

Project Acronym: JISC Academic Social Networking
Version: 1.2
Contact: Dr. Laura James
Date: 22 April 2010



JISC Academic Social Networking Final Report

Dr Laura James
Anne-Sophie de Baets
Dr Ian Boston
Simon Gaeremynck
Tjhien Liao
Nicolaas Matthijs
Oszkar Nagy
John Norman
Harriet Truscott
Christian Vuerings

22 April 2010

Contact: Laura@caret.cam.ac.uk
CARET, 1st Floor, 16 Mill Lane, Cambridge, CB2 1SB

Table of Contents

<i>Acknowledgements</i>	3
<i>Executive Summary</i>	4
<i>Background</i>	5
<i>Aims and Objectives</i>	6
<i>Methodology</i>	6
<i>Implementation</i>	6
<i>Outputs and Results</i>	15
<i>Outcomes</i>	17
<i>Conclusions</i>	18
<i>Implications</i>	18

Acknowledgements

We would like to thank JISC and the Institutional Innovation programme, which made this project possible.

Huge thanks to everyone whom we worked with at Flow Interactive, particularly Paul Heath and Jo Frudd, who remain inspiring and supportive even after the end of the project.

Thanks also to Alan Blackwell and others at the University of Cambridge Computer Laboratory, who kindly loaned us use of their user testing lab, which made our work much easier.

This work would not have been possible without the engagement and time of staff and students at the University of Cambridge, who participated in our studies and user testing.

Executive Summary

The genesis of this project was in the rapid rise of social networking sites such as LinkedIn and Facebook, and their uptake by academics and students. Earlier projects suggested that while using social networking sites exposed various privacy and IP problems, users found their interactions with these sites more rewarding than using institutional software. We wished to bring the affordances of these commercial systems to higher education, but in a fashion that allowed us to adapt them to academic purposes and requirements. We were also interested in the growing use of user-centric design (UCD), as design is an increasingly important factor in software creation in the HE world. We undertook user research into communications related to teaching and learning in the undergraduate, postgraduate and academic staff populations, to identify user needs, and took this forward with design work and user testing, then development work in Sakai3, building upon the gadget architecture of Apache Shindig. To do this, we worked with a commercial user design company (Flow Interactive), and Flow staff supported, worked with and trained our nascent UCD team; this was made possible by their willingness to engage in knowledge transfer, and other institutions should be aware that this kind of activity is open to them.

Although we have not moved forward with Shindig within Sakai at this stage, we acknowledge that the gadget architecture is powerful and believe that a gadget framework (Shindig or Wookiee) will be valuable as the project moves forward. More importantly, we have found the gadget model to be a key factor in designing and developing powerful next generation interfaces, and the entire Sakai3 UI is structured around gadgets, making it easy to work with and modify. Thus the concepts behind our original choice of Shindig have come through in our final and ongoing work, even if the detail did not.

Flow felt that UCD is a scalable process, and were confident of some successful outcome at any budget. This is because for creative work such as design, and the well understood discipline of UCD, do not have "one right answer" - there are many potential reasonably good solutions, even with limited time and resources (compared to not attempting UCD). This means that even modest projects should be able to get good outputs of some sort using these methods.

We have substantially added to the body of knowledge relating to user-centric design practice in UK HE, especially as regards emerging technologies such as social networking and other Web2 and new media systems. As well as experimenting with new methodologies, sharing with others how they might use these methods, and evaluating our experiences, our work has generated some concepts for designs and systems which we have not been able to take forward in this project, but which we hope that we, or others, will be able to take forward in the future. In particular, some of the early concepts from the user-centric design phase are engaging and intriguing, and will be worthy of further exploration. We also produced a range of materials for others who wish to undertake user-centred design in the HE sector, allowing them to benefit from our experience (our main handbook has been used at Georgia Tech, Michigan, and elsewhere).

A full user-centric design process can be a powerful method for uncovering academic user needs and wants, which would not be uncovered by a project team working instead through a prototype&test iteration process (which may be more conventional in HE), and can lead to more usable IT systems in HE which address needs which might otherwise have been missed. To achieve this, a clear research question must be posed, and a team where technology-lead thinking is balanced with design and user-focussed thinking is essential throughout. Strong team working and group communication (using many methods and channels, as sharing design thoughts can be challenging) is also important. Nonetheless, there are real challenges in UCD, in communicating design work to stakeholders who have not been engaged throughout the process in detail, and in empowering stakeholders to assess design outputs effectively.

This project has hugely enhanced the capabilities of a major open source VLE/VRE (Sakai), by adding in scholarly networking functionality, including the ability for academics and students to communicate easily and flexibly, in the engaging manner pioneered by Facebook, but within an academic context and around learning and research content. Sakai3 will have social and collaborative aspects in every part of the system, with seamless creation and sharing content of all types throughout. Complex projects, whether VLEs or other systems, which seek to add social features or to evolve to more socially-oriented versions, should take a holistic viewpoint from the start, and tackle these problems across their systems, rather than attempting to retrofit small social widgets to existing platforms. User research, seeking answers to one or more well thought out research questions, is a strong way to start this sort of redesign effort.

Background

Social Networking websites such as FaceBook and LinkedIn have emerged rapidly. The use of FaceBook by higher education students has attracted particular interest and it is clear that many students are using the FaceBook site for academic work. The relevance of LinkedIn to academic networks is less widely observed, and adoption among staff is lower. This may be attributable to a less-good fit with academic purposes and the lower rate of adoption of new technology among older people; however we believed that the model might need only modest change to be of considerable value in academic networks. These sites (and likely developments at Google) pose dilemmas for the Education Community, particularly regarding intellectual property rights asserted over user content by these commercial sites and concerns about authenticity, authority and identity of the individuals using the sites.

In 2007 the University of Cambridge experimented with exposing secure content from its Virtual Learning Environment to sites like FaceBook and iGoogle, as well as desktop 'Widgets' such as Mac Desktop Widgets, Windows Vista Sidebar Widgets and Google Desktop Gadgets. During the same period the University carried out a 'Learning Landscape' study that investigated, among other things, student use of technology. Our conclusion was that use of, and dependence on, external sites for mainstream learning support has significant drawbacks - such as lack of access to content for marking and record keeping, lack of even student control over content, privacy concerns, etc. However, it was equally clear that students found their interactions with social networking sites far more rewarding and intuitive than their interaction with institutional software.

As part of the University response, we began to explore more Google-like approaches to software design and specifically, we introduced 'Data APIs' modelled on GData, and a gadget dashboard modelled on iGoogle, into the VLE. These innovations went into production in summer 2008. We were exploring how to introduce network-forming affordances when Google announced the OpenSocial APIs in November 2007 and the Apache Shindig project in December 2007. Since then we have been engaged with these efforts with the goal of introducing the affordances of FaceBook and LinkedIn to higher education, but in a fashion that allows us to adapt them to academic purposes and requirements.

Our original proposal was to implement the Google OpenSocial APIs at the University of Cambridge by deploying the Apache Shindig project as a major infrastructure initiative. There were many institutional goals that this initiative was expected to support, but this proposal focused on the student benefits for teaching and learning. By developing a user interface through rapid development cycles with students and staff, we hoped to deliver affordances that are similar to those of FaceBook and LinkedIn, but which were adapted to academic purposes and which do not present intellectual property concerns. Since Shindig is an Apache project, and is already in production at a number of social networking sites, the facility to implement similar infrastructure is available to other institutions today. What was not readily available was a demonstrator of how affordances from FaceBook and LinkedIn could be adapted to academic purposes. We aimed to provide that demonstrator, using the Java version of Shindig, but the availability of a PHP version and the liberal Apache licence, means that the lessons learned could be applied at institutions running Moodle or even commercial software, and thus be of value to the whole sector.

We also had a growing interest in user-centric design, with a view that design is an increasingly important element of software creation in the HE sector; as noted above, the "competition" of commercial Web2.0 offerings includes many systems with strong design and the expectations of users, particularly students, are for web systems where user needs and design are at the forefront of the experience. We wished to explore how user-centric design methods could support this kind of project, in engaging us with users, exploring their real needs, and delivering a tested solution. As such, we chose to embark on a process together with a commercial user experience consultancy, who would both support us through the project and train our team, and to evaluate the success of this exploration.

Aims and Objectives

This project aimed to involve users in the process of developing a number of “social network”-style functionality/applications within an institutional VLE, where the applications are focussed on supporting and enhancing teaching and learning. These applications would be deployed at the earliest possible stages within Cambridge’s production VLE so that the utility and usability (both technical, and social/pedagogical) could be assessed, and feedback pushed back into the next development rounds.

We also intended to explore the utility of user-centric design in the processes of creating software for the HE sector.

Methodology

We began the project with user research examining the areas of teaching and learning communication in our undergraduate body (followed by postgraduate and staff in early 2009). This fed into analysis, development of design personas, and then a process of user-centric application design. In parallel we investigated other social applications in the consumer and education spaces, and researched uses of social tools of all kinds (particularly people’s online profiles) within the University at present. This also fed into the design process. We also built up the technical infrastructure of next generation Sakai (plus related development and demonstrator servers) to enable faster development later on.

The plan was that user-centric design methods would be learnt during the undergraduate research phase with training and practice demonstrations from Flow Interactive, with the next 2 phases (staff and postgrad) allowing CARET staff to demonstrate their new skills and to have this competency confirmed by Flow. The project also attempted to share methods with the HE emerging technology community to enable others to undertake similar work, and if appropriate other institutions user-centric design results would be fed into our final design selection process (as well as Cambridge-based results).

The subsequent workplan was to focus on iterations of design, development, test and evaluation. Software tools and widgets would be developed within Sakai and would be tested in-house as well as with real users in usability/utility assessing user testing. Additional tools would be delivered at each stage so that at the end of the project a wide suite of tools designed to support learning would be available to all our staff and students, who can select the tools most useful to them.

Implementation

The project had several phases as detailed above. We began with scoping and planning, followed by user research to explore campus needs and behaviours. We then moved into an intensive user-centric design phase. After this, we planned to move into development, but in practice were held up by two things: the delay in the underlying software project needed for our development to fit into, and a recognition of the need to document, analyse and archive our user experience process. This, then, was the next period of work, including synthesising our experiences into guidance for others, as well as archiving and recording our experiences of the project itself more completely than had been possible at the time of much of the activity. Then, we had a backend development work phase, in parallel with design review, moving into the actual UI development work. A final round of user interface development and user testing leads into final campus-wide deployment (during Summer 2010). These phases are described in more detail in the following sections. During the project, we shifted our plan from development of a suite of tools to a more integrated design in response to the things we learnt during the project (more information below).

In the background throughout the project, we explored existing social and scholarly networks, and also investigated best practice in various areas related to our work, such as the balance between privacy and publicity, and how relationships work and can be represented in online networks. We were able to work closely with related projects and share in their learning, as this research was not a core part of our work. This work, undertaken as it was throughout our project (not just as an initial “literature search”), was helpful in that we were able to stay up to date with the rapidly moving field of research around social networking, particularly as it applies to academia.

User Research

We undertook user research with close support from Flow Interactive Ltd, a user experience consultancy based in London.

Initial planning work involving Flow and project staff pointed to a slightly broader project goal statement; we had initially identified “communications to support teaching and learning”, but it quickly became clear that from a faculty point of view, these were indistinguishable in many cases from communications to support research. So we chose the following problem statement for our user research:

*we want to develop a system that supports and enhances
the use of social networking for teaching, learning and research*

We identified key user groups to work with as:

1. undergraduate students
2. postgraduate students
3. staff

For each, we wished to work out who to recruit for our study. In user-centric work, the selection of study participants is not only to get a spread across the “obvious” demographic splits: for instance, male/female, young/old, etc, but also we examined the axes which we believe *may affect behaviour* in the relevant area of study: in our case, around communication for teaching, learning and research. One example might be the distance between the person’s residence and their main study/work location - this could well affect communication patterns. We found that this structure was very helpful in aiding us to see beyond the usual limits of demographic considerations. We selected individuals to work with to represent diverse positions (and even extremes) on these axes, where possible selecting people who could represent a good spread of points on the axes; this follows best practice in this area and allows the broadest range of perspectives to be gathered.

We used a combination of surveys and personal contacts to identify potential research participants. We undertook diary studies and workshops with individuals and small groups drawn from each of these user groups, drawing them out about the motivations and challenges in their communications around their academic work. (Other research methods could have been used, but with our time, budget and project constraints, Flow recommended diary studies and interview work.)

We then analysed the results in detail; this phase of work used a great many post-it notes and physical space methods for best results, and generated some of the project’s most memorable photographs and moments. Details of our methods are available in our handbook.¹

We separated each piece of information about an activity we had into a block containing the user’s goal, how they went about achieving their goal, what their underlying motivation was, what frustrations they hit, and what went well for them. We then grouped these, and categorised them, deriving aspects of behaviour which are driving the ways in which the users act around academic communication.

¹ <http://bit.ly/bt69u1>

User-Centric Design

This phase of work was unusual. Because this area of work was very new to the team, and to the sector at large, we opted to undertake the work at Flow's offices, where project staff could be fully immersed in a design environment, able to access expertise during the working day as required. This was an unusual step and we are very grateful that Flow was able to accommodate us in this way. This period was intense, with frequent team meetings, and with two CARET project staff based out of a small office at Flow, but also communicating regularly with the rest of the team back in Cambridge. This enabled us to simultaneously train our staff, enhancing their design and conceptual skills, and to undertake the project whilst benefitting from the experience and mentoring of Flow staff as much as possible.

The first part of this phase was characterised by ideation of concepts, discussion and narrowing down of selected concepts, and the second part of user testing and design refinement.

At this point, we were aware that our project deviated somewhat from the "normal" projects which would follow this kind of UCD process, as we did not have a clear system definition - we knew we wanted to develop tools to bring the best affordances of social networking to academia, and that these would need to fit within Sakai3. However, Sakai3 was still at a fairly early design stage, and it wasn't clear how the project's tools would fit into it; in addition, the scope for those tools was very broad. This means that we are less confident in the outputs of this phase, although overall we feel this was an extremely valuable learning activity, which produced some strong interim outputs, and some reasonable eventual outputs.

The ideation part was undertaken twice; the first involved a new Flow consultant working with the research team in Cambridge to create concepts/features, and then working with the full team to rank the features according to technical difficulty and "business requirements". This exercise did not seem to go well; and so the outputs were 'parked', and the ideation and idea selection processes were repeated separately in London the next week.

The work began with a clean sheet of paper, thinking of a large number of ideas for systems fitting somewhat within our research area. It was surprising that we for the moment stepped back from the user research results, and indeed this early design work could perhaps have been done in parallel to the user research. The personas were used to validate the choice of ideas which we took forward, but did not inform the creative stage itself. Later on, of course, the personas played an important role in user testing and design refinement.

This ideation lead to the creation of over 20 individual concept ideas, from some which were fairly obvious and intuitive, to radical ideas which would take a great deal of design and development. We narrowed down this choice by selecting ideas which seemed tractable and which aligned with the "business" needs of the project and the campus.

We highlight here the three initial concepts which we felt were valuable for the campus and potentially realisable within the project, whilst still retaining powerful original ideas.

Concept 1: Go Outside!

"Building a real life academic relationship by showing you what's happening."

This concept is based on the idea that real world relationships are the driving force behind any sort of digital networking. This concept encourages "getting out" of the system, and gives triggers and pointers on how real life relationships can flourish, and from these real world relationships build a digital network. The main means to do this is with events. This concept helps creating and participating in real world events, and more or less everything is based around these. Network building is based on people being together at events and acknowledging this relationship digitally. Resources and people are concentrated around events, and the system facilitates meetings by constantly offering events to go to, and showing who will be there. The primary persona for this concept is an outgoing fresher who would like to expand his/her network quickly. This concept also tries to acknowledge that most of the information exchange will happen in more intimate "tea room" like situations.

Concept 2: The Boardgame

"Guiding you on your academic career by providing you with first steps and connecting you to go beyond that. Showing you how you can help others."

This concept is based around the idea of help. It is aimed for a persona who needs help, and needs clear instructions in terms of what the next step is in achieving their goals. The metaphor of a board game implies a clear visual representation of "official" steps required to achieve a goal, and giving contextual information about other people or resources which can aid this process (the unofficial way). The target persona can view similarities with other people and can get relevant information around people, events and resources. A secondary idea behind this concept is the metaphor of "ballroom dancing" where 1:1 relationships become important: coaching, pairing people with different levels of expertise within a field can be a way of giving assurances and help for a user. Network building and profiles are based on similarities and help to surround an individual with other people who are in the same situation.

Concept 3: The Switchboard

"Managing and building your network the way you like and making communication more effective."

This concept is based on efficiency, time saving and relevance. It is trying to define a system which can be customised and tailored to the exact needs of an individual user. The primary persona of this concept is somebody who is more senior and does not have very much time, thus wants only relevant information, presented in a way which suit his/her needs. The concept can be seen as a "iGoogle" type widget system which has exactly those elements which are relevant to the user with maximum customisability. Network building is more traditional here, and features are geared towards providing and managing resources and outgoing information. The metaphor of traffic sign is prominent here: traffic signs are artefacts which are put out by somebody more senior or experienced to guide others (imagine things which give you instructions where to find particular resources, BEFORE contacting me etc...). The other part of this metaphor is "the spy" who is the primary persona of the board game concept. He is the one who will scan the available traffic signs to get help or guidance. This concept gives tools to realise the relationship between the traffic signs and the spies.

In both the following sections, the project team needed to represent system concepts using visualisations of screens, but so that the users would be able to perceive the concept, and not get distracted by details such as features. The differentiation between concept and feature, and learning around how to present materials in this way was an important part of the team training from Flow.

Paper prototyping - user testing 1

Once we had selected these three major "extreme" concepts, each tailored to meet the particular needs of one of the user research personas, we refined the concepts into realisable ideas, and then created paper prototypes of each. These concepts were designed to be very different from each other, to capture the essence of what might be needed by each persona type. These were user tested with users representing the various persona types (as far as we could select for these using screener questionnaires) across the three main academic demographics (undergraduate, postgraduate, staff). We were lucky to be able to get access to the University of Cambridge Computer Lab user testing rooms, enabling user testing to occur in one room, with observers able to watch and make notes in real time in an observation room via a one-way mirror.

These paper prototype user tests used fairly open-ended questions, to evaluate how users felt about the concepts, not the detailed user interactions. An example would be asking a user what they think about a screen, and what kinds of activities do they feel would be open to them if they clicked in a certain region of the screen.

Onscreen prototyping - user testing 2

We then used the feedback from this user testing session to refine the designs, and to move them somewhat from being concepts to being screens closer to ready for implementation. On Flow's advice, we used the user feedback and user research data to inform a merge process which combined the concepts into one system which would work for all three personas at this point, following the process Flow advised. We built richer prototypes of an integrated concept, as part of an overall "site map" and tested those using on-screen clickable prototypes. This time, the testing involved consideration of user tasks, and we used a range of campus users, not selecting to match the personas this time.

From these test results, we then made final refinements to create concept screens, shown below.

A dashboard view, focussing on the user's personal content and information:

The screenshot shows a web browser window with the address bar set to 'Academic Networking Test'. The page has a navigation bar with links for HOME, COMMUNICATION, CONTACTS, EVENTS, DOCUMENTS, CALENDAR, and SETTINGS. A search bar contains the text 'a good book on data structures...'. Below the navigation bar, the dashboard is divided into several sections:

- My Events:** A calendar view showing events from Monday to Friday. Events include 'Calculus Supervision' (14:30 - 15:30), 'Talk: Biofuels in the 21st Century' (17:00 - 18:00), 'Talk: Fuel efficiency in 30 minutes' (12:00 - 12:30), 'Seminar: Data Structures II' (12:00 - 15:00), and 'Course: Field Study' (10:00 - 14:00).
- Inbox:** A section for messages. It shows 'You have 4 new messages - READ'. Messages include 'Fred Rolle would like to be added to your contacts - YES / NO' and 'Martina Heiddler would like to subscribe to your feeds - ACCEPT / DECLINE'.
- News:** A section for news items. It shows 'Donna Stuart has finished listening to "Aida"' and 'Dr. Alastair Moorley recommended reading "Medieval Moral Tales"'. A red arrow points from the 'Inbox' section to the 'News' section.
- My Widgets:** A section with a 'My TED talk video widget' showing 'Latest TED talk video' and an 'EDIT' button.
- My Goals:** A section with two goals: 'Reading 5 books on Computer Architectures' (21 people joined) and 'Writing a program in C#' (69 people joined). It also shows '43 Goals Achieved' and buttons for 'BROWSE GOALS' and '+ ADD NEW GOAL'.
- My today's pack:** A section with two items: 'Jeremy Truder @ Calculus supervision' and 'Calculus Lecture Slides' (last updated: 21 May 2009).
- Feeds:** A section with three feed items: 'Fred Nolan' (Reading Geometry II... need: data on French fables commented on Jimmy Reed), 'My Data Structure Notes' (Miriam Hole commented on this Anna Tripp updated this), and 'Algorithms Supervision' (Zhu Hong uploaded "My data" Dr Mia Cortes sent a link Dr Mia Cortes commented...).

An alternative “dashboard” view, showing more of an ‘external’ focus:

The screenshot shows a web browser window with a dashboard interface. The browser's address bar contains "Academic Networking Test" and the search bar contains "a good book on data structures...". The dashboard has a navigation menu with links for HOME, COMMUNICATION, CONTACTS, EVENTS, DOCUMENTS, CALENDAR, and SETTINGS. Below the navigation is a search bar and a "SEARCH" button. The main content area is divided into several sections:

- Inbox:** Shows 4 new messages. One from Fred Rolle asking to be added to contacts, and one from Martina Heidder asking to subscribe to feeds.
- News:** A snippet about Donna Stuart finishing listening to "Aida" and Dr. Alastair Moorley recommending "Medieval Moral Tales".
- Feeds:** Three items: Fred Nolan's comment on "Geometry II...", Miriam Hole's comment on "My Data Structure Notes", and Zhu Hong's upload of "My data".
- Explore relevant upcoming events:** A calendar view for the week of Monday (22) to Friday (26). Events include: Talk: Robotics Today (12:30-13:30), Calculus Supervision (14:30-15:30), Book Club Meeting (14:00-15:30), Job Fair (16:00-17:00), Talk: Biofuels in the 21st Century (17:00-18:00), Talk: Fuel efficiency in 30 minutes (12:00-12:30), Seminar: Data Structures II (12:00-15:00), Seminar: Optimisation (10:00-12:00), Sport: Running Club Event (08:00-09:30), Course: Field Study (10:00-14:00), and Course: Programming I Seminar (09:00-10:00).
- Explore relevant people:** Amy Stuart (writing a program in C#) and Jeremy Matthews (reading "The Machine in Me").
- Explore relevant documents:** Optimisation Slides (uploaded by Janet Digby) and Sorting Algorithms Video (uploaded by Dr. Jack Bohr).

A red arrow points from the "Inbox" section towards the "Explore relevant upcoming events" section.

User interface design review

At the end of the user-centric design phase above, we had a number of tested concepts which had undergone a final wireframe refinement stage. However, based on our own analysis of these designs, we had some doubts as to whether they would work as well for researchers and senior faculty, although they looked superb for undergraduates, and perhaps for early career staff engaged in teaching. We were also doubtful about the strength of the design part of the process as some parts of it had clearly gone less well. We therefore reviewed these designs with a group of senior academics in a facilitated workshop, which highlighted that our initial thoughts were right - although these designs were welcomed as powerful and engaging for undergraduates, and potentially some taught postgrad students, they didn't play so well for research-oriented or more senior staff. We also gave the academics an opportunity to come up with and share their own ideas around what a future system could do; this produced some ideas (notably that intelligent email filtering and automatic response generation would be highly valued) but did not produce anything substantive to rival our concepts.

As such, we did not take these rich concepts forward for final development. We decided to move forward with implementations of chat, messaging, profiles and shared public content, including events, which were aspects of the concepts which were well received universally, and which by themselves offer compelling, social-network-style functionality in the academic environment. These main social components will be deployed on campus shortly. The richer concepts around information display and innovative user journeys we have "parked" for the time being, with a view to returning to them later for further investigation, design and potential deployment. In parallel, further design work in the Sakai3 community around content and groups continues, and once that is nearer completion we expect to see it showing many of the innovative social features and user journeys we have envisaged, designed and user tested. In particular, our user research and thoughts around user-centric design have been integrated into this work and will help steer the overall system in development.

We have now had working pilot systems around chat, messaging and profiles available online for some time, which have been deployed and used by a range of communities; one has been explored by the Sakai community and others interested in VLEs⁴ as part of this project, and another instance was created for an ESRC research project and was used by education researchers at the first phases of their careers around the UK.⁵ The ESRC project in particular provided valuable feedback about what components the early career researchers valued, and how they used this sort of system.

Software development

With some designs, design concepts, and clear requirements settled (notably, that messaging and real time chat were important, and that social features should be integrated throughout the system, not an "add-on"), we were able to set about the task of software development. The latter finding, that "social" features would need to be a fundamental characteristic of the entire system, and could not simply be added in as widgets into an existing VLE, meant that our work took on a larger scope than we had initially intended, with overall system design included in scope as well, rather than simply individual widgets.

During the project, the user requirements drove our development work on Sakai3, affecting and prioritising the user interface software design as well as the backend systems; this mirrors the process used by Sakai3 overall. The clear separation between display and datafeeds, which is built into the concept of Sakai3, was very valuable during our development and test cycles to enable us to deliver rapid changes, and to improve understanding and efficiency of code creation and rework.

Deployment

⁴ <http://3akai.sakaiproject.org/dev> (anyone can create an account; this should not be treated as a secure production system; note the server is occasionally redeployed with associated loss of data)

⁵ <http://www.tlrp.org/tel/spnecre/>

We are now taking this evaluated codebase forward with our designs for other system components to create a solid codebase for production deployment on campus this summer.

Evaluation of UCD processes

One of the goals for this project was to evaluate whether the UCD processes we used were valuable, and whether we needed to make any changes from the normal way that UCD might operate in a commercial setting, to meet the needs of HE.

At the time of writing, we have not fully completed this evaluation; it is likely to continue after the end of the project, to allow us to include the reactions we get on campus to the system we deploy; this “mass market” acceptance is a critical component of understanding whether the group and individual user testing possible to date is able to produce results which reflect the whole campus.

However, we can comment on some aspects. Our relatively inexperienced staff, plus the design team at Flow, found it much easier to conceptualise and evaluate concepts aimed at undergraduates (as student life has been experienced by all of the team) than those which might target academic staff, who can be something of a breed apart. In a future attempt, it might be possible to balance this with better information for the team about academic life, pressures and concerns, or it may help if more senior staff with deep HE understanding were on the design team.

It is also likely that the scope of our project was perhaps too broad. When providing “business goal” input, and scoping the project at the start and at design concept review points, the project management and project director were unable to provide clear and specific information/decisions, as the work was to fit within the framework of Sakai3. At the time, Sakai3 was still in active creation and ideation, and not all of the concepts were thought through; feedback into the communications area, which this project addresses, was therefore not as coherent as might have been desired.

Outputs and Results

A full list of output documents, presentations, videos and code links is provided on the project blog⁶.

User-centric design

We have created documents which include information about our *results*, such as personas, design concepts, and wireframes; as well as descriptions of our *methods* and experiences. Some of these are provided in the specific - the story of what we did; some are abstracted to the general case, where we felt able to do this, to better support others undertaking similar activities. We have produced a range of designs and interactions, at both concept and implementable design levels.

As well as recording what we did, and how others could attempt similar work, we have learnt a great deal about what went well and what went less well. Documenting our processes throughout was valuable as part of our learning experience, as well as enabling us to review our activities more thoroughly. It was time consuming though, particularly during the intense project phases.

Clearer goals at the start of the project would have been valuable; in addition, an illustrated timeline we produced for the user research phase was useful, and we should have repeated this exercise for subsequent phases and the project overall. The planning work at the start of the project covered the user research in detail, but did not look further ahead, which was a mistake. Our timescales were ambitious, given that we were learning at the same time as doing.

Videos and photographs are valuable records of our work, but it is important to recognise the distinction between research data, recordings of sessions which may be educational for our team to review, and recordings of training material which might be of use to others attempting UCD. Some of the video and photographic material from our real user interactions cannot be shared because of confidentiality.

The three training levels we used at each project phase were valuable:

1. we watched Flow consultants undertake a task
2. Flow consultants watched us do a task and provided feedback
3. we did the task by ourselves

In some cases, though, we realised in retrospect that our team's emphasis was on the activity, as we were studying *what we should do*; whereas we possibly should have paid more attention to *how the data was being worked with*. This may have weakened some of our actual work at some points.

The design phase based at Flow's offices was the work we are perhaps least confident in now; the challenges of uncertain and/or too broad problem scope, plus a brand new team member, and time pressure all impacted this work. It was good to have input from many Flow people at this time, each bringing new perspective; but at the same time, this was a challenge, as not everyone knew of the big picture of the project, or what work had been done so far. Training others was not something they were all used to doing as part of their work.

The design work was the least documented part of the project, with design decisions being somewhat opaque in retrospect. The merge process in which we combined the three distinct concepts into one seemed like a weak point. In review, we feel we lost a great deal of design originality at this point - the final wireframes were generally felt to be "disappointing." This may be due to lack of input from Flow, or that our technical/time constraints were too tight for us to make adventurous choices, or because the three ideas were individually too strong to be combined. There may also be implications for our ability to assess design (see below). The speed of this phase also made it challenging for the whole project team to stay in touch, particularly when combined with the remote working, and this may have resulted in weak coordination and goal setting.

The original concept generation stage at Flow was felt to be strong, and we intend to revisit these ideas after the project, to see which of them may be feasible to take forward as separate features or systems, either within the VLE or another system or standalone.

⁶ <http://bit.ly/cymmdZ>

When we reflect on our wireframes now at the end of the project, although they seemed somewhat disappointing at the time, we can see commonalities with other design work done within the Sakai community - both in terms of aspects of the design, and the reaction to the designs as communicated. This leads us to wonder whether our designs in fact were stronger than we realised, and/or whether as a group we were unable to assess the quality of designs.

Software outputs

Although our initial plans included work on OpenSocial applications, to be hosted within Apache Shindig in the backend codebase, we have not yet taken this path. Although OpenSocial provides a great interface to connect a social network of contacts to applications connected to content, it does not provide the social network itself; the profiles and contacts have to reside somewhere. As we wanted to include these within the scholarly environment, rather than hosted on an external commercial provider, we have built these within Sakai3. The functionality of OpenSocial, therefore, is likely to come along slightly later, and as one way of connecting the Sakai scholarly network to external content, rather than vice versa.

The project has substantially contributed to the development of social academic networking software within the Sakai3⁷ codebase, including the backend systems (Sakai3 Nakamura, formerly known as "K2", and built upon Apache Sling⁸) and the user interface.

In particular, the project has informed the use cases surrounding the construction and management of the Social Container within Sakai Nakamura (Sakai3 Server). This implements Social Network support for the OpenSocial API but enhances the use cases to include asymmetrical attributes associated with each connection in the graph, which was one of the major findings of the research phase. The code base for this work can be found in our github repository⁹ with the social graph implementation encapsulated inside an OSGi bundle.¹⁰ Specification for the technical area of work is in Sakai Confluence¹¹. This custom social graph will be connected to a Shindig based Social API via RESTful integration¹² to support Academic Networking Gadgets build on the OpenSocial Gadgets specification. In addition, where tighter UI integration is required, connection to the social graph through UI components conforming to the Sakai 3 widget spec will be possible.

The gadget concept remains strong in all aspects of our software development, and indeed the Sakai3 UI is primarily constructed from gadgets for flexibility.

Deployment

The project originally promised an institution-wide demonstrator of the final system; instead, we are in the final stages of preparation for a summer 2010 launch of our full production institutional VLE/VRE incorporating these features for all users on campus and beyond. We have demo systems displaying varying states of maturity of the different system components (such as backend code, frontend code,

⁷ <http://confluence.sakaiproject.org>

⁸ <http://sling.apache.org>

⁹ <http://github.com/ieb/open-experiments/>

¹⁰ <http://github.com/ieb/open-experiments/tree/master/slingtests/osgikernel/bundles/connections/>

¹¹ <http://confluence.sakaiproject.org/display/KERNDOC/KERN-531+Social+Graph+and+User+Connections>

¹² <http://confluence.sakaiproject.org/display/KERNDOC/KERN-530+OpenSocial+API>

and design) which are available most of the time. The best at the time of writing for others to browse our latest design work is linked in the footnotes.¹³

Community engagement and dissemination

We have attended programme meetings (including presenting at the last minute in response to a request, in a parallel track in Manchester!), run an Assembly about user testing, presented at the Sakai conference in Boston 2009, and blogged about our work. We'll be presenting at the 2010 Sakai conference in Denver in summer 2010 too.

We look forward to continuing to work on academic networking, and to share our results, code and experiences and learning, more widely, after the end of the project.

One unplanned output has been a number of videos and reports around how to use a user-centric design process in the higher education sector. These are available from the project blog.¹⁴

Outcomes

This project has hugely enhanced the capabilities of a major open source VLE/VRE (Sakai), by adding in scholarly networking functionality, not as an add-on, but integrated into every part of the system. Scholarly networking includes the ability for academics and students to communicate easily and flexibly, in the engaging manner pioneered by Facebook, but within an academic context and around learning and research content. It enables scholars to share information about themselves in a public online profile (which is used by others to find and connect to them). Sakai3 will bring "social", people-oriented functionality to all aspects of the VLE, so that individual pieces of content will be surrounded by the people who are working with them, their comments and so on.

We have substantially added to the body of knowledge relating to user-centric design practice in UK HE, especially as regards emerging technologies such as social networking and other Web2 and new media systems. We also have created some understanding of which projects this type of user-centric design process might work best for, and some of the pitfalls to watch out for. As well as experimenting with new methodologies, sharing with others how they might use these methods, and evaluating our experiences, our work has generated concepts for designs and systems which we have not been able to take forward in this project, but which we have shared in the hope that we, or others, will be able to take them forward in the future. In particular, some of the early concepts from the user-centric design phase are engaging and intriguing, and will be worthy of further exploration.

We are only just starting to explore in practice how real users will use the systems we have developed as part of their real academic experience, and expect to take this forward over the next couple of years, learning about how social applications can support teaching and learning in a campus-based HEI, and sharing these experiences. As such our evaluation of our UCD process and the project overall will continue beyond the end of the funding period.

The user research results concept that events are a key component in academic social networking, has enabled us to take forward the JISC EGRET¹⁵ project with the integration of the open source Talks.cam events syndication system.

¹³ <http://sakai3sam.caret.cam.ac.uk/dev>

¹⁴ <http://academic-networking.blogspot.com/>

¹⁵ <http://egret-project.blogspot.com>, and <http://www.caret.cam.ac.uk/page/egret>

Although we have not moved forward with Shindig within Sakai at this stage, we acknowledge that the gadget architecture is powerful and believe that a gadget framework (Shindig or Wookiee) will be valuable as the project moves forward. More importantly, we have found the gadget model to be a key factor in designing and developing powerful next generation interfaces, and the entire Sakai3 UI is structured around gadgets, making it easy to work with and modify. Thus the concepts behind our original choice of Shindig have come through in our final and ongoing work, even if the detail did not.

Conclusions

Social components of VLE/VREs should be integrated throughout the system design, not added on as an “extra”. Sakai3 will have social and collaborative aspects in every part of the system, with seamless creation and sharing content of all types throughout.

A full user-centric design process can reveal user priorities and desires which would not be uncovered by a project team working instead through a prototype&test iteration process, and can lead to more usable IT systems in HE which address needs which might otherwise have been missed. To achieve this, a clear research question must be posed, and a team where technology-lead thinking is balanced with design and user-focussed thinking is essential throughout. Strong team working and group communication (using many methods and channels, as sharing design thoughts can be challenging) is also important.

Implications

Co-training method

The method of working with Flow Interactive, combining actual work by our staff, consultancy work by their staff, and training of our staff, was highly effective, even if our project constraints (such as poorly defined scope, and a complex underlying system which was not fully designed) meant we could not take complete and full advantage of it. We spent less on consultancy than the alternative of a fully consulted project, and improved staff expertise. The work did take more staff time than would be required if we had fully trained staff delivering the project entirely themselves (without consultancy), and this needs to be budgeted when planning this kind of activity. We would engage in this type of co-training again on other projects, and would recommend it to others, providing the cost was well thought through and that a consultancy willing to work like this could be found. Flow were surprisingly open to knowledge transfer, and other institutions should be aware that this kind of activity is open to them.

Applicability of UCD

Flow felt that UCD is a scalable process, and were confident of some successful outcome at any budget. This is because for creative work such as design, and the well understood discipline of UCD, do not have “one right answer” - there are many potential reasonably good solutions, even with limited time and resources (compared to not attempting UCD). This means that even modest projects should be able to get good outputs of some sort using these methods.

(However, the counterpoint to this is that it is harder to validate the effectiveness of a given project activity, as rerunning the task may not produce the same creative outputs even if the method implementation was sound.)

There may be methodology differences when tackling pure open ended work - without a specific solution style in mind - as opposed to targeted problems where the shape of the system to be designed is roughly known. We had a fairly open ended problem (but perhaps did not realise this early enough) and yet the methodology we applied was perhaps better suited to a more scoped task.

On projects which involve collaboration across institutions, there may be additional concerns with UCD methods. In our case, we found that the personas travelled well and were felt to be applicable in multiple institutions; this might, or might not, be a common experience, depending on the research question studied. The requirements of IT systems, and institutional cultures, are less likely to travel

well in our view; this may mean that design outputs might not work well in multiple institutions, depending on the requirements expressed etc.

Choices and decisions around UCD methods

The choice of staff for the different project stages was clearly important; continuity of staff may have value at some points in the process but not necessarily all. It may have been possible to order the research and design phases somewhat differently too. We do not feel we have advice we can give in this area, other than that projects attempting UCD should consider whom to place in which roles when in the project should think carefully about the impact of their choices.

One key challenge of the use of UCD in an HE environment was that of recruiting users; term times and other academic time pressures limited our project timing. Selecting users to match our personas, and fitting into academic timetables, was difficult.

Challenges of stakeholder engagement with design work

In this and in related Sakai3 design work, we found that we hit similar challenges around assessment and communication of design outputs.

Specifically, sharing wireframes of proposed designs or concepts tended to elicit criticism from stakeholders. This may have been because our communication of the design work was poor; perhaps wireframes were not the best choice to communicate a design concept, or they were too “polished” in appearance. It may have been that stakeholders who have not been exposed to the detailed design process have overly high expectations of the outputs, and therefore tend to judge them harshly. It may also be that stakeholders unaccustomed to and inexperienced in assessing or identifying strong design simply cannot manage this task effectively without support or training.

UCD best practices would suggest that rougher visual representations of design work better when seeking user and stakeholder input, as the viewer understands implicitly that they are seeing a work in progress, not a nearly finished system. When seeking buy-in from senior stakeholders, presenting user response (such as user testing results, particularly “real” data such as user quotes or videos) can be particularly effective; we have not fully explored this in a university context.

These observations present challenges for future design work in these environments, and we cannot offer strong solutions at this stage. Nonetheless, there may be mileage in trusting our designers more, working hard to explore multiple ways of sharing design work (and evaluating all of these), and using evidence to supplement presentations of design outputs.

Using larger teams of designers, rather than an individual, may also help them handle the critical responses which we have found to be all too common when seeking feedback from academics.

Social components for VLEs

As we reflect on how our expectations for the social components of our VLE have changed since we started this project, it is clear that our original ideas around socially-enhanced widgets and individual pieces of functionality were wrong. To provide the compelling user interactions of FaceBook and its ilk, we need to embed social activities and information throughout the entire interface. This is in many ways a greater design challenge, requiring high level conceptual and architectural design skills and a strong understanding of the user base and its behaviours and expectations, particularly around privacy, context for online connections or relationships, and innovative and previously unimaginable user activities.

Complex projects, whether VLEs or other systems, which seek to add social features or to evolve to more socially-oriented versions, should take a holistic viewpoint from the start, and tackle these problems across their systems, rather than attempting to retrofit small social widgets to existing platforms. User research, seeking answers to one or more well thought out research questions, is a strong way to start this sort of redesign effort.

Our work has shown that user-centric design can be a powerful method for uncovering academic user needs and wants, which we would not have uncovered through other forms of investigation.