

Evolution of a Knowledge focused Computer Supported Learning System by Ensuring Extensibility through Generalization and Replication

Mr. David White
Dr. Lesley A. Gardner
Associate Professor Donald Sheridan

Department of MSIS, Auckland Business School, University of Auckland,
Private Bag 92019, Auckland, New Zealand.

d.white@auckland.ac.nz
l.gardner@auckland.ac.nz
d.sheridan@auckland.ac.nz

Abstract

If sufficient attention is not paid to the information models on which Learning Platforms are based the ability to deliver rich functionality is hindered. This paper describes the issues and records the experiences of the Cecil development in confronting this and the benefits that have been gained. This system, which has been evolving for 7 years, has been re-engineered to create the foundations for a knowledge engine, which embodies at its core the way we group and classify our knowledge of the world through generalisation and specialisation. These characteristics provide the foundation to deliver a wide variety of solutions in many domains apart from the initial design target of teaching and learning.

By implementing a singular and global identity of all knowledge elements and other information entities to allow logical extensibility the framework for physical extensibility, replication and peer-to-peer interaction, has been established.

1. Introduction

Over the 7 years that the Cecil Learning Management System has been in use at the University of Auckland, much has been learned about the requirements for such systems and thus the underlying principles of system design that

are required to keep extending functionality not envisaged in the early days of the design. These design principles which lead to extensible systems on both logical and physical dimensions, though applicable to the design of all information systems, are examined here within the context of learning management systems.

After a brief background to Cecil this paper focuses on two aspects of extensibility – logical and physical.

The significant lesson we have learnt in the area of logical extensibility has been the need to base the design on highly generalised information structures. This then allows the creation of new relationships between different aspects of the system as needed, without major re-coding being required. Thus by the use of strong object-oriented design techniques, in particular generalisation and specialisation, flexibility is built into the design.

This effect of logical generalisation is illustrated with the following examples.

1. The evolution of topic (knowledge) structures to allow the synchronous and asynchronous sharing of topics is examined and shown to be of fundamental importance in working with topics by individuals and groups as they teach, study and work together. This functionality has recently been refined by significant changes to the underlying data model to allow both synchronous and asynchronous sharing of the topic structures.

2. By addressing design issues at a fundamental level, the resultant system becomes available for much wider use than the original design intended. Thus the micro design expands and creates a more generalised tool. This has meant that Cecil is also used as a project workspace and as a personal knowledge manager.
3. Though the assessment engine was designed to be used for on-line quizzes, it has now had its use extended to criterion-referenced assessment [4] (rubrics), surveys, voting, and indeed to assess *anything* on the basis of criteria.

While the bulk of this paper is focussed on logical extensibility, the need for physical extensibility to support highly distributed disconnected devices (including PDAs, tablets, etc.) is imperative as the use of these devices becomes widespread. We show how by addressing the issues of logical extensibility, for instance the use of GUIDs, physical extensibility is facilitated.

The paper then discusses the upgrade issues that were addressed as these changes were applied to a system supporting continuous operation servicing 20,000 logins per day. As well as issues with sustaining the service during the change it was important that the users were able to see any changes as an evolution and that the migration caused little disruption to their established workflows. This was achieved by, as much as possible, maintaining the previous user interfaces in parallel with the new.

Related work in this field is addressed by considering how the work to transform Cecil opens up possibilities for solving problems such as Management of Ontology evolution and as a Social Support Decision Support [6] system for Discussion and Voting on Topics [8].

2. Background

This section briefly describes the background of Cecil development and use.

Cecil is a computer supported learning system; this means it provides computer services to assist instructors in organising courses and providing resources to students over the Internet. Cecil provides services in four main areas: teaching administration; online course material; web-based testing; instructor-student and student-student communication [1].

In response to movement in education to modes of distance and technology-based learning, the university began to build

significant, innovative, and tangible projects. Cecil is one of these projects and represents a translation of a vision into a new, useful and extensible tool. It takes advantage of Internet technologies to enable onsite and remote access while leaving the design open for enhancements as Internet bandwidth allows.

The Cecil architecture is classic n-tier consisting of a database (SQL Server 2000), an object layer (COM), a web interface (Active Server Pages), and Windows thick clients delivered via Citrix. The object layer contains the business rules and is used by all presentation layers. The system runs on clusters of servers to provide scalability and high availability at an economical cost. The current implementation is based on a central server but the inherent design and technology that underpins the Cecil engine allows for use in a peer-to-peer environment.

Cecil has had at its core, from the outset, the ability to share topic (knowledge) structures between courses and faculty facilitating the sharing and reuse of learning objects. These topic (note that the term topic is used as a user friendly term for knowledge element) structures are the "containers" for all instructional resources including questions, tests, assessments, discussions and all electronic media and are themselves contained in or referenced by learning activities. This offers the teachers and course designers the possibility of creating strongly and coherently structured modules of learning which can be re-used from group to group over time.

The University of Auckland is the largest university in New Zealand and Cecil currently services 20,000 logins per day - an increase of 5,000 from the previous semester.

3. Generalisation as the Key to Logically Extensible Systems

When the initial design of the data structures of Cecil was prepared in 1995 based on previous models of teaching and learning, the designer proposed (to himself) the creation of a highly generalised structure. Through fear of the unknown and the pragmatism of working within the constraints of existing hardware and software platforms, the design implemented was less generalized than the inherent nature of the world requires. As the development work of Cecil has proceeded over the years this decision has meant that creation of new functionality has been hindered by this lack of generality.

This can be illustrated by considering that if the identity of a person and a topic do not inherit from a generalised entity called, say, *Anything*, then when we come to create a discussion about a person, a topic, an activity, a learning object we must associate that discussion with two or more independent structures making the implementation of shared functionality more difficult. This same argument can be applied to relationships between questions, assessments, activities and all other objects in the domain.

Thus as the work on the extension of the topic tree to a topic network proceeded so too did the generalisation of other elements. This was known to the development team as the “*Anythingisation*” of Cecil. During the southern summer of 2001-2 Cecil these changes were made.

4. Evolution of Knowledge Structures

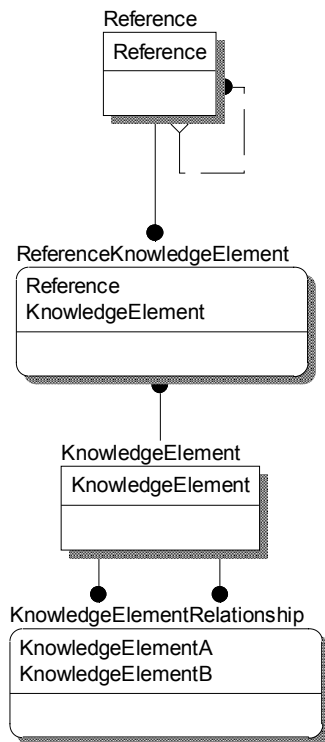


Figure 1. Original design of the Reference Tree and Body of Knowledge Network

The original design had two structures (Figure 1). The first was a network of knowledge elements and the second a tree structure called references (later named topics to clarify user understanding). The term *references* was intended to describe books, papers, Web Sites, Videos and all other manifestations of what was thought of as an abstract Body of Knowledge

(BOK). This distinction between references and topics was made on the mistaken belief that a definitive Body of Knowledge could be defined for a field rather than being just another assembly of topics. Due to the resource constraints of system construction the BOK network was never included in the user interface and as time went by the “references” structure was renamed, by user demand, to the “topic” structure. This was used for working with such BOK’s as the IS’97 [10] definition of the BOK of Information Systems, which was imported to experiment with creating a system for curriculum and course design based on the BOK.




Figure 2. Topic Network

This and other work led to the realization that differentiating between the two structures was artificial and that what was required was to extend the topic tree to a topic network (Figure 2).

This structure of relationship between topics was then extended back to the most generalized element in Cecil *Anything*. This then allows any element defined in Cecil to be related to any other element. This provides extensibility and flexibility for the ongoing evolution of the design.

5. Synchronous and Asynchronous Sharing of Topic Structures

In the previous version of Cecil the topic tree could be shared between multiple courses allowing the re-use and shared use of material between courses. This was totally synchronous so that if a topic tree was shared then changes in one context affected views in all contexts. This is difficult to manage between a few closely cooperating individuals but is totally unmanageable for larger groups. This is especially true when a user, considering that their view is totally independent within a course context, makes a change that has the effect of modifying, without consultation, a topic schema (Figure 3.) in another context.

The experience of synchronous topic sharing led us to design the topic structures so that collaboration was possible as well within private,

group or public structures similar to a file system. Thus the design now supports both synchronous and asynchronous sharing of topic structures. Users are able to work privately within a public or shared group structure and then to publish to others who may or may not subscribe to the publication. This now allows Cecil to support the collaborative growth of bodies of knowledge as extensions and changes to a structure can be discussed using discussions that are themselves contained within the topics. These modifications may then be voted on using criterion referenced assessment mechanisms (tests), which are also topic focused. This ability to associate any elements with any others is made possible as they are all specialised from a single class - Anything.

6. How a Generalized Design Leads to Unintended Uses

As Cecil has evolved other uses for it have been discovered. It has for instance a simple but effective document management component thus departments and groups of researchers within the university are using it for this purpose. The test/question structure is used as a survey tool, for criterion-referenced assessment [4] and as a human resource assessor.

Reflection on this has led to the realization that a course is simply a specialization of a project composed of resources and activities, which focus on and are characterized by a variety of topics. In the case of a course the activities will be lectures, tests, laboratories etc. for another type of project they will be related to that domain. The terms course/project can be further generalized to a *space* or other term, which is applicable in the domain of focus.

6.1 Project Work Space

Following the lead of some Cecil users, the Cecil design and development team now use a number of *spaces* (a generalization of courses) as their virtual office. Whereas previously discussions would occur using email they are now held within the discussions spaces, which are linked to the various topic areas of the project. These topic areas also contain all the electronic documentation of the project that was previously managed by a software source code management system. This has had the profound affect of ensuring that the software developers are using their own work. This has increased

their sensitivity to the needs of the user community of which they are now part. This changed behavior is leading to more focused design work. As well the Cecil users have also been provided with a discussion space into which they may participate in the evolution of the system they are using.

6.2 Problem Resolution Work Space

A further development in this area has come from the realisation that the logging and prioritizing of bugs was in fact a specialization of a discussion. Work is underway to extend the information model to incorporate this useful functionality that can also be used to manage any problem resolution process where discussion and priority setting by groups is required.

6.3 Tests and Questions as Survey and Voting Tools

The powerful assessment engine in Cecil has been extended to create a tool for surveying. The university administration has made use of this for surveying the needs of the 5,000 international students who attend the University of Auckland and to gather the views of the entire staff/student body of 30,000 people in an attempt to address parking issues at the campus which is situated in the city centre. Other researchers, both staff and students, are also using it in their research work.

6.4 Generalized Criterion Referenced Assessment

The use of the Cecil assessment engine for criterion-referenced assessment (rubrics) focused on assessing learning was addressed in "Criterion Based Assessment using the Support of a Computer" [4]. Since that time it has been recognized that the same functionality can be applied for many other different assessments against criteria.

The first project in this area is the re-skinning (new user interface) of the Cecil engine to deliver a system for the assessment of human resources. This work is being done by an independent software development company (RunTheOtherWay) to deliver assessment capability combined with a project workspace, discussion and knowledge/document management to a corporate customer.

Other examples of where this same approach may be applied are the assessment of buildings

for fire risk, flooding potential and evaluation of security risks.

6.5 Personal Knowledge Manager

Knowledge is something that we gather and accumulate within ourselves throughout our lives. The outer manifestations of this knowledge can be seen in the books and other media that surround us in our homes and offices. The tracks of our lives are also seen in the images, audio and video recordings of people, places and things which have been significant for us.

We collect memorabilia including cuttings, writings, drawings etc. related to our world. Other of our personal knowledge resources are scattered over hard drives and within email systems. As Cecil continues its metamorphosis into a peer-to-peer knowledge focused workspace a user interface is being designed to allow people to manage their knowledge and the resources that inform and illustrate their track through the world. Inherent in this work is the ability to share, merge and replicate information with other instances of Cecil, thus creating spaces that range from public to private.

The user interface is being designed as highly graphical and object focused. By having a common identity (GUID) and generalised object structure from which every type of object in the system is specialized makes the design and implementation of the user interface simpler.

By extending and refining the user interface of Cecil and the inherent generalizations a project is underway to create a personal knowledge manager. It is the intention of this project to create a user interface that makes use of multiple dimensions and is usable by small children.

7. Replication as the Key to Physically Extensible Systems

Synchronization across multiple devices is now the norm with systems such as Microsoft Outlook/Exchange supporting replication across PDAs, personal computers and servers. Most LMSs, including Cecil, are based on a centralized server with the users connecting for sessions.

Replication addresses the need for teaching and learning support systems to manage the high degree of disconnectedness that is necessary for autonomous mobile users or in parts of the world where connection to the "net" is expensive

and/or unreliable [7]. This disconnectedness is also significant in situations where users wish to directly, using say infrared or radio, interchange materials.

It also allows for automated updating of a local data store without user intervention removing the need for the user to "poll" the system for changes. For instance at present users must, on a regular basis, log in to see if there are new resources they may wish to download. This is an unnecessary action as the user's own instance of the Cecil engine will be able to automatically check other peers and synchronize, based on rules governing optimum times and constraints on transfer sizes etc., without the need for human intervention.

While the current implementation of Cecil is based on a central server the inherent design and database technology which under-pins the Cecil engine allows for use in a peer-to-peer environment. Thus as the work was done to increase the use of generalization and inheritance structures to allow more rapid extension of functionality the identities of the system elements were converted to GUID's (Globally Unique Identifiers) which then facilitate replication between peers of server farms, personal computers and PDAs.

8. Related Work which is Informing Future Design

8.1 An Environment for Voting and Discussion on Topics

Because of the tight inter-linking between voting (using the assessment engine as modified for criterion-referenced assessment), discussion and topics, the Cecil engine could be used as support a "Social Decision Support Systems" in the context of a participatory democracy. [6].

8.2 Management of Ontology Evolution

As both the discussion and test mechanisms are focused on knowledge elements/topics and all information entities are derived from a common class - *Anything*. Cecil delivers on the design specification called for by Tuoff [8] to enable the "Knowledge Sharing and Learning in Virtual Professional Communities".

The information structures described in this paper are parallel to "... mapping all objects within any C-KLASS system as a subclass of a

node or mapping each to a node, which then all tools can operate upon. Essentially a node is an element of interest” [8]

For the collaboration and mediation on the development of bodies of knowledge within communities, algorithms could be developed which would then be used to decide on the basis of voting and participation in discussions etc which topics would become part of an accepted group body of knowledge.

While various processes to work collaboratively on ontology development have been used, including the Delphic method [9] the availability of tools and environments to augment these processes are at present not integrated. Practitioners make use of word processors, email, questionnaires and specialized databases. With extension to the Cecil user interface and security model, under construction at present, a tool to support this activity will be available.

The need for such an environment is called for in the emerging discipline of Information Systems where there is a shift from the IS discipline being one referencing other disciplines to one which is referenced by many others. [5]

8.3 Enriched Interfaces for Discussions

The design of the Textweaver [7] system has much in common with the design goals of Cecil with respect to the recognition of the need for the use in a disconnected environment and also in the manner in which elements of discussion can be “woven” within each other and with objects that are externally referenced. The scope of Cecil is however beyond the discussion element of Textweaver. The work done on ensuring generalisability ensures that the functionality called for in Textweaver is catered for in the Cecil design.

8.4 Groove

Groove [11] is the system which has been developed by the designer of Lotus Notes, Ray Ozzie, and that provides an excellent example of the peer-to-peer collaboration model towards which the development described here is taking Cecil. The Groove engine does not currently provide the knowledge management aspects of the Cecil topic structures or the ability to make assessments and hold discussions focused on any object of interest.

9. Deployment of the Changes

The changes described above affected the core of a system that is in use 24/7/365 supporting 20,000 user logins per day, producing 1.5 million hits per day and so had huge inherent risk. The final impetus for the change came from the need to create a powerful collaboration environment beginning with discussion and chat. Since the inception of Cecil, based on the principle of buy before build, 3rd party “discussion and chat” products, including Web Crossing and Web Boards, were integrated into Cecil at different times. These proved lacking in functionality, reliability and scalability delivering a mediocre service to the user community.

Early in the design of the collaborative environment the need for a highly generalised environment was recognised. After all, one can have a discussion about *anything*. Thus it should be possible to discuss a person, a topic, an activity, an assessment, a document etc. This requirement gave the impetus to make the changes required.

A decision was made to use GUID’s (Globally Unique IDentifiers) with apprehension over the effect these large identifiers might have on performance of the database because of index width. The choice of GUID’s over integers was made to ensure that an original design goal of synchronisation/replication between peer-to-peer Cecil instances could be achieved.

Through the southern summer of 2001-2 this transformation was undertaken and the system went live successfully for the first semester of 2002. It is performing well both functionally and in supporting a rapidly increasing workload.

10. Conclusions

This paper has described the work undertaken to explore and create, in an operational setting, a knowledge, activity, assessment and learning management system. The effect of constructing such a system based on strongly extensible information structures has been proven to be highly beneficial in delivering evolving requirements.

The paper has shown the evolution of the underlying information structures of Cecil to provide the basis for both logical and physical extensibility. It has shown the serendipitous uptake of Cecil as a survey tool and as a collaborative group-working environment. It is

clear that the underlying model of Cecil, its construction and delivery mechanisms allow a robust multi-faceted tool for assessment, facilitation of materials delivery, knowledge management and group work within a peer-to-peer environment.

The ability to deploy new and highly integrated features has been shown to be greatly enhanced by the work described. This has also led to the application of this design into domains whose functionality overlaps that of teaching and learning.

Now that the internal data design of Cecil has been modified to allow the representation of all objects, within the domains of teaching, learning and knowledge, as a network of generalised nodes, the work of realising extensibility on many dimensions is possible and has begun.

This work is proceeding to the point where *any* of the elements of this network of topics, discussions, activities, learning objects, people etc. (indeed *anything*) can be assessed from *any* viewpoint and discussed within *any* context by *anyone* at *anytime anywhere anyhow* as the *what, where, when, why* and *who* of the universe is explored.

11. References

- [1] Gardner, L., Sheridan, D., and White, D., (2001), Computer Supported Learning. A Large-Scale, Web-based Learning and assessment System to Support Flexible Education, Proceedings of the 34th Hawaii International Conference on System Sciences,. IEEE. 2001 <http://cecil.edu/usage/>
- [2] Sheridan D., (1999), Computer-Supported Learning: A New Zealand Perspective, Tan, F.B, Corbett, P.S. and Wong, Y.Y. (eds.) Information Technology Diffusion in the Asia Pacific: Perspectives on Policy, Electronic Commerce and Education. Hershey, PA: Idea Group Publishing.
- [3] Sheridan, D. (1998), Using the Internet to Support Learning: Implications for Life Long Learning and Self-Assessment. ACCESS: Critical Perspectives on Cultural and Policy Studies in Education, 17(1).
- [4] White, D (2001) Criterion Based Assessment using the Support of a Computer, Proceedings of the 35th Hawaii International Conference on System Sciences,. IEEE. 2002
- [5] Baskerville, R. and Myers, M. (2002) *Information Systems as a Reference Discipline*. MIS Quarterly, 26(1), Carlson School of Management
- [6] Turoff, M., Hiltz S. R., Cho H., Li, Z., and Wang Y. (2002) Social Decision Support Systems (SDSS) Proceedings of the 35th Hawaii International Conference on System Sciences,. IEEE. 2002
- [7] Xin, C., Feenberg A., (2001) Designing for Pedagogical Effectiveness: the TextWeaver, Proceedings of the 35th Hawaii International Conference on System Sciences,. IEEE. 2002
- [8] Bieber M. et al (2001) Towards Knowledge-Sharing and Learning in Virtual Professional Communities Proceedings of the 35th Hawaii International Conference on System Sciences,. IEEE. 2002
- [9] Gruninger M. and Lee, J. 2002 Ontology Applications and Design Communications of the ACM Feb 2002/Vol. 45 No. 2
- [10] Davis, G. B., Gorgone, J.T., Couger, J.D., Feinstein, D. L., Longenecker, H.E., (1997) IS'97 Model Curriculum and Guidelines for Undergraduate Degree Programs in Information Systems, AITP
- [11] Groove, Stanhope, P., Get in the Groove: Building Tools and Peer-to-Peer Solutions with the Groove Platform, John Wiley & Sons; ISBN: 076454893X; 1st edition (May 10, 2002)