more topics that are related to novice users. Next user01, due to his/her prior knowledge of information retrieval, will receive an assignment to implement a meta-search engine, whereas user02 due to his/her background in e-commerce, will have to build an airways reservation system. Both users are provided with the design database, implement database, and forum tasks. Next assuming that for user02, the submission period has expired, the submit task is changed into a late submission task. Finally, the review task will only be assigned to user01.

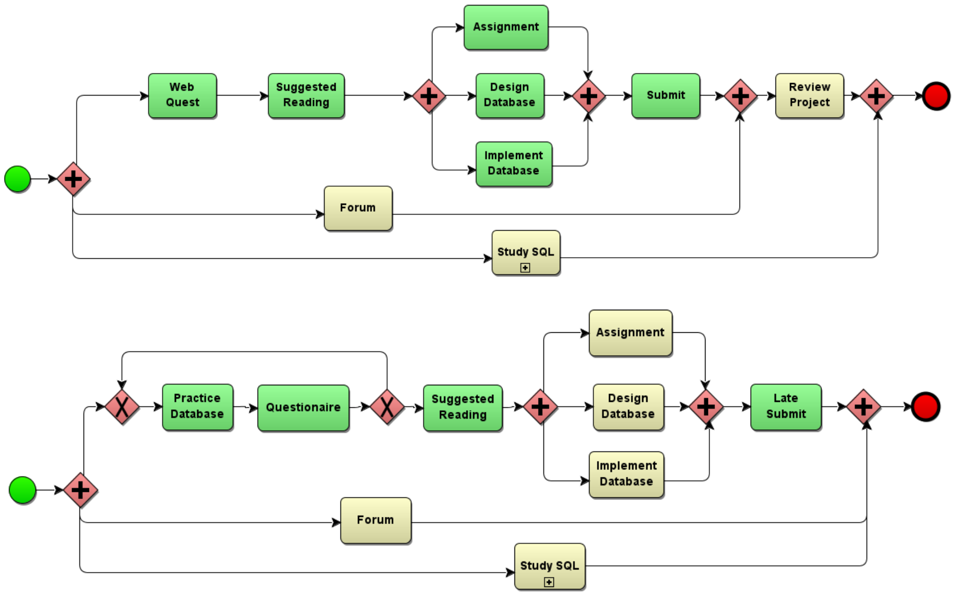

Fig. 8. Case Study for user01 (above) and user02 (below)

4.2. Personalised Learning Portal

Learners interact with the Personalised Learning Activity via a web Portal. The Portal interface is divided into two main areas, the navigation menu on the left hand side of the screen providing access to general information about the course, the tools to be used, the assigned content and tasks, and the content panel on the right hand side. As shown in Fig. 9, once a topic is selected the content panel displays the content with the appropriate navigation options (submenu, buttons). Similarly, once a task is selected the content panel displays the appropriate services with which users can interact (see Fig. 10).

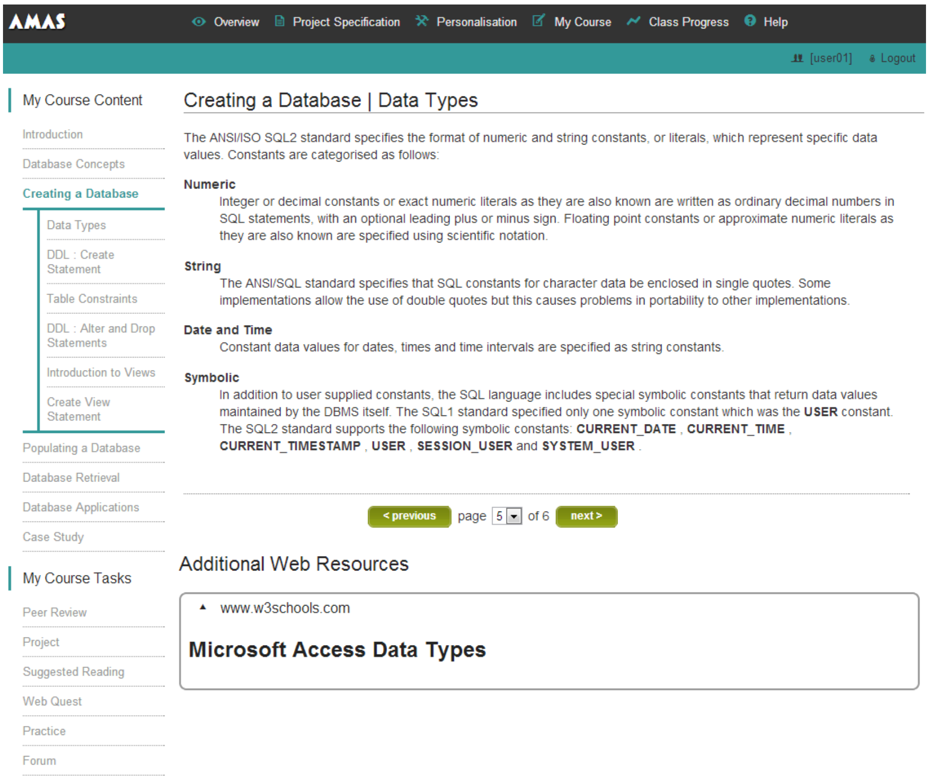

Fig. 9. Personalised Learning Portal for content

In order to realise the different tasks of the Personalised Learning Activity, the services need to be registered with the system. The metadata for the services are stored in a service repository describing the services/tools in terms of a service location (URI), type (e.g., SOAP, REST, Portlet), key-value pairs characterising the service, inputs, outputs, pre-conditions and post-conditions. Both services and tools offered by external and internal providers can be registered. Similarly, metadata about the attributes and characteristics of media content are stored in a content repository. In this case, most of the services are developed as Java portlets deployed in a Liferay Portal server.

In this case study the services used are: 1) a Practice SQL Sandbox Service allowing students to try different SQL commands. 2) a Forum Service allowing for inter-student and tutor-student discussions, 3) a Search Service allowing students to perform web searches based on Microsoft Bing 4) a Bookmarking Service allowing students to keep track of their bookmarked links and submit them as part of the activity 5) a Submission Service allowing students to upload their project reports 6) a Late Submission Service replacing the submission service after a deadline is reached and penalising students for late submissions 7) a Review Service providing students with access to a set of reports that have been randomly allocated to them for review 8) a Recommender Service suggesting selected further reading based on the relevant resources that each student bookmarked 9) a Notification Service notifying students and educators with an email about their allocated tasks and 10) a Questionnaire service that randomly selects SQL questions to test the command skills of some learners.

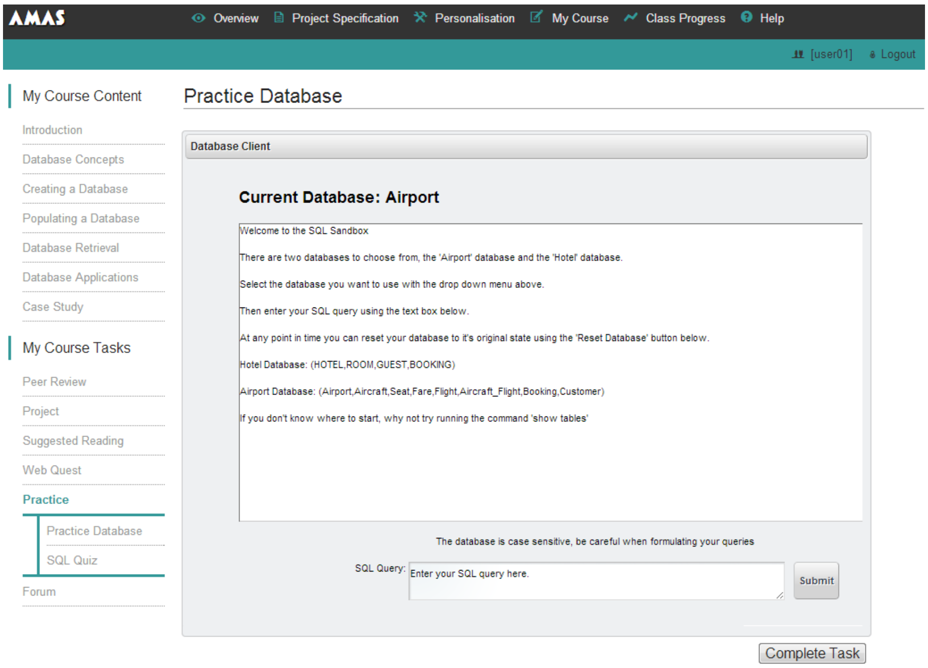

Fig. 10. Personalised Learning Portal for services

5. Evaluation

This is the second year (2012/2013) that the SQL personalised activity has been used as part of the 3rd and 4th year undergraduate courses in Computer Science, Computer Science and Linguistics and Computer Engineering at Trinity College Dublin. Based on the results and feedback from the previous year [17] there were significant updates to the system including a complete redesign of the user interface. New features were also incorporated such as visualizations of the student's progress through the activity. In total we had 101 students take the course this year, all of whom were required to take part in the personalised SQL activity over an 8 week period. Following the completion of the course students were asked to provide feedback through either a paper based questionnaire handed out during a lecture or through an online equivalent of the questionnaire. The questionnaire consisted of 38 statements which the students were asked to gauge how much they agreed or disagreed with each statements using a 5 point Likert scale (ranging from strongly agree to strongly disagree). A SUS [18] usability questionnaire was also used as part of the evaluation.

The primary objectives of the evaluation were a) Students perceived assistance to learning b) Students' perception of the personalisation c) Clarity of usability and d) Controllability of personalisation.

Table 1. Evaluation results for primary objectives

		Positive	Neutral	Negative
Students perceived assistance to learning				
Q1	Upon completion, I felt that I had completed the course objectives	64%	10%	26%
Q2	The activity helped me in learning the subjects covered	52%	16%	32%
Students perception of personalisation				
Q3	The course reflected the questionnaire answers I gave	45%	43%	12%
Q4	The courses generated were easy to navigate	56%	17%	27%
Q5	The content appeared disjoint	49%	31%	20%
Q6	The flow between tasks was appropriate	48%	26%	26%
Q7	Sections contain the content I expected	62%	16%	22%
Clarity of usability				
Q8	The objectives of the generated course were clear	60%	20%	20%
Q9	Upon completion, I felt that I had completed the course objectives	64%	10%	26%
Q10	The courses generated were easy to navigate.	56%	17%	27%
Q11	Interacting with the tasks was intuitive	44%	30%	26%
Controllability of personalisation				
Q12	Have more control on the content included in the customised courses	30.6%	30.6%	38.8%
Q13	Have more control as to how the content was structured	38%	24%	38%
Q14	Have more control on the tools used to complete for a specific task	52%	20%	28%
Q15	Have more control over the flow between tasks in the activity	26.53%	42.86%	30.61%

To gauge how the students perceived the learning activity in terms of its benefit to their learning experience they were asked two questions relating to the course objectives and their ability to learn the subjects covered by the course. The results for these statements are shown in Table 1. To simplify the analysis the results have been aggregated into three categories, positive, neutral and negative. In both cases the majority of students responded positively. However, the number of students that responded negatively (26% and 32% respectively) was substantial. Any activity in which a quarter of the students did not feel that they have completed the objectives has issues that must be addressed. It would seem from the feedback of the students as part of the evaluation that this issue relates to a lack of clarity when expressing the objectives of the activity to the students. Some of the comments provided by students also indicated that for some, the activity was simply seen as extra work that they had to do and that they would have preferred to simply be given a hand-out containing all of the course content.

The student's perception of the personalisation carried out by the system was covered by five of the statements in the questionnaire, the results of which are summarized in Table 1. These statements were designed to elicit feedback from the students on the overall consistency of the learning activity. These metrics can then be used to measure the appropriateness of the adaptive behaviours applied to the activity. The results from statements 3 and 7 provide an indication of whether or not the adaptations carried out by the system are appropriate while statements 4, 5 and 6 look at the same aspect of the system from a slightly different perspective. By asking students about the ease of navigation and flow of the course we can also identify whether or not the selection and sequencing carried out by the system are having any negative impact on the course from the student's perspective. Any mistakes made by the system in the selection and sequencing of content or services would have a direct negative impact on these aspects of the course. A significant result here was the 43% of students who gave a neutral response. It would seem that many students did not fully explore the personalisation controls provided to them where as some students used the feature frequently with one student personalizing the activity 16 times in 8 weeks.

To evaluate the overall usability of the system we first used a standard SUS questionnaire from which we obtained a mean score of 63.9 with a 95% confidence interval of 4.75. This rates the system as 'marginal' or to use the adjective rating scale 'OK' [19]. Ideally we would expect to have a score of 70 or above to be considered acceptable from a usability perspective [19].

In addition to the overall indication of usability provided by the SUS score a series of statements in the questionnaire, see Table 1, were also designed to give further insight into the usability of the system. These results indicate that the general usability issue highlighted using SUS stem from the user interfaces of the individual services and the need to provide more explicit information about the objectives of the activity and its constituent tasks.

The level of controllability by which students could personalise their course was covered by four questions, the results of which are summarised in Table 1. In this case the results indicate that the students have quite balanced and ambiguous opinions. For example 38.8% of students does not want to have more control on the content included, meaning that the existing personalisation options are sufficient, 30.6% found them neutral (or adequate) and another 30.6% want them to be extended. Another 38% of students would like to control further of how the content was structured, 52% would like to have more control on the selection of tools to perform a task (e.g. explicitly select among the available tools or use preferences), and only 26.53% want to further control the sequencing of tasks (e.g. allow users to explicitly select of the learning activity paths).

In addition to the qualitative analysis of the questionnaire results, a quantitative analysis of the system logs was also performed in order to build up a more complete picture of the engagement and effectiveness of the personalised learning activities. From the 101 originally registered students, 88 actually interacted with their personalised learning activities. From these active students 92% completed all of their assigned tasks, whereas the remaining 8% had left behind on their assigned tasks. On average each student interacted 116 times with their assigned tasks and spent in total 7.68 hours on performing them. Similarly, on average each student interacted 171 times with their assigned content and spent in total 4.33 hours on studying their assigned material. In addition, we find that on average each piece of content (171 in total) was visited 88 times by each student over the period of the course. These results indicate that the students found the personalised activity help them to learn the subject.

6. Related Work & Discussion

The application of more activity based approaches in eLearning has been the subject of recent research. The LADiE project developed a set of learning activity use cases based on the experience of educators as part of the projects aim to design a reference model for learning activity authoring and execution [7].

Current integrated learning environments such as Blackboard, Sakai and Moodle provide for the delivery of content and services. However, such environments do not go far enough in addressing the particular needs of a learner via personalisation and suffer from the “one size fits all” problem. Furthermore, despite their support for services, they do not provide any means by which a learning designer can control the sequencing of the services included in their activity. Another common limitation of such learning environments is their closed nature [20], limiting educators to only use services provided by the system.

In contrast, Personalised Learning Environments (PLE) have been a focus for recent research including the EU FP7 ROLE Project [21]. PLEs have developed in response to the restrictive nature of LMS, which limit access to services to only those that are integrated into the system. Similarly, LMS also tend to restrict educators in terms of the pedagogical strategies that they can use [22]. Instead PLEs are focused on the learner, allowing them to construct their own learning environment by selecting services that best suit their needs. As such, PLEs have tended to be applied in non-formal or self-regulated learning contexts [21] rather than more formal contexts in which there might be a requirement for the application of structure to the activity. As PLEs take an adaptable rather than adaptive approach to the construction of the learning environment the focus has been on supporting the learning in the discovery of appropriate services from the large repositories of available services. This has been achieved through a combination of search and recommendation [23] allowing the learner to search for appropriate services while using recommendation to support the learner by tailoring the recommendations to the needs of the individual. This personalisation of the suggested services has been based both on the competencies of the learner [24] and also on pedagogical considerations [24].

On the other hand, specialised adaptive hypermedia systems such as GALE [25] and ADAPT2 [26] handle content adaptivity but fail to address the requirement for services, (ADAPT2 provides limited support for services, treating them as a special type of content allowing very limited capabilities with respect to the type of control flow that can be used to sequence the services). Similarly, GALE provides a general purpose adaptive hypermedia engine that has attempted to provide more sophisticated personalisation, however, it has tended to concentrate on adaptive content selection and composition.

The IMS Learning Design (LD) specification [27] can be used to describe pedagogically driven learning activities using a platform independent language.

However, the specification itself provides only basic support for adaptivity and use of services, supporting only three types of services, namely: an email, a discussion and a search service. Additional services can be added, however the actual implementation of these services is left up to the platform. Where LD-based systems have been extended to support the "adaptive selection" of external services, as is the case with the GSI [28] and the Gridcole [29], their support is limited to the instantiation of abstract service definitions and manual selections, not taking into account the learner's needs.

From the perspective of adaptive services and workflow, eFlow [30] provides an approach for the dynamic composition and enactment of composite services. ADEPT [31] also supports the modification of processes during execution, both at definition and instance level. YAWL supports the dynamic selection of worklets [32] at runtime based on a set of rules that are written by the workflow designer. AgentWork [33] provides the ability to modify process instances by dropping or adding individual tasks based on events and rules. In addition, CAWE [34] is an adaptive workflow system that supports adaptation based on the individual user, the contextual properties and the device they are using. Finally, C-BPEL [35] supports the adaptive selection of services to instantiate the activities in a workflow at runtime. However, most of the workflow approaches outlined do not perform adaptations upon an abstracted and standard workflow language as BPMN, but rather upon concrete implementations that are tied to specific technologies such as for YAWL and WS-BPEL. That means educational designers need to be experts in these languages in order to design a complete and executable learning activity. Services and exchanged data are also hard bound to the workflow, therefore not allowing the dynamic resolution of tasks to services based on their descriptions. There are also even less approaches allowing the dynamic adaptation of workflow instances based on the just in time evaluation of rules. In addition most of these systems they do not consider adaptation from a personalisation and customisation perspective. As a result a user model is not captured and it does not play a significant role in the adaptation process. In addition, there are even less approaches considering the domain specific characteristics and design principles of learning activities.

When compared to the AMAS approach, AMAS takes an activity based approach in learning, which addresses the particular needs of learners via the adaptation and personalisation of learning activities (workflows). Learning content and tasks are integrated in a unified manner. The sequencing of learning content and tasks is fully controlled by the learning designer. Similarly, students can control different aspects of their personalised course by using instrument tools and preferences. Finally, the AMAS approach supports real world applications that are driven by pedagogical strategies and learning objectives specified in narrative descriptions. Depending on the design of the pedagogical strategy, the generated activities can support effectively both self-directed and group-directed activities within a more formal, non-formal and self-regulated learning context.

7. Conclusions

Learning experiences are improved when they are personalised to meet individual needs and stimulate preferred modes of learning. There is also an increasing requirement to combine traditional content-focused learning with more interactive learning activities to provide richer, more dynamic learning experiences.

This paper has presented the AMASE framework to personalised learning activities and the findings from its application in an extensive user trial. The approach has been placed in the context of related work in this area through a discussion of Personalized Learning Environments (PLE), Adaptive Educational Hypermedia Systems (AEHS), and approaches to adaptive service composition. The AMASE approach and framework was then described consisting of a high-level overview of the key components. The architecture was detailed with a core focus on the contextual models and the different phases to deliver personalised activity-based learning. The resulting architecture was implemented and applied in an extensive user trial with actual students. In the trial, learners used a web-based learning portal to interact with automatically generated personalized learning activities in the domain of SQL.

The findings from the user trial have indicated that the majority of students believed the personalised activity help them to learn the subject. However, through the practical application of the approach some shortcomings were identified and lessons learned that will be used to improve the approach moving forward. Overall, the AMAS approach provides novel methods and tools for the delivery of personalised activity-based learning. It presents highly structured and interactive learning experiences, which contrast with many of existing state of the art approaches. The AMAS approach has been applied and evaluated in a practical context with actual students. Results show that the personalised activity-based approach to learning has been positively received by users. Key areas have also been identified as the focus for the future evolution of the approach. That includes comparing adaptive versions of a course versus non adaptive as well as monitoring the level of engagement of a learner as means to identify potential problems and accordingly provide notifications, suggest advices or re-adapt a course.

Acknowledgement. This work was funded by Science Foundation Ireland via grant 08/IN.1/I2103.

References

1. Brusilovsky, P.: Adaptive navigation support. The adaptive web. pp. 263–290. Springer-Verlag (2007).
2. De Bra, P., Aroyo, L., Chepegin, V.: The next big thing: Adaptive web-based systems. Journal of Digital Information. 5, (2006).
3. Bonwell, C., Eison, J.: Active Learning: Creating Excitement in the Classroom. ERIC Digest. AEHE-ERIC Higher Education Report 1. , Washington, D.C (1991).
4. Lave, J., Wenger, E.: Situated learning: Legitimate peripheral participation. Cambridge University Press, 1st edition (1991).
5. Conlan, O., Wade, V.: Evaluation of APeLS – An Adaptive eLearning Service Based on the Multi-Model, Metadata-Driven Approach. In: De Bra, P. and Nejdl, W. (eds.) Third International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH2004). pp. 291–295. Springer Berlin / Heidelberg (2004).
6. Conlan, O., O'Keeffe, I., Brady, A., Wade, V.: Principles for Designing Activity-based Personalized eLearning. Seventh IEEE International Conference on Advanced Learning Technologies (ICALT 2007). 642–644 (2007).

7. Jeffery, A., Conole, G., Falconer, I.: LADiE Project Final Report, http://misc.jisc.ac.uk/refmodels/LADIE/www.elframework.org/refmodels/ladie.html.
8. Dalziel, J.: Implementing Learning Design: The Learning Activity Management System (LAMS). In: Crisp, G., Thiele, D., Scholten, I., Barker, S., and Baron, J. (eds.) Learning. pp. 593–596. ASCILITE (2003).
9. Wade, V.: Challenges for the Multi-dimensional Personalised Web. 17th International Conference on User Modeling Adaptation and Personalization (UMAP). 3 (2009).
10. Staikopoulos, A., O'Keeffe, I., Rachael, R., Walsh, E., Yousuf, B., Conlan, O., Wade, V.: AMASE: A Framework for Composing Adaptive and Personalised Learning Activities on the Web. In: Popescu, E., Li, Q., Klamma, R., Leung, H., and Specht, M. (eds.) Advances in Web-Based Learning - ICWL 2012. pp. 190–199. Springer Berlin Heidelberg (2012).
11. O'Keeffe, I.: Adaptive Composition of Personalised Learning Activities, (2012).
12. GRAPPLE: Generic Responsive Adaptive Personalized Learning Environment, available at http://www.grapple-project.org/, (2011).
13. Bull, S., Kay, J.: Open learner models. Advances in Intelligent Tutoring Systems. pp. 301–322. Springer Berlin Heidelberg (2010).
14. Baresi, L., Heckel, R.: Tutorial introduction to graph transformation: A software engineering perspective. Graph Transformation. 2505, 402–429 (2002).
15. Walsh, E., O'Connor, A., Wade, V.: Supporting Learner Model Exchange in Educational Web Systems. 7th International Conference on Web Information Systems and Technologies (2011).
16. Cumberlidge, M.: Business Process Management with JBoss jBPM, (2007).
17. O'Keeffe, I., Staikopoulos, A., Rafter, R., Walsh, E., Yousuf, B., Conlan, O., Wade, V.: Personalized activity based eLearning. 12th International Conference on Knowledge Management and Knowledge Technologies. (2012).
18. Brooke, J.: SUS: A quick and dirty usability scale, pages 189--194. Usability Evaluation in Industry. (1996).
19. Bangor, A., Kortum, P., Miller, J.: Determining what individual SUS scores mean: Adding an adjective rating scale. Journal of Usability Studies. 4, 114–123 (2009).
20. Dagger, D., O'Connor, A., Lawless, S., Walsh, E., Wade, V.P.: Service-oriented e-learning platforms: From monolithic systems to flexible services. Internet Computing, IEEE. 11, 28–35 (2007).
21. Mikroyannidis, A., Connolly, T.: Responsive Open Learning Environments at the Open University. E-LEARN 2012 - World Conference on E-Learning in Corporate, Government, Health-care & Higher Education. , Montreal, Canada (2012).
22. Steel, C., Levy, M.: Creativity and constraint: Understanding teacher beliefs and the use of LMS technologies. Australasian Society for Computers in Learning in Tertiary Education (ascilite). 1013–1022 (2009).
23. Modritscher, F., Krumay, B., Helou, S. El, Gillet, D., Govaerts, S., Duval, E., Nussbaumer, A., Albert, D., Dahn, I., Ullrich, C.: May I suggest? Three PLE recommender strategies in comparison. in Proceedings of the The PLE Conference. (2011).
24. Berthold, M., Lachmann, P., Nussbaumer, A., Pachtchenko, S., Kiefel, A., Albert, D.: Psycho-pedagogical mash-up design for personalising the learning environment. Advances in User Modeling. 7138, 161–175 (2012).
25. Smits, D., De Bra, P.: GALE: a highly extensible adaptive hypermedia engine. Hypertext 2011. pp. 63–72 (2011).
26. Brusilovsky, P., Sosnovsky, S., Yudelson, M. V: Ontology-based Framework for User Model Interoperability in Distributed Learning Environments. Exchange Organizational Behavior Teaching Journal. 2851–2855 (2005).
27. Consortium, I.M.S.G.L.: IMS Learning Design Specification, http://www.imsproject.org /learningdesign/index.html, (2003).

28. De-la-Fuente-Valentín, L., Pardo, A., Kloos, C.D.: Generic Service Integration in Adaptive Learning Experiences using IMS Learning Design. Computers & Education. 57, 1160–1170 (2011).
29. Bote-Lorenzo, M.L., Gómez-Sánchez, E., Vega-Gorgojo, G., Dimitriadis, Y.A., Asensio-Pérez, J.I., Jorrín-Abellán, I.M.: Gridcole: A tailorable grid service based system that supports scripted collaborative learning. Computers & Education. 51, 155–172 (2008).
30. Casati, F., Ilnicki, S., Jin, L.J., Krishnamoorthy, V., Shan, M.C.: Adaptive and dynamic service composition in eFlow. Advanced Information Systems Engineering. pp. 13–31. Springer (2000).
31. Dadam, P., Reichert, M.: The ADEPT project: a decade of research and development for robust and flexible process support. Computer Science Research and Development. 23, 81–97 (2009).
32. Adams, M., Ter Hofstede, A., Edmond, D., Van der Aalst, W.: Worklets: A service-oriented implementation of dynamic flexibility in workflows. On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE. 291–308 (2006).
33. Robert, M., Ulrike, G., Erhard, R.: AgentWork: A Workflow System Supporting Rule-based Workflow Adaptation. Data & Knowledge Engineering. 51, 223–256 (2004).
34. Ardissono, L., Furnari, R., Goy, A., Petrone, G., Segnan, M.: The Context Aware Workflow Execution Framework. International Journal of Autonomous and Adaptive Communications Systems. 5, 58–76 (2012).
35. Ghedira, C., Mezni, H.: Through Personalized Web Service Composition Specification: From BPEL to C-BPEL. Electronic Notes in Theoretical Computer Science. 146, 117–132 (2006).

Athanasios Staikopoulos is a Research Fellow in the School of Computer Science and Statistics of Trinity College Dublin since 2008. He holds a BSc, an advanced MSc and PhD in computer science from University of Birmingham UK. His research interests include Model Driven Engineering, Domain Specific Languages, Conceptual Modelling and Service Oriented Computing with a focus on adaptive and personalised systems. Athanasios has previously worked on various EU and SFI research projects such as AMAS, ALIVE and CLAMS.

Ian O'Keeffe is a Research Fellow in the School of Computer Science and Statistics at Trinity College Dublin, Ireland. He obtained his PhD in 2012, at Trinity College Dublin following the completion of his Bachelor degree in Engineering. His research interests stem from the domains of technology enhanced learning and digital content management with a focus on adaptivity and personalisation.

Rachael Rafter is currently a Research Scientist in INSIGHT: Centre for Data Analytics, University College Dublin (UCD), Ireland. She has worked as a postdoctoral researcher both in Trinity College Dublin, and in UCD where she was awarded a PhD in Computer Science in 2010. Her research interests broadly focus on recommender systems, personalisation, human computer interaction and machine learning.

Eddie Walsh is currently a Research Fellow in the School of Computer Science and Statistics, Trinity College Dublin, Ireland. He received a BSc and MSc from the University of Limerick, and a PhD from Trinity College Dublin. Eddie's PhD research focused on the integration and exchange of user models to support the personalisation of

user experiences. Eddie has previously worked on various research projects in the areas of user modeling, interoperability, and technology-enhanced learning.

Bilal Yousuf obtained his MSc from Dublin City University in 2004, specialising in electronic commerce. He completed his BSc from Dublin University in 2001 in Computer Engineering. Following on from working as a consultant in systems integration, Bilal returned to academia and is currently pursuing his PhD at Trinity College Dublin. His research interests lie in the field of Visual presentation and Exploration and Visual storytelling.

Owen Conlan is an Assistant Professor in the School of Computer Science and Statistics of Trinity College Dublin since 2004. He holds a Ph.D. in Computer Science with a specific focus on personalisation and its application to Technology Enhanced Learning (TEL). Since his doctoral studies Owen has conducted research into adaptive systems, specifically in the domain of TEL and Digital Humanities. This has led to the narrative approach to personalisation which has underpinned systems such as APeLS and AMASE. He has also performed research into the overlap between personalisation and visualisation, producing highly immersive and adaptive visual environments such as ELEKTRA and 80Days. Owen is the coordinator of the FP7-funded CULTURA project, which investigates next generation research environments for humanities researchers. He has published more than 100 scientific and professional papers.

Vincent Wade is a Professor of Intelligent Systems at the School of Computer Science and Statistic, Trinity College Dublin and holds BSc., MSc and PhD degrees in Computer Science. In 2002 he was awarded Fellowship for his contribution to research in Adaptive Systems. He is Director of CNGL - a world leading research centre focused on Global Intelligent Content and is Academic Director and Founder of the Learnovate Centre for Learning Technology. Vincent has authored almost 300 publications in peer reviewed international journals and conferences and has led over 18 European Union/International Research projects, with significant research funding awards (50m) over the last 12 years. Vincent also previously founded three (spin out) technology companies (Empower The User, Emizar and Wripl) in the areas of Soft Skills Training and Simulation, Customer Care Management, and Personalised Web Services.

Received: December 27, 2012; Accepted: December 10, 2013