Innovative use of telematic tools to support a professional community of practice

Mike Malloch, KnowNet,
mike@theknownet.com,

Graham Attwell, KnowNet and Pontydysgu,
graham@theknownet.com

For presentation at Online EDUCA Berlin, 2001
November 29, 2001

In this paper, we describe a group of projects developing web-delivered tools and environments to support the development and sharing of knowledge by professional communities of practice. These projects are innovative and interesting in three areas: software innovations, theoretical intent, and tight, responsive coupling of social research with software design and real-world usability.

We begin by discussing a larger programme of development and experimentation, of which these projects are part. We explain what is unusual about our approach and where it comes from, and list the key ideas and software architectures we are striving to develop. An important part of that discussion is a delineation of some fundamental limitations of existing knowledge management and collaborative working software. Limitations of social and usability research into online collaboration and learning are also mentioned, in a brief discussion of our own approach to unifying social research and software development.

A description of three of our current projects is then presented. After introducing each project, we discuss some of the innovative developments already being used, and our plans for the next year of these projects. Later papers will explore our experiences using these tools to enhance the emergence of communities of shared knowledge and discourse.

We conclude with a summary of our key aims, ideas and implementations, and call on like-minded developers, researchers and user communities to get involved.

Background – KnowNet's long-term mission and experiment

KnowNet is a small company entirely devoted to the long-term development of fundamental architectural innovations for internet-mediated collaborative knowledge development and learning. It was formed in 2000 by researchers and software developers with long experience of trying to develop useful collaborative environments in existing software frameworks. After many disappointments with existing efforts, the founders of KnowNet determined that fundamentally new ideas, architectures and tools were going to be needed, but that such basic and speculative development work was too ambitious for publicly funded projects but too innovative for large commercial organisations.

KnowNet was formed to act as an autonomous, self-funding organisation dedicated to advancing that mission.

There are several ways in which the development and research we undertake is notably different from others we have encountered

What is unusual about our approach to collaborative experiments and software development?

There are several genuine innovations in our overall approach to both software development for, and research into, collaborative work and learning environments.

Deeply informed mission to address deep issues

Our key distinguishing feature is our total commitment to a long-term aim: creating fundamental architectures for online collaboration in the process of discovering what shape those architectures and implementations should take.

Many developers and researchers have long known that better architectures are necessary and achievable – that it is technologically possible to implement a much more 'semantic' web of discourse and knowledge to overlay the current web of locations, pages and messages. But no-one can predict in advance what that web will look like since (a) this is a social, psychological and historical question, not a technological one, and (b) it entails a 'knowledge economy of scale', and we are very far removed from either large-scale uptake of innovative architectures, or even exhaustive efforts to explore the space of technological/social innovation towards it. Previous radical innovations in users' practice with computer and internet technology make it abundantly clear that entire working practices can change, that users cannot imagine working without technologies that they could not have imagined using only years before.

It will be very difficult, then, to design internet architectures for collaborative knowledge development. If it were easy to see what was required, and what markets might emerge, large organisations would presumably have done so by now. Nor does it seem likely that open-source movements or governmental initiatives can sustain the combination of vision, thoroughness and flexibility needed to achieve this goal. Based on our own experience, we strongly feel that compact, committed and talented entities are required.

Why do we believe that KnowNet can make real advances towards this difficult goal? We outline below some of our own experience.

Real experience of ambitious development projects

The KnowNet team includes talented developers who have 'been there', striving hard to develop useful and innovative solutions for online collaboration in the context of publicly funded projects. This involved both ambitious and radical software development, attempts to leverage or systems-engineer useful solutions from existing tools, and sustained efforts to understand user communities, and to animate and engage them in both types of solution.

Many others have worked hard on similar projects, and know too well what an arduous exercise that is, and how few the tangible outcomes. Integration and interaction is too

shallow for users to feel that they are engaging in a thorough and powerful environment of resources, ideas, issues, tasks, actions, people and discourses. Our experience strongly indicates that thorough, useful and satisfying online knowledge development is possible, but only if some developers can concentrate on the creation of basic architectures and tools to underlie the creation of powerful unified collaboration / content environments.

Experience of the standards-development process

Developers and educators have sometimes attempted to pool resources in order to facilitate the development of powerful architectures for online collaboration and learning. One of our own developers was deeply involved in two such efforts: the IMS Project and the IEEE Learning Technologies Standards Committee.

Joint efforts, combined with the development of (abstract) interoperability architectures, seemed to us and many others to be the most rational way out of the 'tragedy of the commons' – the situation developers found themselves in, where all knew that fundamental architectural frameworks needed to be in place to enable anyone to make progress on their specific solutions, but no-one had a direct interest in producing those frameworks.

But, some years later, these frameworks are still needed, and still not in sight of being developed. Why? Our experience indicates that there are three main handicaps of the traditional standards-development process in addressing the problems on online collaboration and learning:

1 – No examples of non-interoperating solutions. Specification of interoperability architectures needs mature examples of the products and systems that must interoperate. But no-one really knows what kind of systems might emerge once interoperability standards unleash the creativity of ingenious small developers, while everyone expects that functionality will explode far beyond that of existing systems.

2 – Lack of a framework for research. The level of integration and extent of variety-of-services that would characterise powerful collaboration environments are both high, and in a way their benefits are 'all or nothing'. Likewise, it is hard to see the architectural implications until there is an experimental platform on which to do research into users' actions, expectations and needs. Worst, if an experimental product's key functionalities only become evident to users in an environment offering many other products or services, it becomes prohibitively expensive to test experimental software ideas.

3 – No model for prototype implementation. It is a mistake to expect commercial software vendors to create products exhibiting the elegantly interoperating functionalities of the 'blue sky' collaboration and learning environments we desire. Nor have standards development bodies seen it as their role to implement prototypes, test environments or development frameworks, while ad-hoc collaborative development efforts are handicapped by participants' other responsibilities.

Our experience with the standards movements strongly motivates us to a different approach: instead of designing standards and architectures in the abstract for systems no-one has any experience with, we have chosen to work hard to erect (at least) one exemplar platform for experimenting with deeply integrated new collaborative

functionalities, and to design, experiment and reflect in a tight cycle of social research, architectural design and practical deployment. Our intention is to take any insights or architectures gained to standards organisations when experience allows informed discourse on the issues involved.

Intense, informed user experience

Other members of the KnowNet team, while having no software expertise, have been intense users and systems-integrators of existing collaboration and knowledge management environments. Their experience has mirrored that described above: frustration with the shallowness of current environments alongside a strong sense that much more useful technologies could be deployed.

It is interesting to contrast the frustration of trying to use existing systems 'in anger' with the apparent usefulness of these systems in 'demo mode'. Scale of use is one important issue; email and conventional groupware systems become very labour intensive and awkward when traffic is heavy, resources and conversations become lost in traffic. Thoroughness of integration is another; users become frustrated at having to re-express ideas and textually establish relationships when it feels as if 'pointing at' and gesturing would be more appropriate. Another crucial factor is quality of interface – almost all experimental software projects fail to develop thorough interfaces with the 'fit and finish' required to sustain the illusion of ubiquitous action-affordance; real users almost always give up before inducing the opaque user actions which seemed natural to those who implemented them.

Our experience as users of groupware, knowledge management, content management and online learning environments again leads us to believe that another approach is needed, one which includes – among other things – careful attention to learning from real-world use while guiding users into new functionalities with thorough implementations and active participation in user communities.

Theoretical foundation and outlook

KnowNet's founders are, by their natures as well as their academic training, philosophically inclined. Our experiences, described above, have always been coloured by (and often informed by) philosophical and theoretical concerns.

Among the published works which have informed our own thinking are the work of Winograd and Flores on software design and electronic discourse, Brown and Duigid on communities of practice and document ecologies, and socio-cultural activity theory (Engestr m and Cole, 1993). Kapetelinen and Nardi (1997) have produced guidance that is being incorporated into the design approach.

Aside from official theorizing, we feel it is vital for developers and researchers to fully understand both each others goals and methods, and the larger issues entailed in the social life of collaboration, of software design, and of their relationship.

Innovative ideas

At KnowNet's founding, its principals carefully and deeply thought through how a small company could make an impact on this very large problem, and also on what key ideas could motivate some breakthrough applications. In a later section, we discuss some of the ideas we came up with, and why we think they could lead to interesting and truly innovative application spaces.

Mission-centred approach

A final – but crucial – lesson we have extracted from our experience in the field is what we call our 'mission-centred' approach. While we are utterly committed to giving our clients the best possible systems using technologies already developed, and to refining the usability, robustness and integration of our existing tools, our mission is to constantly piggyback new generations of exploratory functionality on top of them, and to continue this mode of development indefinitely. Our goal is to understand what communities of practice can do with electronic technologies, and to help fundamentally re-invent those technologies so that communities can collaborate effectively in terms of the meanings, actions, objects and discourses natural to them, and so that learners can learn by active engagement and creation in a web of ubiquitous meaning and conversation.

Tight loop between theory, research, real-world use, architecture and software design

One comment we often hear from members of the communities we provide solutions to is their amazement at our responsiveness – in new or changed functionality - to their experiences, suggestions and usability problems, and our informed concern with understanding the way they do their work. This is a key principle. For the reasons outlined above, we feel that *no-one* understands what online technologies and communities of practice can do; *no-one* has a clear vision of the kind of software architectures and interfaces required, or of the way people can work and learn in webs of knowledge. Our understanding can only improve by *seriously* trying out new kinds of architecture and interface in a fully integrated environment, and by carefully researching the processes involved as we develop and deploy software and users try to accomplish their work. Our engineering can only learn to address the right problems if we couple its development very tightly, and very responsively, with careful real-world usage experiments.

This cannot be like traditional software design methodologies, since the fundamental specification is exactly what we want as an outcome, not a starting point. Nor can traditional usability testing be appropriate, since our users have never before been able to do what we want them to be able to do. User communities, research work and software engineering must all learn together.

We pride ourselves on the unusual asset that our architects, theorists, implementers, animateurs and researchers all deeply understand this total cycle.

Sustainable, thorough, quality implementation

It is unusual for research projects to create beautiful, usable, thoroughly implemented software. In fact it is almost unheard of, and that is one reason we feel it necessary to work outside the framework of academic research.

It is absolutely essential that the experiments in functionality that we undertake do not forsake quality of implementation in the quest for innovation. Users simply will not feel comfortable enough to fully engage in virtual communities and knowledge networks fronted by half-implemented interfaces or backed by half-realised architectures.

We strive to create highly usable, robust and beautiful interfaces, and constantly refine our designs in searching for seamlessly intuitive combinations of natural action-affordances and rich content. We think we have already constructed some best-in-class interfaces, for instance our very thoroughly realised virtual resource folders and interactive xml-document viewer. For some of our other interfaces, for instance our powerful tools for sharing arbitrary structured metadatas, or for authoring semantically structured XML documents, we have yet to find the natural way to engage users, but are working hard to do so. In the web of ubiquitous knowledge and discourse we dream of, the interfaces must be intuitive, immediate and powerful, sustaining an illusion that a universe of knowledge and actions on it is present-to-hand as users work. We are committed to constantly improving our interfaces and deploying highly professional experiences for our users.

This is an extremely labour-intensive process, though - especially when deploying through the web. The key question for any developer wanting to explore next-generation collaboration environments is how this expensive process can be paid for. We list, below, some of the reasons why we believe the most effective model is that of small, determined organisations like KnowNet, drawing income from practical interim knowledge-development projects.

Not a vendor's restricted vision

Why haven't commercial software vendors built the architectures we need? It seems clear that, once such architectures and associated software become available, they will provide communities of users with a host of new abilities, and that must mean a large market, eventually. But our own bitter experience convinces us that – at least until the possibilities are better defined – commercial vendors will not address the necessary foundation issues. Large vendors do not need to; they are confident of being able to capture a large share of any such market by deploying their own speedily developed solutions long after all the exploratory work has been done (the history of web technologies being an obvious case in point). Smaller vendors simply cannot afford to risk development effort to explore new functionalities that involve so many architectural dependencies (especially given the risk of larger vendors taking all the market afterwards anyway).

Not a project's restricted scope

A great deal of public funding is spent to address precisely the kind of 'blue sky' development work we have described. Many hundreds of projects with closely related

agendas have run their course, and many hundreds of others are current or about to be funded. Why hasn't this public expenditure produced more results?

Our experience is that publicly funded projects are too small, too insular, too limited in scope, too accountable on small-scale goals and too little rewarded on more fundamental ones. Worse, the dynamics of funding and accountability make it very rare for projects to build effectively on previous work or collaborate effectively with similar projects – much effort is wasted constantly re-inventing the same wheels to produce the same half-finished outcomes.

Not one institution's restricted resources and perspectives

There are numerous large but non-commercial institutions which one might feel have an interest in undertaking or underwriting the developments we seek. Large universities or governments, for example, have much to gain from more powerful collaboration environments.

Our experience leads us to doubt that such organisations can mobilise the necessary resources, or effect the necessary in-house culture changes, to be effective in fundamental architectural / experimental work like this. More likely is a slightly larger-scale repetition of the wheel-reinventing and half-finished solutions that plague small projects.

It must be pointed out that, in this context, 'half-solutions' do not 'half-work'. As solutions they invariably prove almost useless, though as experiments they could be very valuable if it were not for lack of information about the processes and obstacles leading to the 'failure'. It must also be mentioned that, in the organisational and project-management cultures likely to apply, this potentially useful information is usually thrown away in an effort to meet accountability standards.

Not a standards organization's restricted mission

Finally, we re-emphasise that standards bodies are very unlikely to muster the concerted effort and purpose required to develop real software, though several crucial roles must and will be played by standards bodies in the long-term process of evolving the ubiquity and interoperability required of truly usable solutions.

Nor can we, as yet, point to the open source movement as a likely development base. Open-source movements are very good at refining and reverse-engineering existing kinds of software, but are very, very bad at fundamental innovation.

We have described, above, the reasons for KnowNet's coming into being, the reasons we feel that we have unusual assets to offer, and the reason we feel that small companies like ours will lead the way in developing ground-breaking architectures for collaborative knowledge development. We have not yet given any detailed description of why we feel that existing collaboration environments are fundamentally unable to make breakthroughs in uptake or usage by dispersed communities of practice or online learners. We address that question in the next section.

What is wrong with existing collaboration software?

Before delineating the limitations of existing software, let us praise what is good about it.

What is *right* about existing collaboration tools

It is evident that existing tools for online working and knowledge sharing have made a wide impact, if not a deep one. We divide our discussion of these successes into two categories:

Ubiquitous, standards-based tools

Almost everyone in the industrialised world now has some access to tools and content-delivery environments based on industry standards: email and web pages are familiar to everyone, and have changed many people's working practices.

The great merit of these tools and environments is that they are based on thorough suites of interoperability standards, and so their scope reaches far beyond individuals' organisations to allow the spontaneous emergence of virtual communities of interest and, to some extent, practice.

To a lesser extent, yet to be standardised but widely available, peer-to-peer instant messaging, chat, video conferencing, internet telephony and shared whiteboard applications are used by many people to enhance synchronous communication at a distance.

Some new standards, like XML, are beginning to enable solutions like shared newsfeeds and interchangeable representations for web documents.

Groupware and knowledge management tools

For more than a decade, special commercial environments have existed to aid people in organisations to communicate and to organise knowledge resources.

Groupware is a term used to describe products, like Centrinity's FirstClass and Lotus Notes, which provide a suite of collaboration tools like shared conferences (shared email folders), calendars, resource spaces and chat rooms.

Non-commercial products also exist for shared bulletin boards and topic-centred discussion arenas, but these are less functional, with less powerful administration, customisation and moderation tools than the commercial products. A standards-development or open-source movement to create free standard environments and tools for groupware would be a very useful step forward.

Knowledge Management is a more recent term applied to environments in which companies can organise corporate document repositories and present a front end combining access to these with groupware functionalities, often also allowing users to customise their view of these resources.

Content Management is a related term applied to tools for expressing, archiving, versioning, managing and delivering large repositories of documents.

These environments provide a significant step forward from 'standard' email and web pages, mostly by organising messages and content so that users can easily access shared versions when they need them. Standard email does not allow shared copies of messages, and forces each user to undertake the task of organisation independently. Standard web

pages make very poor document repositories, since they are designed for appearance rather than content, and are hard to edit or version-manage.

We ourselves are avid users of such systems. We would not deny their usefulness for certain purposes if carefully maintained and moderated, and all our portal solutions include the functionalities for groupware and knowledge management mentioned above.

In addition to groupware and knowledge management environments, there has recently been a proliferation of 'virtual learning environments' (there are many similar terms, such as 'managed learning environment' and 'e-university'). These are, essentially, a combination of some groupware and some knowledge management functionality, generally not thoroughly implemented, along with some basic tools for generating course notes. Our experience indicates that a great deal of human intervention is required when employing these tools in educational settings, and that very little value is added to the learning process regardless of the effort expended; what is useful in disciplined work environments seems to be much less so in learning.

Why is online collaboration and learning still so sparse?

In spite of the universal availability of email and web pages, and even with the careful deployment of groupware, knowledge management and 'virtual learning' environments, communities of users do not enjoy the rich interactive experience we feel is essential before the internet can become a medium for knowledge work and learning.

We briefly list below some of the problems as we see them.

The 'place' metaphor is weak

Whether exhibited in 'site', 'location' or 'discussion space', the overwhelmingly dominant metaphor for engaging with web content and messaging is place: 'Where' do you want to go today? 'Where' should I send this report? 'Where' should I look for this piece of information? 'Where' did I put that message about last month's meeting?

Note that, as natural as this sounds when thinking about the physical activities associated with knowledge work and learning, it is in fact profoundly unlike the terms that naturally describe the activities of knowledge work and learning themselves.

What, why, who, how… these are more natural terms. In fact, when analysed closely, we lack explicit terms for most of the cognitive and inter-subjective motifs of topic-setting and shifting, contextualising, assertion, reflection, consensual edifice-building, etc. This should not be surprising – people are individually and collectively too embedded in these processes to be aware of them. But we must not allow the spurious salience of the place metaphor to stop us from reflecting on the absurdity of applying it to the design of environments for manipulating knowledge in a non-located information-based medium

Part of the problem with the place metaphor is that it forces us away from seeking 'what' we want to know, or expecting appropriate background resources to be made salient to us by the system when we express 'what' we are engaged in with respect to 'what' issue. We may know that the system is too dumb to do that in existing technologies, and that our best bet is to look a 'place' where we can start tunnelling through hyperlinks to get at a fact or opinion. But this need not always be the case.

Another problem with the place metaphor is that knowledge – even as expressed in web pages or documents – is always distributed, and thus never to be found in one 'place'. If knowledge could be expressed so that a dumb computer could assemble what we want to look at from distributed ingredients, instead of people having to repeat this task with long sessions of mouse-clicking, life would be much easier. If authors could express themselves in the context of such a system, so that background and related materials were automatically available to their readers, a lot of repetition and circumlocution could be avoided.

It is not just with respect to documents and crafted content that the place metaphor is inappropriate. Electronic messaging, though it currently seems naturally about 'sending' to a 'place' (or an 'account'), could be much more powerful if notions like 'about' and 'assertion' could organically create the system's links between 'messages' (brief ad-hoc contributions to a conversation), content and activities, as well as notions like 'to whom' and 'in reply to which item in someone's in-box'.

Please note that we are not implying any reliance on the representational techniques and complicated inference engines of Artificial Intelligence. Sadly, we also have bitter direct experience of the futility of that approach. The alternatives we propose are simply to allow interfaces for people to structure and contextualise their writing, resource-collation and messaging in meaningful but sharable categorical terms that can be processed by computers, to allow messaging and content to be discoverable and collatable automatically from their semantic properties, and to allow content and discourse to have common interfaces in each others' context.

Expression of meaning is still textual

Web pages and email messages are both based on simple standards for the transmission and display of text characters. This is why universal standards have succeeded for them; no-one had to agree about how to express meanings, just characters. The html standard has a very few, very broad categories like heading level and list-item, and the email standard allows a subject separate from body, but there are no conventional ways to denote more specific semantic types or gestures.

It would be very hard to reach agreement on standard ways to 'tag up' meanings, and we do not propose that standards should exist for this, but we note that there are some useful ingredients in place for at least standardising a non-textual medium for expressing meanings:

- the XML standard allows the exchange of arbitrary tree structures, and of the grammars on which they are based, in a widely accessible way
- the Dublin Core metadata standard allows the expression of basic descriptive assertions
- various domain-specific sub-standards are in development (usually based on XML) for expressing meaningful kinds of gesture, for instance chemical formulae and mathematical equations
- technology exists that could allow a standardised platform for expressing, sharing, and discussing a variety of community-specific semantic structurings while using

common tools for searching, organising and displaying structured content and metadata

What might it be like if there were conventions for expressing meanings directly, rather than as strings of text characters? Computers are very *bad at processing meanings*, but very *good at storing and processing abstract structures*. We would not expect computers to *do* anything with meanings (outside of restricted areas like simplifying formulae and querying databases), but they can display and organise structures and categorical types to human users on the basis of rules, and they can allow humans to make searches, organise content and specify relationships in terms of abstract types and structures. Fine-grained expression of meaning would be cumbersome in such terms, but broader-grained expression could be more convenient than prose to generate (think of the difficulty authors face in turning the 'mind-map'of ideas and relationships they are trying to convey into linear prose or bullet-points).

We list below some of the consequences of using text characters (organised into files or messages) as the standard for conveying meaning.

Hard to find appropriate content

The text-as-standard becomes an obvious problem for most people when they are trying to search the web. Standard web searches must be based on strings of characters found in the target documents; as we all know, thinking of useful search strings is a dark art and not always a useful one.

'Metadata' is a term used for categorical or semantic assertions made about pieces of content. Many web users are becoming familiar with the term, as frustration with character-string searches grows.

Of course, services have emerged which attempt to apply some categorical structure to web content:

- many pages have been sorted into broad categories by search engine staff, or by volunteers of the open directory project
- link sites, which painstakingly attempt to provide a common entry point for content related to a particular issue
- many authors and organisations now embed Dublin Core metadata in their web pages, sometimes alongside other metadata terms
- community or organisation metadata repositories, like link sites, can apply semantic terms to web resources (KnowNet's portal offers one example of such a repository: users and groups can apply any metadata they like to any resource or bookmark)

These services suffer from being based on structures that are either shallow (large-scale categorising efforts, Dublin Core) or sparse (community repositories). It would be much easier to leverage the human effort of categorisation and semantic tagging if documents were expressed with some semantic structure in the first place. For many documents or messages, structured authoring could provide enough 'self-metadata' to allow powerful searching.

Note that the same problem applies to email or groupware messages. In heavy traffic, important message content quickly becomes lost, and text-string search is often incapable of retrieving it (think of the head-scratching you've done to try to remember a phrase or exact term from a message so that you can search for it in a folder or conference).

Content does not interoperate meaningfully

Web content is widely used for educational purposes, but a constant problem in the design of collaborative learning environments is that web pages are monolithic and stupid. Monolithic in that they can only be delivered all at once, replacing whatever content had previously been displayed. Stupid in that they cannot send information to their environment (whether about user actions, or even just about their own properties). It is possible to impose special constraints on web-page designers to make it somewhat easier to assemble them into learning sequences and environments, but this method defeats the point of standards, and is in any case fraught with technical difficulties.

A second problem is familiar to those who have created educational web content. Whereas the dream of hypertext was to allow a universe of content to interact, putting an omniscient tutor over every student's shoulder as she acted on your content, the reality is that linked or background content has to be hand-crafted (or at least hand-linked) by the author.

How much better it would be if authors had only to create that part of their content which is unique to their purpose, with content from other sources accessible on the basis of its semantic properties, and able to be woven into displays and sequences on the basis of its structure.

How much more interactive the user experience could be if small chunks of content from whatever source could be woven into an integrated and managed client environment so that their actions within the content could be processed by the user's choice of software in the context of the user's preferences, peers and other activities being undertaken in the session.

Hard to link related content and conversation

A common problem in the design and use of groupware systems for learners is that it is very hard to sustain electronic discussions around resources.

In face to face communication, even by telephone, humans have a wealth of techniques for indicating topic and for 'pointing' at objects, resources, themes or goals. In electronic messaging, these techniques are not available, and other devices must be found, exploiting the graphical user interface. Embedded URLs are very useful for pointing to a resource in the first instance, but very bad at tracking the resources or versions involved in a sustained discussion. URLs are given the same salience in groupware systems as any other text within a message body; relationships and links are not foregrounded in the user interface.

If it is difficult to point to a web resource from a discussion, it is even more awkward for users to point to

- one discussion thread from another

- one piece of content from another

- a discussion thread from a piece of content

Two-way links are, at present, almost impossible to create.

're;'is a weak basis for discourse structuring

The 'reply button' is, of course, the generator of most electronic message traffic, as the frequency of the prefix 're:' in email subject headings attests.

Unfortunately, the progress of discussions is often lost in this succession of 're:'s'. The reason is that, while there is only one type of 'reply button', users will in fact be expressing many different kinds of response, some of which move the discussion forward or away from the original thread; this variety and movement is lost in a homogeneous chain of 're: re: re…'.

This makes it hard to follow threads in heavy-traffic sites, and very hard to return to old threads or enter them as a newcomer.

Note that it is technologically possible to design groupware that provides special ways of replying with conversational gestures.

A second source of information loss in the reply chain is that, while messages typically contain many and diverse ideas and points, it is only possible to reply to the entire message, not its sub-content,

This makes it impossible to follow the strands of threads as they branch.

Structured messages and structured linking of new message content to existing content and discussion would allow individual points to be organised transparently within message content and responded to individually. Instead of homogeneous linear threads, it would be possible to create diverse branching trees and webs of message content and links.

Content creation, deployment and management is technical

Because HTML is designed to describe the *appearance*, as opposed to the *meanings*, of content, authoring web pages is to some extent a technical process, even when graphical tools are used. Deployment (posting to the web) and management (editing, controlling navigation, version-control) are even more technical; usually authors and editors have no direct control over their content at all once it enters the deployment workflow.

This creates many problems for authors. It also separates those with knowledge of content from control over the navigation and link structure of larger content blocks, and makes flexible organisation of structure and content very difficult.

Readers, too, lose control over content needlessly. If content were authored, stored, organised and delivered in terms of its meaningful properties, users could flexibly control the way they want to view it. Currently, users have very little control even over the appearance of a web page they are viewing.

Structure is sparse, shallow and mostly accidental

Content and discourse on the internet is, of course, structured already. Roughly, the categories are site, page, groupware conference, discussion thread, etc.

But this structure is

- inflexible
- largely accident of history of content and message creation, and thus uninformative
- shallow both in tree-depth and in semantics – there is little meaningful range of structuring decision that can be applied to standard web content and messages

The 'action' is missing from 'interactive'

Learners and knowledge workers can look at content, but what can they *do* with it?

Because content is monolithically unstructured, and its link-structures are painstakingly hand-assigned during authorship, there is nothing in the nature of the content itself that invites action from a user, or that can be actively related to, or aggregated with, other content. 'Interactivity' has come to mean either action on remote databases through data entry forms, or action on highly expensive islands of animation or branching movies. If content comprises vertical books, users prefer to read it on paper.

It is the constant frustration of workers trying to elicit usership of e-learning platforms that their users find so little to do with the content they are presented with that they engage very shallowly and infrequently.

Discussion divorced from content

Conventional systems make a sharp distinction between conversation and resource (the popularity of listserver archives as a public technical resource reveals a low-tech example of the internet community spontaneously disagreeing with this distinction). We believe that a large part of a community's 'know how' is knowing what people talk about when they talk about a field or a work practice.

If both content and messaging were authored and accessed in structured ways, the distinction would naturally blur. Content could grow in small chunks with rich relations to networks of discussion and content.

Key issues and ideas

Sharing & structuring: Metadata, XML & semantic expressiveness

We feel that the key requirement for structured authoring, metadata and messaging to succeed is an architecture for sharing and reconciling diverse structural 'grammars' and semantics.

Our experience with metadata for online learning resources makes it very clear that there is no 'correct' metadata schema. Users will only feel comfortable and empowered when

they can directly influence, and immediately override, the structures they use to express or describe their meanings.

This of course creates a problem for interoperability. But that is only a technical, architectural problem – we can design systems to help maximise both coverage/interoperability and richness/expressiveness. Such systems will not be automatic. They will have to provide good interfaces for users to discover, discuss, refine, reconcile and decide on schemas, and for information professionals to design mappings and canonical schemas for common purposes. Tools must be able to seamlessly deal with a wide range of schemas (for instance our own advanced metadata tagging and search tools configure themselves on the basis of XML schemas chosen by users).

Users must have a wide range of structured editing tools present to hand in their appropriate contexts, together with a rich ability to find and link in related content.

Structured communication

Electronic communication is changing the way people work, enabling remote collaboration and accessible, self-documenting discussions. As uptake and usage of email and electronic groupware increases, though, the limitations of current protocols and tools are creating real problems for sustained collaboration.

One problem concerns the unstructured nature of email. While text messages can carry important meanings when their context is shared between writer and reader, they are harder to understand outside of that immediate context. Nor is it straightforward to search for those meanings as time passes after their initial receipt, or to combine meanings from multiple messages into new ideas or assertions, or to link those meanings to other electronic environments. At present, everything depends on text, and on hierarchical organisation imposed on message storage (whether in email clients or groupware environments).

This problem only exhibits itself at larger scales of usage and higher expectations of sharing and building on an organisation or community's discussions. For simple, lightweight use, simple email is effective. At slightly larger scales and slightly higher functionality expectations, conventional groupware makes a big difference, but, even when very carefully moderated and maintained, groupware systems become less useful and accessible as they grow in success.

Structured messages could be very useful in overcoming the frustrations users experience with large numbers of text messages. By 'structured', we mean that different aspects of the communicative act are represented distinctly in the messaging and message-browsing system. In particular, structured messaging could replace the concept of 'thread' (which organises messages by a sequence of 'reply' acts) with richer, multidimensional concepts like 'raising issues', 'solving problems', 'tracking progress', and 'linking related material'. Structure in the messages could be represented semi-graphically to message readers and composers; more importantly, the relationships among messages - and between messages and other resources and documents - could be graphically represented.

By providing an environment in which many small components and content-nodes were already available for re-use, embedding and cross-linking, and in which simple 'gestures' could replace textual elaboration, message composers could eliminate much of the longwinded re-expression and periphrasis that makes context-setting so difficult in existing email.

Readers of structured messages could avoid having to scan long passages of redundant or familiar text; they could get the 'gist' from graphically depicted structure and relationships, and 'drill in' to the aspects of content and context that they needed.

Perhaps most usefully, message structure would automatically create a form of structured metadata allowing rich querying and alternative virtual organisation of communities' valuable discussions.

KnowNet has been building structured document editing, browsing, and interactive-annotation environments. We have plans to extend the annotation interfaces and architectures to full structured messaging, embedded within structured document repositories and structured resource collection.

We propose to develop a prototype of this structured messaging system, building on the tools and environments we have created already (XML-editing on and offline; structured document browsing; structured annotation; shared structured resource collections; structured metadata searching and tagging). We will provide a web-based messaging system with a graphical message-browsing interface (network / graph / mindmap metaphor rather than threading) and a drag-and-drop graphical interface for creating structured messages from text components, links and list-choice-in-context (structured metadata) gestures. We will use our web-XML-editing infrastructure to provide users of the system with a wide, and appropriate, range of message types and 'reply' gestures.

Freeing users to create webs of knowledge, practice and discourse

The key problem in any attempt to experiment with new internet architectures is ‘where will all the content come from’? The web exploded with content because of its very simple standards – anyone could create a web page (though it was very much like early desktop publishing in requiring naïve authors to attend to fine-grained page design). If we are to expect interestingly structured, interoperable content, authors are required to change the way they write; we have found this to be a very difficult learning curve for most professionals.

Our key idea is to design the environment to reduce the average ‘chunk-size’ authors must create, and maximise the value of small structured bits of knowledge in the right context.

This is why we feel it is key to

- provide a structured messaging environment, so that message content can become core content
- provide a very thorough and responsive interface for finding materials and relating them richly to new content, whether message content or ad-hoc ‘nuggets’ of knowledge.

- provide a transparent system for sharing structures, so that users can explore rich semantic structurings for their own uses without locking themselves into non-interoperable modes of searching and organising.

Three current projects

The communities

KnowNet is currently running a number of experimental projects for knowledge development in communities of practice. These are based on our current web portal environment, with considerable development of more fundamental software innovation planned for next year.

Three of these projects are described below.

E-VAL

E-val is a two-year project funded by the EC Leonardo fund. The project, which is co-ordinated by KnowNet and involves partners form 8 countries, aims to develop ICT tools to support the process of evaluation.

The tools will be tested in the different partner countries. There are 4 different zones planned for the E-val system. The first zone will facilitate the development of a shared knowledge base on the theory and practice of evaluation. The second zone will offer tools for those developing evaluation plans and programmes. The third will be a communications interface between evaluators and project staff. The fourth zone will provide tools for the analysis of the outcomes of evaluation and for programme evaluation based on an overview of the outcomes of different projects within a programme.

CEDRA

The CEDEFOP research arena (CEDRA) was set up in 2000 to promote knowledge sharing and development in vocational education and training in Europe. CEDEFOP is the European agency responsible for research and development in vocational education and training.

Changes in production processes, the introduction of new technology and changing work organisations are leading to an appreciation of the importance of knowledge development and communication to economic and social development. CEDRA is seen as an experimental activity promoting innovative projects and activities and working in collaboration with the VET research community.

It supports research networks working around different themes. These include the learning organisation, Information and Communication Technologies and vocational education and training, the learning region, work process knowledge, and work experience as a learning strategy.

CEDRA also supports the development of new communication tools and resources including the web based European Research Overview (ERO) providing access to

resources, and the ICT based tools being developed to support this web site and the ICT research network portal site.

At present, Cedra is comprised of the following three strands:

1. Sharing information & resources - a platform for the exchange of information about research projects, resources, activities and events

2. Thematic areas - focusing mainly on 'learning strategies and processes'

3. Knowledge development methodologies - an area for reflection on research and development methodologies.

Vet and ICT Research Network

The ICT and VET Network is a virtual European community of researchers in the field of vocational education and training (VET), undertaking an experimental electronic collaboration to develop resources, research and knowledge on the uses of information and communications technologies in VET.

Its aims are:

1.To promote research on policy and practice in the uses of ICT for education and training in Europe.

2.To pilot the use of collaborative software tools to support knowledge sharing, knowledge development and research activities.

3.To carry out an experiment using shared resources for developing knowledge about ICT and education and training in Europe.

In the initial phases of the project, in particular, relating to the experiment, the network will comprise people with a background in research in ICT and education and training.

In the long term, however, membership is open to other researchers, practitioners and policy makers working in the field of the use of Information and Communication Technologies in education and training.

Current implementations and environments

KnowNet’s current portals provide the full range of groupware functionalities, closely integrated with resource management and structured document-based discussions. Our portals give a customised front-end that combines a specialised groupware desktop (based on FirstClass, from Centrinity) with access to shared resources, structured search tools, interactive documents and document-centred discussions.

Basic shared messaging, conferencing and scheduling

For deployment in our portals, FirstClass groupware is customised for each portal (and each user). Users can access all the functionality of their FirstClass conferences either through the web interface provided in the portal, or with a special standalone client application (free to end users)

Conferences

Conferences are shared discussion spaces. Groups of users can setup as many conferences as they wish around specific themes. This allows 'contextualisation' of conversations where as messages in a particular conference relate to one theme. By replying to specific items, messages are 'threaded' on topics enabling easy navigation through a conversation between a number of users.

FirstClass documents can also be created and edited within conferences - aiding versioning of shared documents and negating the need for users to download document attachments.

As messages are stored in conferences and not personal e-mailboxes - colleagues joining projects have instant access to previous conversations and documents.

Groups

The use of Groups within FirstClass simplifies the administration of users, conferences and calendar. FirstClass User Groups control the visibility of conferences to users (via a 'Model Desktop') and users can be members of more than one group - thus inheriting the conferences on multiple Model Desktops. Conferences and Calendars can also be members of a Conference / Calendar Group and inherit the permissions of that Group.

Shared calendars

Shared calendars can improve the work flow within a project by acting like conferences where a group of people can view and contribute to the same calendar. Items in a calendar can be simple 'To Do' items or scheduled 'Events' - which can included booking of facilities, invitations to external colleagues. Calendar entries can include pictures, voice messaging and attachments.

Adding items to a shared calendar automatically adds the item to the contributor's personal calendar.

Synchronous chat

The FirstClass Chat facility enables synchronous conversation between users logged onto the server. Two types of Chat enables a 'Private' conversation between two or more users that is invisible to others unless invited to the Chat or a group chat where users can 'jump in' to an ongoing conversation without being invited. The text-based chat can be downloaded or posted into a FC message to enable those not present for the synchronous chat to contribute later.

Shared resources and descriptions

Our portals offer a very thoroughly implemented interface and architecture for sharing resources.

The ease of resource-sharing is a key constraint on the usability of groupware environments. While most groupware systems support straightforward uploading of documents and URLs into shared areas, users tend to become frustrated by the difficulty of locating these documents and resources when usage grows.

Our resource-sharing system allows extremely flexible organisation and search of resources, based on our Metaphase architecture for structured metadata sharing. Each user and group has a shared bookmarks folder – a virtual common file system. A web-based drag and drop interface makes the creation and maintenance of subfolders, resources and pointers easy, while putting a rich and powerful set of permissions-based collaborative management features at all users' disposal. Because our shared bookmarks use hot pointers rather than file copying, the same resource can be organised into many different views without creating clutter or versioning problems. However, though our shared folder system is very powerful and useful, it is not the key feature of our system. Most importantly, our system puts the power of xml-based structured metadata to work for users.

Metadata - structured methods for describing or documenting resources – is easy to access as users browse through resource folders (Dublin Core metadata is shown on rolling over a resources row in a folder, other metadatas are one mouse click away). And, of course, metadata allows very powerful and flexible searches to be made and shared. Our system allow arbitrary, multiple metadata schemas to be used to describe resources and folders, and provides powerful interfaces for tagging, searching and browsing. Because metadata can be richly structured, it allows multidimensional organisation and cataloguing. Because we put all this power at all users' disposal, many different views and interpretations of the resources can co-exist in the same system.

Combining the ease and richness of our shared bookmarks with the limitless organisational and documentation power or our metadata systems allows a flexible way to allow shared resource bases to remain accessible as system usage grows

Shared interactive-xml key document repositories

Our portals include a number of XML document repositories, together with a range of editing, viewing, and annotation tools to allow users to create, employ and develop knowledge in a rich interactive environment. Our current examples are largely for research communities exploring vocational education and labour markets in Europe, but the same techniques can be applied to any community which requires collaborative knowledge sharing, knowledge-based discussion, or online collaboration.

Currently we use our own offline editing environment, but we will be creating other environments to make it as easy as possible for users to create structured documents or small chunks in context with a rich range of embeddable or linkable resources.

We can provide a number of renders from the content in these repositories. Currently we deploy a very easy-to-use and flexible graphical viewer and a plain print render.

Structured messaging from 'inside' documents

A key functionality available in the XML document viewer is the ability to browse or join discussions hosted within the document structure itself. The discussions are also available from within our groupware conferences, and can be copied to any email address.

Structured search

Our resource folders and xml documents allow users to apply metadata to any object or content-chunk in the system. Our search tool allows users to specify and combine searches based on that metadata. Search results are displayed in a compact and powerful result-folder interface and tightly integrated with the resource folders to allow drag and drop actions between them.

'Edit this page' ease of content-management by users

Our users have access to the creation and editing of plain web content with the ease of 'edit this page' buttons in the pages themselves, based on Manila from Userland.

We will be extending this ease of editing by employing the next phase of our XML structure editor to allow users to graphically manipulate and create site structure from within their browser.

What we hope to learn from the projects

We are hopeful that, because of the small determined groups of researchers who are the targets of these projects, we can maintain a high degree of interaction with the communities.

We hope to learn:

- what functionalities our communities make use of, and what extra functionalities they could use
- how people with shared community of practice can create knowledge in concert
- how to make our interfaces more useful and our implementation more responsive
- how to create cleanly extensible architectures, and foresee interoperability issues in future functionality

Future plans

In the near future, we will be pushing forward several lines of design and development which will feed back into the projects, adding functionality and informing our long-term planning:

- re-designing and re-deploying our structured metadata interfaces for describing and searching using arbitrary metadatas
- adding a new interface for managing and discussing the structural schemas underlying metadata and XML documents
- improving the facility for semantic structuring and assertion-making in our XML editors
- implementing a first-phase structured messaging environment, utilising reply gestures and structured message-types, initially for project management semantics
- designing a new, web-based XML editor, embedded in a search and link environment for assembling meanings that include or relate to existing content
- experimenting with other motifs for XML content creation, using qualitative data-gathering techniques
- experimenting with report-generation and querying of qualitative data harvested from structured content and structured discussions
- trying to organise collaborations with like-minded developers and thinkers
- if you find some of the ideas or developments we have discussed above interesting, please contact us