

Adaptive Learning Environments and e-Learning Standards*

Alexandros Paramythis and Susanne Loidl-Reisinger
Johannes Kepler University, Linz, Austria

alpar@fim.uni-linz.ac.at

loidl@fim.uni-linz.ac.at

Abstract: This paper examines the sufficiency of existing e-Learning standards for facilitating and supporting the introduction of adaptive techniques in computer-based learning systems. To that end, the main representational and operational requirements of adaptive learning environments are examined and contrasted against current e-Learning standards. The motivation behind this preliminary analysis is attainment of: interoperability between adaptive learning systems; reuse of adaptive learning materials; and, the facilitation of adaptively supported, distributed learning activities.

Keywords: adaptive, e-Learning, standards, personalisation, interoperability

*This is an extended version of a paper presented in the 2nd European Conference on e-Learning (ECEL 2003), November 2003.

1. Introduction

In recent years we have witnessed an increasingly heightened awareness of the potential benefits of adaptivity in e-Learning. This has been mainly driven by the realization that the ideal of individualized learning (i.e., learning tailored to the specific requirements and preferences of the individual) cannot be achieved, especially at a “massive” scale, using traditional approaches. Factors that further contribute in this direction include: the diversity in the “target” population participating in learning activities (intensified by the gradual attainment of life-long learning practices); the diversity in the access media and modalities that one can effectively utilize today in order to access, manipulate, or collaborate on, educational content or learning activities, alongside with a diversity in the context of use of such technologies; the anticipated proliferation of free educational content, which will need to be “harvested” in order to “assemble” learning objects, spaces and activities; etc.

There exist currently several systems which employ adaptive techniques to enable or facilitate different aspects of learning (Brusilovsky, 1999). An important observation one can make going over the related literature is that a dichotomy appears between typically commercial, standards-based e-Learning systems on the one hand, and (typically research prototypes of) adaptive learning environments (ALEs) on the other, with little, if any, standards compliance. It is argued that this dichotomy is, in part, due to the lack of sufficient support for adaptive behaviour in existing e-Learning standards.

In support of this argument, this paper explores the concept of adaptivity in the context of computational learning environments. Furthermore, it attempts a high-level assessment of the sufficiency of existing e-Learning standards for driving the convergence of the two strands of systems outlined above. The intention is to provide a preliminary assessment of the adequacy of existing e-Learning standards for specifying, and guiding the implementation of, adaptive behaviour within learning environments.

The motivation for seeking standardization in adaptive e-Learning is directly linked to cost factors related to the development of ALEs and adaptive courses thereof (e.g., higher initial investment, higher maintenance costs) and the low level of reuse possible in the field today (due to proprietary models and representations of system knowledge, adaptation logic, etc.) (Conlan *et al.*, 2002a). Our rationale can be briefly outlined as follows:

- To protect the high investment necessary for the development of adaptive learning material, one has to ensure that the latter is not bound by proprietary standards and formats. This is a main prerequisite for enabling the transfer of such material to new environments.
- Taking this concept one step further, one may need to ensure that different learning environments can interoperate in the context of adaptation. A typical exemplary setup might involve one environment holding an individual user’s model and interaction / learning history, and another acting as a content repository.

- At the same level, but worth individual mention, is the case of content discovery and aggregation. This introduces an entirely new dimension, as content “characterization” through metadata provided by its initial author / designer, can now be augmented with aspects relating to the use of that content by individuals and groups, and collected as part of the adaptation “cycle”. Furthermore, by combining findings from several compatible systems, which serve the same adaptive course to a multitude of users, it would be possible to make improvements to the course itself. These could be effected wither in a fully automated way, or in a “semi-automated” one, in cases where it would be preferable that no modifications are made to courses without prior approval by human experts.
- Departing from the “traditional” treatment of the learner as a solitary, mostly passive receptor of information, one would also need to account for adaptive support in the context of collaborative learning activities. Such activities may be carried out from within the same or “compatible” learning environments, which, in turn, points to a different level of interoperation requirements between such environments.

The rest of the paper is structured as follows. The next section, “Background”, outlines the main concepts of adaptive personalization in learning environments. The following section, “Adaptation and e-Learning standards”, starts with a brief account of the landscape of related e-Learning standards, and goes on to discuss how these can accommodate adaptivity, and where extensions or entirely new standards are required. Finally, the paper is concluded with an account of the main points put forward and their implications.

2. Background

2.1 What is adaptive learning?

The term “adaptive” is associated with a quite range of diverse system characteristics and capabilities in the e-Learning industry, thus making it is necessary to qualify the qualities one attributes to a system when using the term. In the context of this paper, a learning environment is considered adaptive if it is capable of: monitoring the activities of its users; interpreting these on the basis of domain-specific models; inferring user requirements and preferences out of the

interpreted activities, appropriately representing these in associated models; and, finally, acting upon the available knowledge on its users and the subject matter at hand, to dynamically facilitate the learning process. The preceding informal definition should differentiate the concept of adaptivity from those of tailorability / configurability, flexibility / extensibility, or the mere support for intelligently mapping between available media / formats and the characteristics of access devices. Please note that in several places in this paper, the term “adaptation” is used as a synonym for “adaptivity”.

Adaptive behaviour on the part of a learning environment can have numerous manifestations. Instead of attempting to exhaustively enumerate all of these, we will provide a high-level categorization, which suffices for the analysis in the following section. The broad and partially overlapping categories that we will be referring to are: adaptive interaction, adaptive course delivery, content discovery and assembly, and, finally, adaptive collaboration support. Each of these categories is briefly qualified below, followed by an overview of the models and processes that are typically instated in adaptive e-Learning systems.

2.2 Categories of adaptation in learning environments

The first category, *Adaptive Interaction*, refers to adaptations that take place at the system’s interface and are intended to facilitate or support the user’s interaction with the system, without, however, modifying in any way the learning “content” itself. Examples of adaptations at this level include: the employment of alternative graphical or colour schemes, font sizes, etc., to accommodate user preferences, requirements or (dis-) abilities at the lexical (or physical) level of interaction; the reorganization or restructuring of interactive tasks at the syntactic level of interaction; or the adoption of alternative interaction metaphors at the semantic level of interaction. Although interface adaptations can be thought of as generally independent from the material or “content” delivered through a learning environment, this is not usually the case with learning *activities* - the major differentiating factor being the emphasis on ensuring and optimising “content” attainment in the former case, versus the emphasis on supporting a process in the case of activities. The dependency of learning activities on interface adaptations is a natural consequence of the fact that the interface encapsulates the

very “tools” for carrying out an activity, be it interpersonal communication, collaboration towards problem-solving, etc.

The second category, *Adaptive Course Delivery*, constitutes the most common and widely used collection of adaptation techniques applied in learning environments today. In particular, the term is used to refer to adaptations that are intended to tailor a course (or, in some cases, a series of courses) to the individual learner. The intention is to optimise the “fit” between course contents and user characteristics / requirements, so that the “optimal” learning result is obtained, while, in concert, the time and interactions expended on a course are brought to a “minimum”. In addition to time and effort economy, major factors behind the adoption of adaptive techniques in this context include: compensating for the lack of a human tutor (who is capable of assessing learner capacity, goals, etc., and advising on individualized “curricula”), improving subjective evaluation of courses by learners, etc. The most typical examples of adaptations in this category are: dynamic course (re-)structuring; adaptive navigation support; and, adaptive selection of alternative (fragments of) course material (Brusilovsky, 2001).

The third category, *Content Discovery and Assembly*, refers to the application of adaptive techniques in the discovery and assembly of learning material / “content” from potentially distributed sources / repositories. The adaptive component of this process lies with the utilization of adaptation-oriented models and knowledge about users typically derived from monitoring, both of which are not available to non-adaptive systems that engage in the same process. At this point, we would like to make an explicit distinction between the perspective of the individual learner wishing to locate relevant material within a (possibly constrained) corpus, and the perspective of the author or “aggregator” who undertakes the task of putting together a course from existing materials and targeting a specific audience – or, seen differently, collecting and tailoring material for accommodating specific user / context characteristics. Although adaptation may very well be suitable in both perspectives, in the context of this paper we will be focusing on the first one, i.e., the assembly and contextualisation of material that is intended for an individual learner. This allows us to consider the more complex scenarios that emerge when one’s personal learning and

interaction history can be utilized to infer criteria for content selection and processing.

The fourth and final category, *Adaptive Collaboration Support*, is intended to capture adaptive support in learning processes that involve communication between multiple persons (and, therefore, social interaction), and, potentially, collaboration towards common objectives. This is an important dimension to be considered as we are moving away from “isolationist” approaches to learning, which are at odds with what modern learning theory increasingly emphasizes: the importance of collaboration, cooperative learning, communities of learners, social negotiation, and apprenticeship in learning (Wiley, 2003). Adaptive techniques can be used in this direction to facilitate the communication / collaboration process, ensure a good match between collaborators, etc.

2.3 Models in adaptive learning environments

All of the above categories of adaptation in learning environments are based on a rather well-established set of models and processes. The rest of this section presents brief accounts of some of the models that one typically encounters in ALEs.

- *The domain model:* Since most current ALEs are focused on adaptive course delivery, the domain-, or application-model is usually a representation of the course being offered. However, in those cases where more general learning activities are supported, the domain model may additionally contain information about workflows, participants, roles, etc. The most important aspect of adaptive-course models is that they are usually based on the identification of relationships between course elements, which are subsequently used to decide upon adaptations (Brusilovsky, 2003).
- *The learner model:* The term learner model is used to refer to special cases of user models, tailored for the domain of learning. The specific approach to modeling may vary between adaptive learning environments. Nevertheless, there is at least one characteristic shared by practically all existing systems: the model can be updated at interaction time, to incorporate elements or traces of the user’s interaction history. In other words, the learner model in ALEs, not only encapsulates general

information about the user (e.g., demographics, previous achievements, etc.), but also maintains a “live” account of the user’s actions within the system.

- *Group models:* Similarly to user / learner models, group models seek to capture the characteristics of groups of users / learners. The main differentiating factors between the two are: (a) group models are typically *assembled* dynamically, rather than “*filled in*” dynamically, and (b) group models are based on the identification of groups of learners that share common characteristics, behaviour, etc. As such, groups model are used to determine and “describe” what makes learners “similar” or not, as well as whether any two learners can belong to the same group. This dynamic approach to identifying groups and user participation in them is already used widely in collaborative filtering and product recommenders, and bears great promise in the context of e-Learning.
- *The adaptation model:* This model incorporates the adaptive theory of an ALE, at different levels of abstraction. Specifically, the (possibly implicit) adaptation model defines *what* can be adapted, as well as *when* and *how* it is to be adapted. The levels of abstraction at which adaptation may be defined, range from specific programmatic rules that govern run-time behaviour, all the way to general specifications of logical relationships between ALE entities, that get enforced automatically at run-time. The most widely known ALEs today (e.g., NetCoach (Weber, and Brusilovsky, 2001), AHA! (De Bra *et al.*, 2002b), InterBook (Brusilovsky *et al.*, 1998), etc.) use adaptation models that generically specify system behaviour on the basis of properties of the content model (such as relationships between content entities).

Although there would be probably little contention as to the enumeration of the models encountered in ALEs, the related literature reports a proliferation of approaches in their representation and utilization within different systems (Brusilovsky, 2003). It is argued that this is one of the major stumbling blocks that stand between adaptation and the e-Learning mainstream today. Awareness of this problem has given rise to several research efforts, aimed at standardizing as much of the adaptation modelling process as possible, on the basis of existing standards (see, e.g., the “Workshop on Adaptive E-Learning and

Metadata” carried out under the auspices of the WM2003 conference - <http://wm2003.aifb.uni-karlsruhe.de/workshop/w05/>). The “reuse” of existing e-Learning standards and their “retargeting” for use in the context of adaptation, which is also a premise of this paper, is intended to: (a) facilitate the smooth and gradual transition from existing non-adaptive learning environments and courses to their adaptive counterparts, and (b) enable the graceful downgrading of adaptive content and activities when delivered over, or supported by, a “traditional” learning environment.

3. Adaptation and e-Learning standards

There currently exist numerous organisations, consortia, etc., that are working in the area of e-Learning standards. For instance organisations like the Dublin Core Metadata Initiative, the IEEE, the IMS Global Learning Consortium, the Alliance of Remote Instructional Authoring and Distribution Networks for Europe, the Aviation Industry CBT Committee, the Advanced Distributed Learning Initiative, etc. are dedicated to, or have committees and working groups active in, the establishment of e-Learning standards.

It is beyond the scope of this paper to enumerate all entities involved in the establishment of e-Learning standards, or the standards themselves. Instead, the authors have opted to make selective references to some of the standards, where such references are relevant to the ongoing discussion. Nevertheless, it should be noted that the core of standards that have been analysed and are referred to in the subsequent sections are the various specifications of IMS¹, ADL SCORM², the set of standards previously known as “PAPI”³ (henceforth referred to simply as PAPI), and the AICC specifications⁴.

In the following, we first delineate the main problems not addressed by today’s standards and then proceed to identify what we consider as necessary additions / enhancements to them, as well as point out requirements that necessitate the evolution of new standards.

¹ <http://www.imsproject.org>

² <http://www.adlnet.org>

³ <http://jtc1sc36.org/>

⁴ <http://www.aicc.org>

3.1 Adaptation-oriented “domain” modelling

Current standards and concepts for educational metadata focus on content-centred approaches and models of instructional design. Scenarios that concentrate on how to structure and organize access to learning objects are mirrored in concepts such as content packaging. Standards focus on search, exchange and re-use of learning material, often called content items, learning objects or training components. The Learning Object Metadata specification, in particular, aims at metadata to facilitate the generation of consistent lessons composed of de-contextualised and distributed learning objects (e.g., consistence in the level of difficulty). Its vision is to enable computer agents to automatically and dynamically compose personalized lessons for an individual learner. The IMS Learning Design specification goes a step further, by providing a conceptual model that enables authors to describe processes and activities including social interaction. The MASIE Centre Report (MASIE Centre, 2002) identifies four main uses of metadata today: categorisation of content, generation of taxonomies, reuse, and dynamic assemblies. All uses are directly or indirectly relevant to adaptation / personalisation.

As already mentioned, current, generic ALEs that support adaptive course delivery require an additional level of information about the entities that make up a course, namely the interrelationships between the entities (Brusilovsky, 2003). The primary goal in seeking standardisation in this dimension is to make it possible to have declarative definitions of relationships and concepts, leaving their procedural interpretation and implementation to each ALE. Using these, different systems may choose to provide different adaptive features or support different types of personalisation, much in the same way that systems differ in how they present standardised modules.

(De Bra *et al.*, 2002a), for example, address the definition of higher-level concept relationship types and the automatic translation of instances of such types into lower-level adaptation rules for the AHA! adaptive e-Learning system. Some of the relationship types discussed therein denote direct relationships between concepts and learning elements (e.g., concept A *is a prerequisite for* concept B, element X *exemplifies* concept C), while others bear a clear adaptation / knowledge inference flavour

to them (e.g., element Y *when read provides knowledge towards* concept D, or, element Y *when read indicates interest in* concept E).

At a lower level than De Bra, we also need to be able to define “assets” associated with “learning objects / elements” which can have standardised relationships to each other and to the enclosing object. Consider, for example, two *mutually exclusive elaborations* of a given concept, one being *brief* and the other *detailed*; contrast that with two *complementary elaborations* of a given concept, the first being a *required brief* reading, while the second being an *auxiliary amendment* to the first. This also implies the possibility to define learning elements that are (more or less) atomic chunks of learning material, distinct from “pages” and with arbitrary granularity (e.g., a paragraph).

Currently, defining relationships such as the ones described above, can be achieved through the use of Learning Object Metadata, if the following conditions are met:

- A “vocabulary⁵” is developed defining the relationships between concepts, as well as the characteristics of these relationships (e.g., transitivity), so that their interpretation by application software is not open to interpretation.
- Every learning entity that is an individual “concept” has an associated LOM-compliant metadata record.
- The entity’s metadata specify the entity’s relationships with other entities, using the aforementioned relationship vocabulary and the entities’ identifiers.

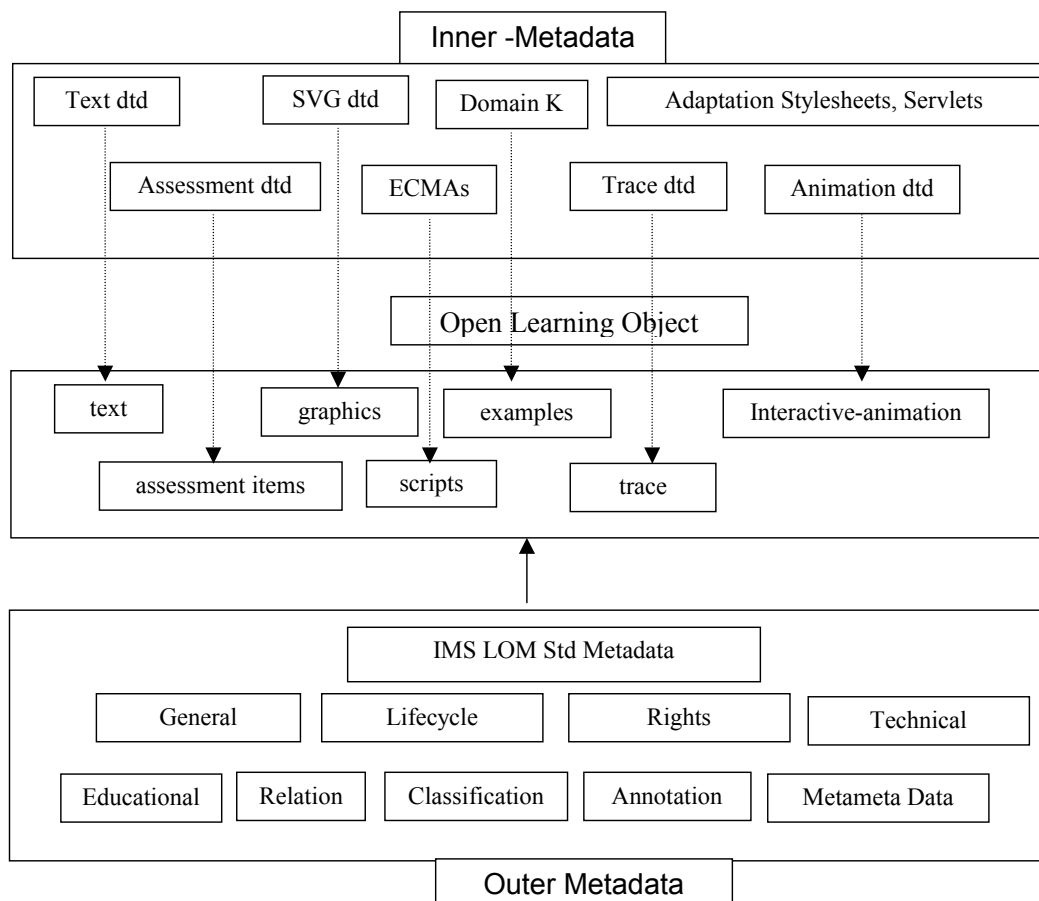
This approach has the benefit of compliance with current standards, and requires only the introduction of a new, adaptation oriented vocabulary for relationships. A similar approach would be to introduce dedicated (optional) adaptation-specific constructs in the main course description. The latter, however, would evidently require modifications to standards commonly used to define courses, which may be considered a much higher (as compared to the above approach) “entry cost” for introducing adaptation in e-Learning standards. A third option would of course be to keep adaptation-related information / metadata separately than the description of the course itself. This has the benefit of rendering the two rather independent, but would most likely prove problematic in terms of course maintenance. This is especially the case as far

⁵ Alternatively referred to as a “value domain”, for which the “permissible values” are well specified.

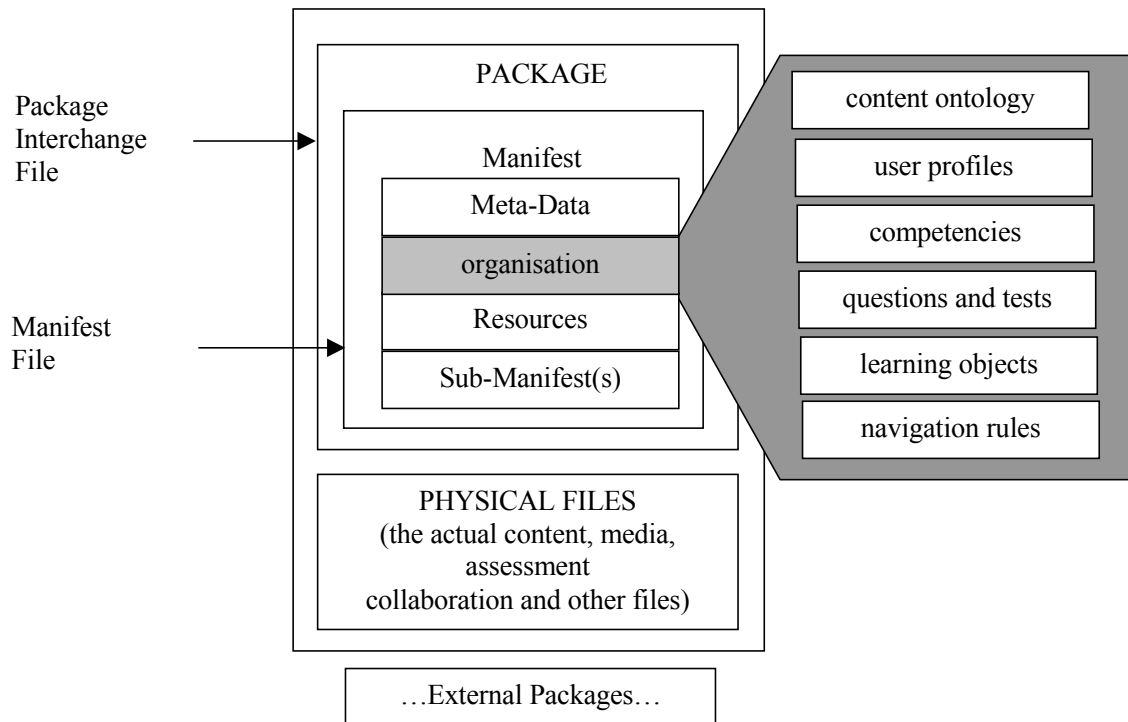
as “synchronisation” between the two is concerned.

Thus far we have discussed the case of characterising relationships between existing course objects / elements. However, as pointed out in (Brusilovsky, 2003), some types of adaptation require a model that is different than (although connected to) the main course model. For example, a model of course concepts and their semantic relations may need to be maintained “separately” from the model of physical course-material organisation (e.g., files, navigation hierarchy). Apparently, whether the two are separate or not, there must exist associations from one to the other, so that the system knows which concepts correspond to given resources, and vice versa. Standardisation in this direction would evidently necessitate new standards: such concerns are beyond the traditional approaches to organising and describing course material and activities.

Examples of ALEs that extend existing standards to support adaptive course delivery include OPAL, OLO and KOD, among others. OPAL (Conlan *et al.*, 2002), which delivers content personalized to the learner’s cognitive and presentation learning preferences using aggregation models based on ADL SCORM. OLO (Rodriguez *et al.*, 2002) and KOD (Karagiannidis *et al.*, 2001) (see Figure 1, (a) and (b) respectively) both address the topic of extending the metadata that accompanies “packaged” learning objects, with the intention to facilitate adaptation. Although the projects take considerably different routes, they are largely motivated by the same objective, to augment the “traditional” metadata with additional elements that are vital when one is to decide upon, and apply course-oriented adaptations. Furthermore, both projects attempt to “integrate” adaptation metadata with the traditional course information (e.g., KOD incorporates the adaptation logic –rules– in an extension to the *organization* element of IMS CPS).



(a) Open Learning Object and Inner Metadata (from Rodriguez *et al.*, 2002)



(b) The IMS Content Packaging Standard and proposed extensions by the KOD project (from Karagiannidis *et al.*, 2001)

Figure 1: Proposals for adaptation-oriented extensions to the metadata accompanying “packaged” learning objects.

3.2 Learner and group modelling

Learner modelling in existing standards is addressed at a rather coarse-grained level, although all related specifications have explicit provisions for the evolution of a learner’s model, or profile, over time. An example of specifications in this strand is the IMS Learner Information Package specification, which incorporates the results of “top-level” educational activities, in addition to relatively static information about the user (e.g., demographic).

Although this information is of paramount importance for e-Learning systems, the coarse-grained level of detail renders them of limited use in the context of ALEs. The main underlying problem is that ALEs require a “history” of the user’s interactions, in order to be able to tailor themselves to the particular needs of the individual user. Furthermore, this “history” is more often than not closely associated with the domain model itself (e.g., the course model). Consider, for instance, the very common desideratum (in ALEs) of basing adaptations on the user’s familiarity with a given concept. This requires the establishment of a new set of relationships, which codify a learner’s “status” with respect to a learning entity or concept. Such relationships may refer

to directly observable learner behaviour (e.g., whether a learner *has read*, or *has not read* a node in the learning material), or to inferred status drawn from multiple sources, including results of exercises, etc. (e.g., *knows*, *does not know*, or *is ready for*).

Arguably, the only standard available today that has extended provisions for modelling fine-grained user activities is PAPI. The PAPI standard reflects ideas from intelligent tutoring systems where the performance information is considered as the most important information about a learner, and also stresses the importance of inter-personal relationships (Vassileva *et al.*, 2003). The strengths of PAPI in relation to ALEs stem from its support for representing learner activities in quite structured manner and in as great detail as necessary. Further to the above, PAPI provides a variety of bindings (multiple codings, APIs and protocols), which facilitate its employment in different scenarios within ALEs.

Although PAPI might be more appropriate for modelling users in the context of adaptive ALEs (as compared, for example, to IMS LIP), it is far from being adequate in all its dimensions. Dolog and Nejdli (2003), for example, report on recent work carried out in

the context of the EU/IST Elena project⁶, towards the development of the first version of a learner profile to support simple personalization techniques. To cater for omissions or weaknesses in each individual standard (as identified through scenario development and analysis), the RDF-based learner profile they propose is based on subsets of both IMS LIP and PAPI.

Whichever standard (or combination of standards) one might use as a basis for standardisation in ALEs, there exists an additional issue that needs to be addressed. Specifically, it would be necessary to agree upon ways of deriving portions of the learner model from the domain / course model (at least for as long as the learner is “taking” a course), as well as upon when and how such detailed information gets “summarised” into the more coarse-grained models that exist today. This is of particular importance in the case of ALEs that employ what are known as “overlay” models, to relate the learner’s current progress in a course, with the course model itself.

The discussion, thus far, has been restricted to the modelling of learner interactions in the context of encountering and assimilating course material. The conclusions drawn, however, are applicable to learner activities at more general scopes. For example, by recording users’ social interactions and allowing for their characterisation by the users themselves, it becomes possible to adaptively facilitate a wide range of interpersonal exchanges, as well as targeted collaborative work.

It may be argued that such learner “history” information is an internal concern of ALEs, and, since it does not need to be specified prior to the deployment of learning material, it is not subject to standardisation. This, however, would most likely preclude use of the aforementioned information in adaptive behaviour other than course delivery. Consider the following examples in support of this view:

- An intelligent learner support agent sets out to discover auxiliary learning material for a given user. Having access to detailed information about what the user has already learned (or, what the user has not learned yet) the agent is far more likely to discover more contextually relevant items than would be possible otherwise.

- A newly created course is characterized by its authors as “fast” and “introductory”. Nevertheless, in practice, students need to spend three times the anticipated time and effort before they can get an acceptable level of familiarity with the material; additionally, upon completion, students are capable of solving problems from an associated repository at all levels of difficulty. It should be clear that selecting this course purely on the basis of its associated metadata might lead to serious mistakes (e.g., in the process of content filtering). Adding information from its actual use provides a more “informed” view of the course and has the potential to lead to better personalization as a direct consequence.

Maintaining detailed information about a user’s activities within an ALE also gives rise to a new opportunity in terms of group identification and modelling. Specifically, if one can refer to learner activities in a standardised way, then one can also identify dimensions of activities that should be used as predictors or measures for determining group membership. For example, one could identify that learners are to be grouped along the dimension “willingness to interact with peers”, which is to be inferred from (among other things) the user’s active participation in on-line discussion fora.

Unlike the case of learner modelling, group modelling as discussed in this paper is only cursorily covered by existing standards. In fact, PAPI seems to be the only specification that provides sufficient support for describing the characteristics of groups. However, the very features of PAPI that constitute its strengths in the case of learner modelling, turn into potential stumbling blocks in the case of group modelling: PAPI is mainly oriented towards the activities / performance of individuals or groups; however, it makes no explicit provisions for describing the characteristics / attributes that are shared between the group participants. As a result, semantic information over what actually qualifies a person as a member of a group can only be indirectly modelled.

3.3 Adaptation modelling

The issue of modelling the behaviour of any adaptive system has two complementary but distinct dimensions, which we will examine separately: the specification of adaptation logic, and the specification of adaptation

⁶ <http://www.elena-project.org/>

actions. The former is responsible for relating information available in one or more models and assessing whether adaptations are required. The latter refers to specifying the very actions that need to be effected by the system for a given adaptation to be achieved.

Attempting to standardise the way in which adaptation logic is expressed would be, in the authors' opinion, rather premature at this point in time. Existing approaches include simple rule-based engines, case-based reasoners, etc., all the way to powerful logic-based reasoning engines. Given this wide range of approaches in use, it is apparently unrealistic to aim at a single specification that could accommodate them all. On the other hand, developing a range of specifications should be undertaken only after evolution in the targeted approaches has reached a critical level of stability, ensuring validity and endurance of the specifications over time.

Unlike the case of adaptation logic, adaptation actions constitute a well-researched and rather "crystallised" field, especially as far as Adaptive Hypermedia Learning Systems are concerned (Brusilovsky, 2001). Furthermore, recent research (Paramythis and Stephanidis, 2004) has proven the feasibility of formalising and declaratively specifying (using an XML-based language) adaptation actions to be effected as part of an adaptation cycle. It is argued that such efforts could easily be extended, so as to arrive at a standard that allows for flexibility as far as adaptation logic in concerned, and defines a concrete way for coupling that logic with an extensible set of adaptation performatives for ALEs.

Of the existing standards, the only one that supports the explicit representation of dynamic behaviour on behalf of the system is the IMS Learning Design (LD) specification. In more detail, Levels B and C of the specification under discussion introduce the concepts of properties, conditions and notifications, which can be used to specify arbitrarily complex dynamic behaviours for a system. The main setbacks in employing the IMS LD for modelling adaptation in ALEs are rooted in the fact that specification of dynamic behaviour is achieved through the definition of programming flows (including condition variables), enriched with event semantics:

- The approach can be considered rather low-level: Specifying complex adaptive behaviours is tedious and error-prone.
- Conditionals may only refer to variables or states that exist in the context of a

single IMS LD document (which makes it impossible to consult models external to the document).

- Dynamic behaviours cannot be defined at the system level (and applied in more than one contexts, or for more than one sets of learning materials / activities).
- The dynamic behaviour specified cannot be reused: there is tight coupling between the behaviour itself and the artifacts to which it refers.
- And, finally, the behaviour specification lacks semantic-level information which would allow an ALE to modify or affect it in any way.

Despite the above shortcomings, the IMS LD may be a very appropriate vehicle for introducing adaptive capabilities in non-adaptive e-Learning systems. Specifically, an adaptation engine can be introduced in an LD-capable system, which would effect adaptations by generating or augmenting LD specifications "on the fly". In other words, such an engine would translate adaptation logic and actions into IMS LD compliant constructs, which would then be delivered to the user. By going through this process dynamically (at run-time), the system would also be able to incorporate into the generated constructs, current information derived from adaptation-specific models.

3.4 Standardisation at the level of adaptation components and services

The majority of ALEs are designed to exist as stand alone systems. As a result, little or no attention is paid to exposing or utilizing adaptation-oriented components services to / from the "outside world". PLS (Conlan *et al.*, 2002b) and KnowledgeTree (Brusilovsky and Nijhaven, 2002) are two of the few examples of departure from this rule. Both of the aforementioned systems are designed to source content and functionality (such as Learner Management, Collaborative Tools, Testing Services) externally, not encapsulating all functionality into a monolithic core. However, they take quite different approaches to this sourcing.

PLS utilises a standards-based API (based on ADL SCORM 1.1) to interface with other compliant systems. In this way it can integrate with a Learner Management System (LMS) and pass user and assessment information back and forth between the systems. The PLS service is based on the notion that an adaptive

content provider should be a service provider rather than a repository for extraction of content. Communication between PLS and a learning portal (or LMS) is achieved by enhancing the SCORM Runtime Communication API as used in SCORM v1.1.

The KnowledgeTree framework, on the other hand, is intended to facilitate interoperability and reuse at the level of distributed, reusable learning activities (with the emphasis being on learning activities, as opposed to learning objects). KnowledgeTree, like PLS, goes into the realm of run-time communication and interoperability standards, seeking to standardize the ways in which different specialized subsystems supporting aspects of the (adaptive) learning process can communicate and exchange information that would allow them to be aggregated into a “whole”.

KnowledgeTree considers the standards-based model as not appropriate for adaptive distributed content and argues for a 3-component model (portal – content – student model server) (See Figure 2). PLS is structured to work within existing courseware management systems (CMS) that are completely static and thus consider adaptive services to be the main providers of adaptivity. It is assumed that the adaptive selection and structuring of content can only be done by a service. In contrast, KnowledgeTree allows for different kinds of portals – some can be as static as existing CMS, but some can be adaptive. In this vision, an adaptive portal can provide different adaptive support such as, for example, as adaptively selecting the best of existing static or adaptive content and adaptively arranging it for the student.

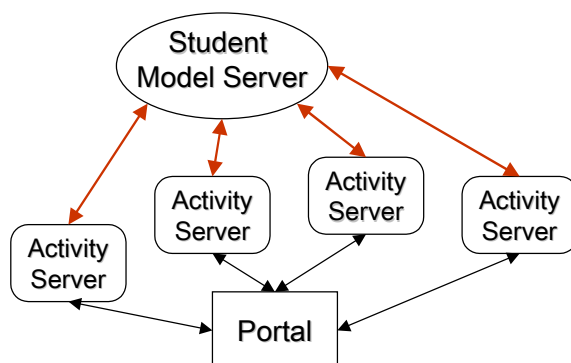


Figure 2: Main components of the KnowledgeTree distributed architecture – portals, activity servers and student model server. Adapted from (Brusilovsky and Nijhavan, 2002).

The two ALEs discussed above, despite taking two alternative routes to enabling service-, or component- based reuse of adaptation-oriented functionality and content, call our attention to important omissions in existing standards. Specifically, existing specifications necessarily refer to interchanges between components involved in the “traditional” interaction cycles between ALEs and their users. Adaptation, being outside this interaction cycle cannot be covered by these standards. Rather it introduces a new, distinct set of communication goals (or “reasons” for communication), as well as requirements in terms of the exchanges between system components. To compensate for this fact, we would need to engage in efforts to either: (a) enhance existing service-level specifications to explicitly account for the notion of adaptivity / personalization (PLS model), or (b) introduce new specifications that are “vertical” to, and independent from, existing ones, and are specifically intended to enable adaptation-oriented interchanges between those components / services that participate in the various phases of the adaptation cycle (KnowledgeTree approach).

4. Conclusions

This paper has attempted a preliminary assessment of the adequacy of existing e-Learning standards for supporting the introduction of adaptation techniques in e-Learning systems. The analysis, however cursory due to space limitations, has pointed out that existing standards do have some provisions for adaptation, but require substantial extensions to accommodate common practice in ALEs. Our findings can be summarised as follows:

In terms of domain modelling (i.e., modelling of courses, or learning-related activities), existing e-Learning standards do not suffice to capture the rich semantic structure that underlies static learning materials, or single- / multi- participant processes. Several alternative approaches are possible for integrating such semantic-level information into the metadata structures that typically accompany learning materials, as has been clearly demonstrated by recent research efforts. Of these approaches, the ones that seem to be most promising are those that seek to formalise / standardise ways for semantically articulating relationships and properties of the “units” from which materials and activities are composed; alternative adaptation methods and techniques can then be devised on the basis of the available information. More restricted in scope and

viability seem to be the approaches that attempt to integrate the adaptation “logic” itself with structural metadata (e.g., course metadata in the case of the KOD project). The main potential drawbacks of the latter approaches are: (a) only the logic that is already incorporated can be applied in adapting / personalising materials, and (b) the incorporated logic is of a procedural nature, and as such of no particular semantic value, making unfeasible its use as input towards the application of alternative adaptation methods / techniques.

The modelling of individual learners as users of an ALE is also not sufficiently covered by what is already available. The “PAPI” set of standards seems to be a step in the right direction, but is not adequate by itself. Not only that, but it is currently “competing” against a host of other related specifications, which may be more likely to adopt by organisations that seek to comply with a series of standards / specifications that come from the same body (e.g., IMS LIP is more likely to be supported by organisations complying with IMS specifications in general). ALE-oriented standardisation (not to mention “de-fragmentation” of learner profiling work) would greatly be facilitated by the convergence of these specifications, with (portions of) PAPI providing explicit support towards adaptation. The modelling of groups of learners, on the other hand, is still in an embryonic stage in terms of standardisation. Although there exist today ALEs that do identify and represent such groups on the basis of learners’ performance, preferences, interests, goals, etc., the representations used are intended only for internal “consumption”. Although PAPI can be used to convey information about groups of learners, it lacks the expressive capacity to capture the very information that would be of particular interest in ALEs: the common attributes / characteristics of learners (derived, e.g., from their history of interaction with the ALE) that have resulted in their being classified as a group. Specifications that address this shortcoming would have to be developed anew.

Adaptation modelling in the context of ALEs is a more complex issue, because it potentially involves both adaptation “logic” and the “actions” that result from the application of that logic in relation to the various static and dynamic models maintained by an ALE. It has been argued that seeking standardisation at the level of the “logic” might be futile at present, mainly due to the proliferation of

approaches that exist currently, and the vast differences between them (e.g., formalism used, computational models, etc.) A more realistic goal might be the standardisation of adaptation actions (i.e., the “things” that can be adapted within an ALE, and the “ways” in which they can be adapted). This would enable the employment of widely differing approaches to “logic”, while unifying, to some degree, the representation of modifications that can be adaptively effected in learning materials, computer support for inter-personal activities, etc. Furthermore, IMS LD has been discussed as a potential standard for instantiating adaptation logic and actions, at a given point in time and for a given user, or group of users. As in the case of group modelling, standardisation efforts in this direction would have little to build upon currently.

Standardisation at the level of adaptation components and services has only recently been addressed (at the level of *de facto* standards) in the context of research efforts. Two different approaches were presented and discussed in this paper: the extension of specifications that deal with the interchange between components of LMSs, and the introduction of new specifications that are explicitly intended for enabling the exchange of adaptation-oriented information / services between the major parts of an ALE. It is argued that these two approaches are not necessarily contradictory, or, for that matter, mutually exclusive. In fact, they seem to have complementary advantages and drawbacks: the former signifies an only partial departure from existing specifications, but fails when it comes to distinguishing between functionalities that can be exposed as services from specialised sub-systems; the latter, while addressing the problem just mentioned, must be approached very carefully in order to ensure that it does not interfere with, or render unusable, lower-level specifications that are already in wide use. Given the above, it is argued that a combination of the two approaches would be the better grounded alternative for any future endeavours in this direction.

In closing, we would like to touch upon a few topics that we feel are inevitably intertwined with any effort to expand upon current standards and specifications in the direction of adaptation / personalisation. To start with, it is argued that extensions to standards / specifications should happen in a way that keeps the “entry cost” of employing adaptation facilities in the development of e-Learning

materials, to as low levels as possible (mainly in terms of invested resources). An example of what would constitute, in the authors' opinion, a gradual and non-taxing path towards such employment, would be as follows. Authors should be able to provide an existing course with "traditional" metadata to an adaptive system, and get basic adaptation facilities (resulting from a "default" interpretation of the course structure and material by the system). Later on, authors could progressively add "adaptation metadata" as a stepwise approach to enabling / providing more advanced adaptation features.

Secondly, it is important that future extensibility of (new or enhanced) standards is seriously taken into consideration. It can be anticipated that the progressive uptake of adaptive methods and techniques in e-Learning systems will give rise to new adaptation patterns, and an even wider range of approaches than are in use today. Where possible, therefore, new standards / specifications should provide all the necessary extension points that would allow for the progressive enrichment of the respective models.

Finally, the adoption of the new standards or extensions proposed in this paper is, in our opinion, highly dependent upon the development of authoring tools that facilitate the creation of compliant resources. The creation of high quality-, standards compliant-learning material is already a quite demanding goal. The introduction of adaptation facilities will inevitably impose an additional "burden" on content creators. In order to bring the related cost / benefit ratio to non-prohibitive levels, it is necessary to have tools that: can assist authors in converting "static" material; support the authoring of adaptive content; enable the specification of adaptively supported activities in ALEs; etc.

References

- Brusilovsky, P. (1999) "Adaptive and Intelligent Technologies for Web-based Education", *Special Issue on Intelligent Systems and Teleteaching, Künstliche Intelligenz*, 4, pp19-25.
- Brusilovsky, P. (2001) "Adaptive hypermedia", *User Modeling and User Adapted Interaction*, 11(1/2), pp87-110.
- Brusilovsky, P. (2003). "Developing Adaptive Education Hypermedia Systems: From Design Models to Authoring Tools", in Murray, T., Blessing S., & Ainsworth, S. (eds.), *Authoring Tools for Advanced Learning Technologies*, Kluwer Academic Publishers, NL.
- Brusilovsky, P., Eklund, J., and Schwarz, E. (1998). "Web-based education for all: A tool for developing adaptive courseware", *Computer Networks and ISDN Systems* (Proceedings of Seventh International World Wide Web Conference, 14-18 April 1998), 30 (1-7), 291-300.
- Brusilovsky, P., and Nijhavan, H. (2002). "A Framework for Adaptive E-Learning Based on Distributed Re-usable Learning Activities" in *World Conference on E-Learning in Corp., Govt., Health., & Higher Ed.* 2002(1), 154-161, [online], AACE, <http://www.aace.org/dl/index.cfm/fuseaction/View/paperID/9357>
- Conlan, O., Dagger, D., and Wade, V. (2002a). "Towards a Standards-based Approach to e-Learning Personalization using Reusable Learning Objects" in *World Conference on E-Learning in Corp., Govt., Health, & Higher Ed.* 2002(1), 210-217, [online], AACE, <http://www.aace.org/dl/index.cfm/fuseaction/View/paperID/9365>
- Conlan, O., Wade, V., Bruen, C., and Gargan, M. (2002b). "Multi-model, metadata-driven approach to adaptive hypermedia services for personalized e-Learning" in De Bra, P., Brusilovsky, P. and Conejo, R. (eds.) *Proceedings of Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (AH'2002), Málaga, Spain (2002), 100-111.
- De Bra, P., Aerts, A., and Rousseau, B. (2002a). "Concept Relationship Types for AHA! 2.0", *Proceedings of World Conference on E-Learning in Corp., Govt., Health, & Higher Ed.* 2002(1), 1386-1389, [online], AACE, <http://www.aace.org/dl/index.cfm/fuseaction/View/paperID/9568>
- De Bra, P., Aerts, A., Smits, D., and Stash, N., (2002b). "AHA! Version 2.0, More Adaptation Flexibility for Authors", *Proceedings of the AACE ELearn'2002 conference*, October 2002, pp. 240-246.
- Dolog, P., and Nejd, W. (2003). "[Challenges and Benefits of the Semantic Web for User Modelling](#)" in *Proceedings of AH2003 workshop at 12th World Wide Web Conference*, Budapest, Hungary, May 2003, [online],

- <http://www.learninglab.de/~dolog/pub/semanticwebandum.pdf>
- Karagiannidis, C., Sampson, D., and Cardinali, F. (2001). "Integrating Adaptive Educational Content into Different Courses and Curricula", *Educational Technology & Society*, **4**(3).
- MAISE Centre (2002). "Making Sense of Learning Specifications & Standards: A Decision Maker's Guide to their Adoption", [online], e-Learning Consortium, Working Group to make Sense of our Standards and Specifications (S3), http://www.masie.com/standards/S3_Guide.pdf
- Paramythis, A., and Stephanidis, C. (2004) "A Generic Adaptation Framework for Hypermedia Systems", in Chen, S. Y., and Magoulas, G. D. (eds.) *Adaptable and Adaptive Hypermedia Systems*, Idea Group, Inc., forthcoming.
- Rodriguez, O., Chen, S., Shi, H., and Shang, Y. (2002). "Open Learning Objects: the case for inner metadata", *Proceedings of the 11th World Wide Web Conference (WWW2002)* (Alternate Track on Education), Honolulu, HI, May 2002, [online], <http://www2002.org/CDROM/alternate/693/>
- Vassileva, J., McCalla, G., and Greer J. (2003). "Multi-Agent Multi-User Modeling", *User Modeling and User-Adapted Interaction*, **13** (1-2), 179-210.
- Weber, G. and Brusilovsky, P. (2001). "ELM-ART: An adaptive versatile system for Web-based instruction" *International Journal of Artificial Intelligence in Education*, **12** (4), Special Issue on Adaptive and Intelligent Web-based Educational Systems, 351-384.
- Wiley, D. (2003). "Learning Objects: Difficulties and Opportunities", [online], http://wiley.ed.usu.edu/docs/lo_do.pdf (retrieved through associated discussion page: <http://www.reusability.org/blogs/david/archives/000066.html>)

