The Design and Implementation of a Mobile Learning Resource

Mike Sharples, Dan Corlett, Oliver Westmancott,
Educational Technology Research Group,
School of Engineering,
University of Birmingham,
Edgbaston,
Birmingham B15 2TT

Email: handler@eee.bham.ac.uk

Tel: 0121 414 3966

Fax: 0121 414 4291

Keywords: interactive learning environments; personal technology; mobile learning

Abstract

The convergence of mobile communications and handheld computers offers the opportunity to develop technology that will assist individuals and groups to learn anytime, anywhere. We describe the theory-informed design, implementation and evaluation of a handheld learning device. It is intended to support children to capture everyday events such as images, notes and sounds, to relate them to web-based learning resources, to organise these into a visual knowledge map and to share them with other learners and teachers. A working prototype system, for children aged 9-11, is discussed and evaluated, as an exemplar of personal mobile systems for life-long learning.

1. Background

Over the past ten years, educational researchers and practitioners, policy makers and politicians have mapped out a new landscape of learning as a situated and life-long activity. The defining features of contextual life-long learning (CoLL) (Sharples, 2000) are that:

- learning is not confined to pre-specified times or places, but happens whenever there is a break in the flow of routine daily performance and a person reflects on the current situation, resolves to address a problem, to share an idea, or to gain an understanding;

- formal education cannot provide people with all the knowledge and skills they need to prosper throughout a lifetime. Therefore, people will need continually to enhance their abilities, in order to address immediate problems and to participate in a process of continuing vocational and professional development.

A consequence of this reconceptualisation of learning is that the environments where CoLL occurs cannot be pre-specified, but are created through the activity of learning. Nor can the environment be decomposed into elements that are independent of the learner (Roth, 2000) but instead are dynamically constructed by learners interacting with their surroundings. For example, a student on an archaeology field trip finds a piece of pottery and thereby creates a micro-environment for learning that is fundamentally bound to a context that includes time, location and the student's knowledge, skills and available resources.

The last decade has also seen a revolution in communications and computing technology, with the installation of digital cellular phone networks, and the development of mobile computers and digital cameras. These three technologies are now converging, into personal digital assistants (PDAs) that can enable people to access internet resources and run experiments in the field, capture, store and manage everyday events as images and sounds, and communicate and share the material with colleagues and experts throughout the world. By happy coincidence, there is a natural alliance between learning as a contextual activity and the new personal mobile technology, so that it is becoming feasible to equip learners with powerful tools to support learning anytime, anywhere.

Although the component technologies to support CoLL are now available, we have found no detailed discussion of the design of personal mobile technologies for life-long learning. Companies including IBM, Microsoft and Hewlett Packard are promoting 'anytime, anywhere' learning with laptop computers. A study commissioned by Microsoft of school students given laptop computers found that they used computers at home for a wider variety of learning tasks than a comparable control group with desktop machines (Walker, Rockman, & Chessler, 2000). A project by Philips has developed a prototype personal communicator and organiser for children, based on the results of participatory design sessions with children aged 7-12 (Oosterholt, Kusano, & de vries, 1996). Druin and colleagues are developing, in collaboration with Alburquerque elementary school children, a generic interface for children using a "pan and zoom" metaphor (Druin, Stewart, Proft, Bederson, & Holland, 1997). A team at Simon Fraser University led by Inkpen (Inkpen, 2000) is carrying out a participatory design study with children to develop handheld computers for collaborative learning. The Classroom 2000 project at Georgia Institute of Technology (Abowd, 1999) has developed technology to enable students in a lecture theatre to read slides from the screen onto Personal Digital Assistants. Projects to design mobile technologies for adult learning include FieldNote from the University of Kent which integrates handheld tools for data collection and re-use, including a GPS device, into a system that enables fieldworkers to capture and share information (Ryan, Morse, & Pascoe, 1999). Fischer and colleagues are investigating software systems to support life-long learning that allow users to learn as they design artefacts following their unique interests and needs, but these are developed for desktop machines not mobile devices (Fischer & Scharff, 1998). Sharples (2000) sets out a general framework for the design of mobile technology for lifelong learning and gives a brief overview of the HandLeR project.

This paper provides a detailed account of a project at the University of Birmingham to develop a Handheld Learning Resource (which we have given the generic name HandLeR) that can assist people of all ages in their personal learning throughout a lifetime. This is an ambitious aim and to make the project more manageable we have specified three scenarios that span the range of users and types of learning. The scenarios are summarised below.

- An 11 year old child is on a school field trip. She captures images of an historic site, annotates these with notes and sketches and organises them into a visual idea map. The child extends the idea map by adding pages from an internet guide to the site and its history, accessed through a high-speed wireless network around the building. She then merges parts of her map with those of the other children to create both a personal and a group website of the visit.
- A radiologist is in her first year of specialist training in neuroradiology. She attends a case meeting and receives the MR (magnetic resonance) images for a case through a wireless network to a tablet computer. She marks up the images and describes the case using terms from a structured Image Description Language and then compares her description and diagnosis with those of fellow trainees and the consultant radiologist. She can call up a database of related cases, viewing the images and their structured descriptions. At home, she reviews the day’s work, discusses difficult cases with remote colleagues by voice and shared data, and posts queries to a bulletin board for discussion the next day.
- A senior citizen is recalling and organising a lifetime of memories. He revisits favourite places and old friends and captures snippets of conversation and images through a handheld or wearable camcorder that automatically adds time and location data to each item. Later, he arranges these into a digital album, adding spoken commentary and images from local websites.

These scenarios can be seen as snapshots of a continuing process of technology-mediated contextual learning. A central aim of our research is to enable learners to integrate these learning episodes across time, to support their growth and transformation of knowledge.

Research is already in progress at the University of Birmingham to develop a Magnetic Resonance Imaging (MRI) Tutor (in collaboration with De Montfort University, University of Sussex and the Institute of Neurology) (Sharples, Jeffery, Teather, Teather, & du Boulay, 1997) and a wearable camera that will capture images as the user moves around, add time and location data and enable users to annotate and share the images. This paper describes a project to realise the first scenario, of providing technology to support children learning on a field trip. The aim of the project was to produce a functioning demonstrator of a HandLeR that will enable children, aged 9-11, to capture learning events in the field, to annotate, share and organise them into resources for learning, and to communicate directly with other learners or teachers. This paper describes the design and implementation of the system and reports a small formative evaluation of it with children aged 10 and 11.

2. Design Methodology

The project followed a methodology of socio-cognitive engineering (Sharples et al., 2000a) which aims to analyse the complex interactions between people and computer-based technology and then transform this analysis into usable, useful and desirable socio-technical systems (technology in its social context). The methodology has been successfully applied to the design of a broad range of human centred technologies, including a Writer's Assistant (Sharples, Goodlet, & Pemberton, 1992), a training system for neuroradiologists (Sharples et al., 2000b), and a system to support requirements capture from customers for electronic test equipment. Socio-cognitive engineering draws on the knowledge of potential users and involves them in the design process, but it is critical of the reliability of user reports and it extends beyond individual users to give an analytic account of cognitive processes and social interactions, styles and strategies of working, and language and patterns of communication, so as to form a composite picture of human knowledge and activity that can inform system design. Our aim is to design human-centred systems that are based on a sound understanding of how people think, learn, perceive, work, communicate and interact.

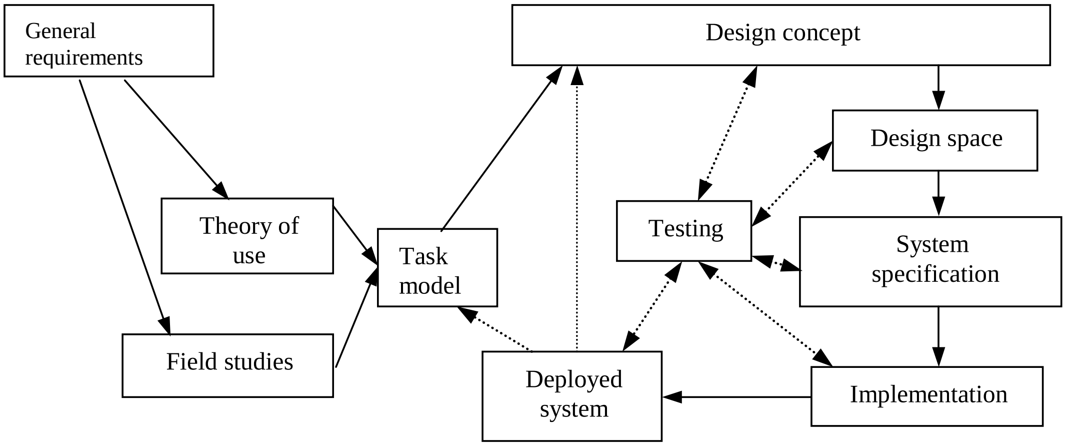

Figure 1. Overview of the flow and main products of the design process.

Figure 1 shows an overview of the design process. It starts by specifying the General Requirements and constraints for the system to be designed. It sets out the type of activities to be supported by the new technology (such as learning and knowledge management), the domain (personal contextual learning) and any general constraints (such as time and budget available for the system design). This provides parameters for two parallel studies: an investigation into how the specified activities are currently performed in their normal contexts, and a theory-based investigation of the underlying cognitive and social structures and processes. The outcomes of these two studies are synthesized into a Task Model whose aim is to describe the activity system (Engeström, 1987) that may include the main actors with their tools and resources, their physical, social and cultural context, external representations such as notes and diagrams, the rules and conventions that influence the activity, the distribution of labour, and the terminology and patterns of discourse.

The Task Model provides the bridge to a cycle of iterative design that includes: specifying a design concept; generating a space of possible system designs; specifying the functional and non-functional aspects of the system; implementing and deploying the system. Although this cycle is based on a conventional process of interactive systems design, it gives equal emphasis to cognitive and organizational factors as well as task and software specifications. The outcome is an implemented technology with guidelines for its use. When deployed, this will create a transformed socio-technical system with new activities to be supported and problems to be addressed.

3. General requirements

The general requirements for technologies to support contextual life-long learning (Sharples, 2000) are that they should be:

- *highly portable*, so that they can be available wherever the user needs to learn;
- *individual*, adapting to the learner's abilities, knowledge and learning styles and designed to support personal learning, rather than general office work;
- *unobtrusive*, so that the learner can capture situations and retrieve knowledge without the technology obtruding on the situation;
- *available* anywhere, to enable communication with teachers, experts and peers;
- *adaptable* to the context of learning and the learner's evolving skills and knowledge;
- *persistent*, to manage learning throughout a lifetime, so that the learner's personal accumulation of resources and knowledge will be immediately accessible despite changes in technology;
- *useful*, suited to everyday needs for communication, reference, work and learning;
- *easy to use* by people with no previous experience of the technology.

These requirements provided initial constraints on the design. The requirement for a highly portable device means that the technology should be light and capable of being carried and operated on the move. This suggested a pen tablet computer, with a high resolution colour screen. The technology to capture sounds and images and to communicate should either be built into the device or distributed around the body (for example as a buttonhole camera and a headset). The latter would require either a festoon of wires (conflicting with the requirement to be unobtrusive) or wireless connections between the components (the technology for this is now available, but expensive).

The requirement for it to be highly portable and available anywhere, indicates wireless communication, either through cellular telephone or wireless Local Area Network (LAN). Wireless LAN technology provides high speed, up to 11Mb/s, and relatively inexpensive communications. With a maximum range of about 100 metres from a base station it is suited to use within a building or campus. For locations such as parks, museums and historic buildings that are regular venues for school field trips then wireless LAN is a feasible means of wireless communication.

The requirement that it should be persistent, to enable a learner to accumulate and manage learning throughout a lifetime, dictates either that a child should be equipped with a single device for a lifetime of learning (not an option where hardware is developing so rapidly) or that the learning environment is separated from its physical instantiation on a particular device.

This latter approach is taken by standard personal computer operating systems and applications where, for example, Microsoft Office 2000 can run on a wide variety of desktop or portable computers. A person's learning environment could reside on a web server and could be synchronised with a handheld or desktop computer (in a similar manner to synchronising a personal organiser such as the Palm Pilot through a single button press). The advantage of this approach is that a child could run their own personal learning environment on a games console or interactive TV at home, a desktop computer in school and a handheld device on a field trip or a journey between home and school. A museum or field centre might provide a low specification handheld device for visitors which would enable them to capture learning events, synchronise their learning environment and communicate with their teacher or colleagues.

To adapt to a learner's changing skills and knowledge, the system must be able to maintain a profile or model of the learner which can determine the way in which the accumulated knowledge and learning material is stored and then presented back to the learner in new contexts. This presents a major research challenge. Most attempts at developing computer models of a learner's knowledge have concentrated on specific topic areas and learning over a short period of time. Research into cognitive and skill development over long periods (Eraut, 1994) indicates that learning is not simply a monotonic accumulation of facts and knowledge, but involves deep conceptual change and reconceptualisation. A life-long aid must either be able to detect, model and support such reorganisations of knowledge, or provide tools for the learner to manage this process.

Lastly, for the system to be useful and easy to use, the technology must present an appropriate and intuitive system image (Norman, 1986). A system image is the combination of product design, interface, and interaction design that hides the complexity of the internal electronics and programming and presents a "notional machine" which matches the user's tasks and understanding. The telephone system, for example, masks most of its complex electronics and communications behind the system image of a "speaking tube". A central issue is what should be the system image of mobile technology for contextual life-long learning. The traditional "desktop metaphor" based on office equipment such as files and folders is not appropriate for people learning in many locations and contexts. The system image, interaction design and interface for HandLeR have been derived from a task model informed by theories of situated learning and interview-based design studies with 9 and 10 year old children.

4. Theory of Use

The Theory of Use analyses theories of learning, cognition and social interaction that could inform the system design. For a more detailed discussion of a Theory of Use for personal learning technologies, see (Sharples, 2000). We can summarise contextual life-long learning in terms of "3 C's" of effective learning: *construction*, *conversation*, and *control*.

Successful learning is constructive process (Brown & Campione, 1996) that involves seeking solutions to problems and relating new experiences to existing knowledge. Central to learning is conversation, with teachers, with other learners, with ourselves as we question our concepts, and with the world as we carry out experiments and explorations and interpret the results (Pask, 1976). Learning is most successful when we are in control, carrying out an active and continuing cycle of experimentation and reflection (Kolb, 1984).

A theory of learning which combines these aspects within an all-encompassing framework is Pask's Conversation Theory (Pask, 1976). Conversation Theory is an elaborate and difficult construct that spans epistemology, educational technology, and cybernetics. It describes learning in terms of conversations between different systems of knowledge. Pask was careful not to make any distinction between people and interactive systems such as computers – with the great advantage that the theory can be applied equally to human teachers and learners, or to computer-based teaching or learning support systems. What follows is a distillation (and, inevitably, a gross simplification) of the theory, sufficient to provide a basis for system design.

Let us begin with a person engaged in some activity in the world, such as carrying out an experiment, or solving a problem, or exploring a park or museum. As the person performs the activity he or she tries out new actions, reflects on how these work and makes decisions about what to do next (Figure 2). The person is actively constructing an understanding of the activities. There is continual interaction and adjustment between the person's thoughts and actions. In order to gain from that experience, to perform it differently or better in future, the learner needs to form a description of themselves and the activities, to explore and extend that description and to carry forward the understanding to a future activity. That is the minimum requirement for any person, or any system, to learn: it must be able to converse with itself about what it knows.

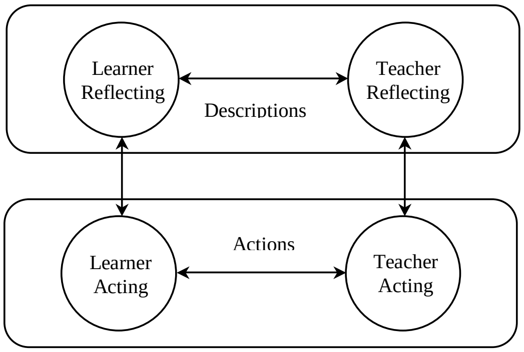

Figure 2. A framework for conversational learning

A more effective form of learning is when people can converse with each other, by interrogating and sharing their descriptions of the world. Suppose that two people – for example a student and a teacher, or two students – are working together on a project. Both people are interacting with the world and conversing at the level of actions – "look here", "what's this?", "do that". They are also conversing at the level of descriptions by exchanging reflective descriptions of their knowledge: "what do you think of this?"; "why do you do that"?

We can say that the two people share an understanding if Person A can make sense of B's explanations of what B knows, and person B can make sense of A's explanation of what A knows. Thus, it is through mutual conversation that we come to a shared understanding of the world. Learning is a continual conversation: with the external world and its artefacts, with oneself, and also with other learners and teachers. The most successful learning comes when the learner is in control of the activity, able to test ideas by performing experiments, ask questions, collaborate with other people, seek out new knowledge, and plan new actions.

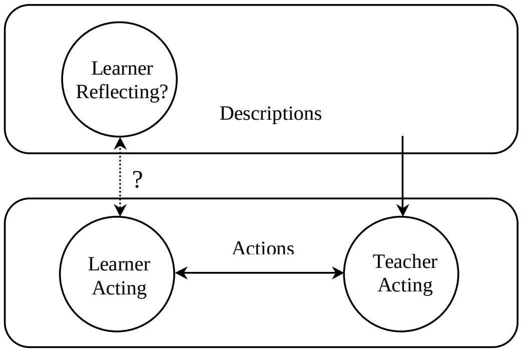

Figure 3. Computer-aided instruction

What place is there for technology within this conversational space? One possibility is for the computer to take the place of the teacher (Figure 3). That is traditional computer-aided instruction, and the difficulty is that it only covers part of the conversational space. The computer can hold a limited dialogue at the level of actions: "look here"; "what's this?"; "do that", but is not able to reflect on its own activities or its own knowledge. And because it cannot hold a conversation at the level of descriptions, it has no way of exploring students' misunderstandings or helping them to reach a shared understanding. Research in Intelligent Tutoring Systems has gone some way towards remedying this, but is still at the stage of hand crafting individual systems and a long way from developing a computer that can converse freely about its own knowledge.

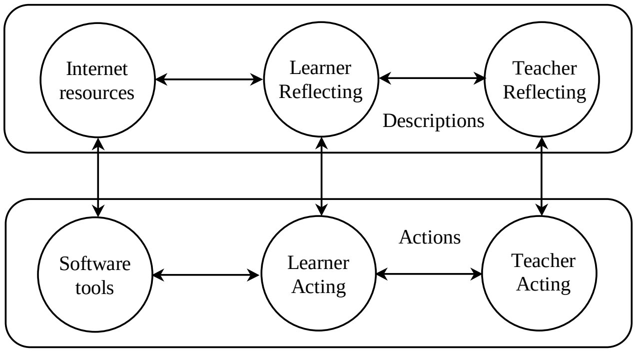

Figure 4. Technology for conversational learning

Alternatively, the technology could provide an environment in which conversational learning takes place, that enables conversations between human learners (Figure 4). Networked computers can extend the range of activities and the reach of a discussion, into other worlds through games, software models and simulations and to other parts of this world by using the computer as a means of communication, through phone, email and computer based discussions. Children are comfortable with this use of technology, since it is exactly how they use mobile phones, computer chat rooms and multiplayer games. The technology provides a shared conversational learning space.

The main implication of this Theory of Use for design of personal learning technologies is that it suggests a system image based on support for multiple conversations, as opposed to one derived from learning as the transfer of knowledge. Neither a “classroom” nor “library” metaphor evokes learning as conversation.

5. Field Studies

The field studies were based on interviews and questionnaires with 219 children aged 7 to 11 in two schools, designed to explore their learning habits, and likes and dislikes of computers and learning technology. The most relevant results from the survey are summarised below, with the figure indicating the percentage response to that question. Informal interviews were also carried out with teachers.

The sampled children preferred to work in a group (67%) rather than on their own, but they preferred to keep the products of their own work private (64%) rather than have it displayed to other children. Their preferences for seeking help with a problem were, in descending order, a friend (53%), a teacher (39%), a book (38%) and parents (37%). These findings suggest that personal technology should support group work, whilst enabling children to own the results. 44% of the children kept scrapbooks and (from an open response) the main contents mentioned were photos, newspaper cuttings and artwork.

The children’s greatest dislike of using computers was their slow speed. When asked in an open question what computers should do in the future, the most frequent response was that it would talk to them, followed by it being intelligent and having a personality. The children also drew pictures of what they thought that a personal handheld or wearable computer should look like. There were many instances of talking computers and computers with personalities; physical attributes such as arms, legs, eyes and ears were also commonplace. A selection of these is shown in Figure 5.

Figure 5. Children's Drawings of 'computers of the future'

6. Task Model

The Task Model draws together findings from the theory of use and field studies to form a composite picture of the activities and contexts that the technology might support and raises issues, including limitations, conflicts and breakdowns in traditional ways of working that technology could address.

The studies for HandLeR (including a related investigation of adult learners (Vavoula & Sharples, 2001)) have suggested the need for a system that can support personal learning projects across multiple contexts, including school, home and outdoors. It should offer facilities for capturing and annotating everyday events, relating new experience to stored knowledge through appropriate external representations, conversing with peers and experts, and managing personal learning projects.

The field studies raised an important issue regarding ownership of personal learning technologies and their data. One of the teachers suggested that the most useful application for a HandLeR would be the inclusion of a personal profile of the child. Different levels of the profile could exist to allow different people (teachers, children, and parents) to see different information. To the children, a HandLeR is a means to undertake personal learning projects, free of the influence of adults. To teachers, it can be a means of monitoring the children's progress in learning and gathering diagnostic information.

A HandLeR can be seen as a 'boundary object' (Star, 1989), an object or construct that is claimed by a group of actors with divergent viewpoints. A library, for example, provides a shared space and system of classification that can serve people's disparate needs. Alternatively, a boundary object may provide a common irritant, a garden hedge between two disputatious neighbours. Personal learning technologies, because they interpose between informal and formal contexts of learning, may become a focus for conflict between children and teachers (when children bring their personal computers in to the classroom) and between differing approaches to pedagogy (supporting individual learning projects or delivering a common curriculum). These issues cannot be relegated to ‘user preferences’ or ‘contexts or use’, since they lie at the heart of the system design. In designing HandLeR, we have taken the approach that the primary owner of the data (though not necessarily the physical device) is the learner and that the user profile is a means by which a learner can configure the software and create a sense of ownership of the system.

7. Design Concept

Whereas the Task Model provides an account of how people currently perform the required activities, the Design Concept envisions how these activities might be performed with the aid of new tools and technologies, to address limitations in current modes of working or offer new opportunities.

The Design Concept for HandLeR is informed by the Conversational Framework outlined in Section 4. The system distinguishes between operations at the levels of action and description. Operations for action enable a learner to capture and annotate events as images, sounds and written notes, to perform experiments and to converse with teachers and learners. Operations for description enable the learners to manage the captured learning events, relate them to previous actions, merge them with learning resources available on the web, and to create a composite map of current and previous knowledge. The two distinct modules promote a deliberate cycle of action and reflection.

7.1. Learning objects and idea map

The outcome of each captured event is stored as a “learning object” (LTSC, 2001). Learning objects are small items of learning material, such as a test, a simulation, or a lecture. They are generally stored in a structured XML format with metadata tags that describe attributes of the learning such as topic and prerequisites. The objects can be deposited in a bank of learning resources and then retrieved to satisfy a specific learning query, or they can be assembled into a teaching programme.

Learning objects were designed to provide reusable components of computer-based instruction, but the general approach, of storing, annotating and retrieving small items of learning applies equally to contextual learning. The metadata tags will need to be extended to include, for example, a more detailed definition of context of learning than the broad LTSC terms such as “Primary Education”. This would include the time and physical location of the learning, and possibly other elements such as co-learners present. Given sufficiently sophisticated technology such as GPS and wireless communication, some context information could be generated automatically. This would provide basic data to enable indexing and retrieval of the experiential learning objects.

The learning objects can be accessed and presented along any of the contextual dimensions: for example as a timeline, or a spatial map of where the learning events occurred. The system also provides the means to interact with the learning objects at the level of descriptions, by linking them by conceptual relation into an “idea map”, showing nodes representing learning objects connected by links indicating conceptual relations. This notation (and similar ones such as "concept maps" (Novak, Gowin, & Johansen, 1983) and "mind maps" (Buzan, 1989)) have been widely used in education to visualise and explore the conceptual relations between items of learning. The idea map might be generated automatically based on similarity between the type and content of the learning objects (for example, keyword matching of the content of notes). Generally, though, learners will create explicit links between objects, building up a personal map of their learning. Since the learning objects are XML entities, the map also enables a learner to create a personal website on a learning topic by creating, adding and linking learning objects.

7.2. System architecture

The system architecture provides a schematic of the main components of the system and their interconnections (see Figure 6). The Internet and HTTP protocol provides a common communications backbone for the system, allowing the components of the system to be physically distributed, so that a learner may operate a variety of mobile or fixed devices to access personal learning resources stored in a fixed server, or may load a copy of the resources onto a personal machine for learning in the field.

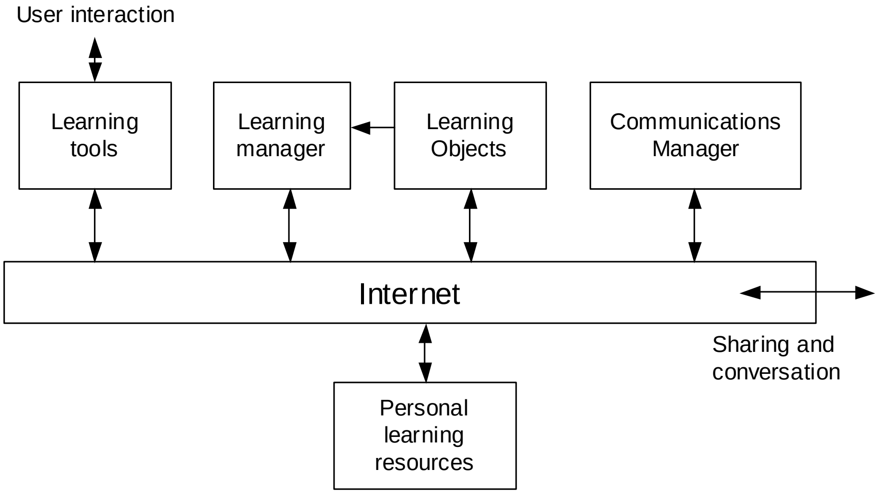

Figure 6. System Architecture

The user interacts with a set of *learning tools* – such as an integrated camera, notebook, sketchpad, and mobile phone – that enable the capture and annotation of learning events, the management of learning resources, and conversation with other learners and teachers. The *personal learning resources* consist of links to learning objects organised into an idea map, and also the user's learning profile, personal details, calendar and contacts list. The individual *learning objects* are XML pages that will typically be distributed across the Internet. Some will have been created by the learner, some by co-learners, and some will be provided as teaching resources. The metadata information must be sufficiently detailed to distinguish between accredited teaching materials and the results of an individual’s contextual learning, and also identify the owner and give permission to share or modify the objects. The *learning manager* stores a local cache of learning objects and deploys software agents to search for, filter and organise the objects. The *communications manager* creates direct voice and data communication with other learners and handles the sharing of resources.

7.3. System image

A central problem that the Design Concept must solve is how to present a coherent, appropriate and attractive system image. The only ubiquitous technology for learning is a notebook and pen. These could provide a basis for a system image, but a minimalist “blank notebook” is unlikely to appeal to young learners. Another possibility is a ‘virtual world’ metaphor that presents a simulated environment with familiar locations and objects to support learning, such as a “library”, “lab”, and “classroom. One fundamental problem with the virtual world metaphor is the representation of abstract items such as minds, concepts and ideas. It also imposes a traditional typology of teaching tools and locations that may preclude new contexts for learning.

An appropriate system image must explain itself and be instantly recognisable by any user. It should distinguish between activities at the level of action and description. The metaphor should be consistent across the system tools. It should also be timeless, unaffected by changes in fashion and working practice. The system image developed for the children’s HandLeR is that of an avatar with a humanoid body. It provides easily recognisable body parts that either act as direct links to the main tools (eyes for the camera, mouth for communication) or can hold recognisable icons (notebook, palette for painting). It also offers an animate agent that could act as a guide or mentor to the learner. The avatar could be any human-like image such as an image of a pop or sports star. For simplicity and general appeal we chose a cartoon rabbit for the initial prototype (Figure 7).

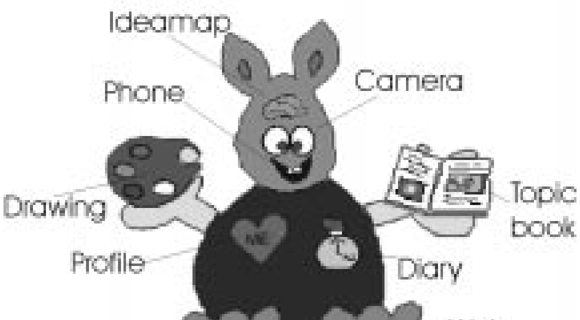

Figure 7. Rabbit avatar annotated with button actions

The intention was that the avatar should provide an intuitive means for the child to interact with HandLeR at the level of actions: capturing, annotating and communicating everyday events. It also provides a convenient interface object, the rabbit’s brain, to link to the idea map.

7.4. Physical Concept models

The physical appearance of the device is an important aspect of the design concept. The device should be light enough to be held in one outstretched hand. It should be possible to operate it on a flat surface or while carrying it. It should be capable of operation by left and right handed people. It should have simple controls to operate the tools such as camera or phone, and it should look appealing to the intended users. A physical concept model for pre-teen children was produced in solid foam (Figure 8). It is designed to resemble a video games controller, with curved edges for ease of holding with either hand, a touch panel screen, and two large buttons to operate the integral camera, sound recorder and communications.

Figure 8. Physical concept model of a children's HandLeR

8. Design space, specification and implementation

The Design Concept produces a general architecture and set of constraints on the design, but does not specify how these will be implemented. It leaves many possible options for implementing the system architecture and for designing the interface and user interaction.

The systems architecture shown in Figure 5 was implemented as a set of modules communicating via HTTP (Figure 9). A Java client handles the user interaction and manages copying of images and text between modules. It also provided custom tools for interacting with the "rabbit" avatar and the idea map. The personal learning resources, including node and link information for the idea map, are stored as an SQL database. A proxy web server retrieves individual learning objects, stored as local or Internet web pages. An FTP and HTTP server manages the transfer and sharing of the learning resources.

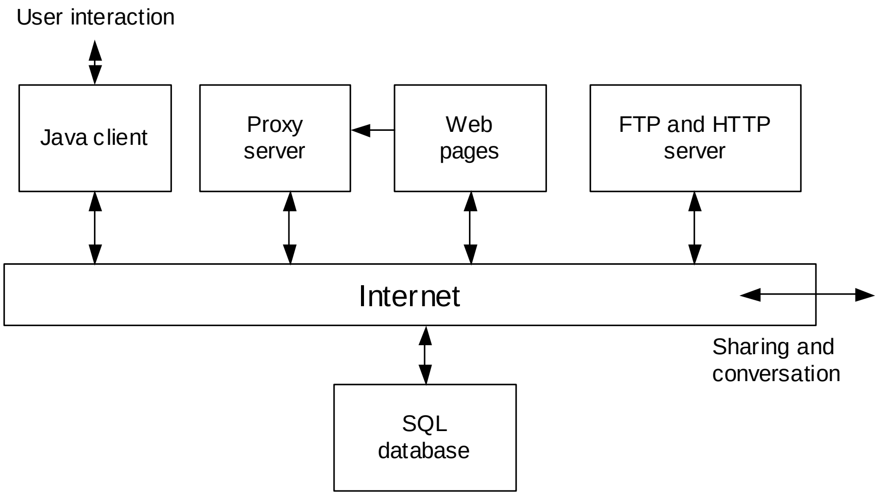

Figure 9. Implementation of the systems architecture

We employed the OVID (Object, View and Interaction Design) design methodology to explore and refine a space of possible designs and to implement the design choices. OVID has been developed by the IBM Ease of Use Group (Roberts, Berry, Isensee, & Mullaly, 1998) to bring rigour to the design of interactive systems and to provide a bridge between interface design and programming. The aim of OVID is to assist a designer to create interfaces that meet user requirements and are easy to use.

OVID fits within the more general methodology of socio-cognitive engineering, providing a principled approach to system specification and implementation. The main steps to designing with OVID are Requirements Analysis; Modelling; Design; Prototyping, Evaluation and Implementation. OVID supports iterative design, so any of these steps may be revisited to revise and refine the design. A detailed description of the design and implementation of HandLeR in OVID, and the benefits and limitations of the OVID method for designing novel interactive systems is given in (Corlett, 2000).

The Modelling stage is central to OVID. It is informed by the Analysis (within the context of socio-cognitive engineering this includes the General Requirements, Theory of Use, Field Studies and Task Model) and in turn it informs the rest of the design. Using a subset of the Unified Modelling Language (UML), OVID incrementally builds a model starting with the user's *perceived model* of his or her environment. A *designer's model* is then drawn up which identifies the task *objects* and the *views* which will be associated with them. Interaction and state diagrams with state tables are drawn to ensure completeness of the design. This model is finally augmented with the necessary system details to become the *implementer's model*.

We began the modelling by creating a space of possible designs in the form of storyboards, sketches and outline specifications. These were discussed in relation to the requirements and design concept, with designs being rejected and amalgamated until the process produced a small set of specifications for the system tools and operations. These main system tools for drawing, writing and communicating were modelled in terms of recognisable objects (such as pens, erasers and paintbrushes) and activities through single sentence task descriptions such as “child writes in topic book”.

An issue that arose during this stage was how to classify and group the tools. Some tools serve multiple functions; for example, the pad could be used for writing notes or drawing idea maps. It was decided to organise the tools according to their purpose in supporting learning, and to create a generic container “topic book” that that can be used to assemble and annotate notes, drawings and images.

The interface was designed to provide explicit support for conversation at the levels of action and reflection (see Figure 10). The main screen (Figure 10 a) shows the “rabbit” avatar and a series of buttons for help, search, adding items to the topic book and communicating with other devices. Applications selected by tapping on the avatar’s body parts appear on the right of the screen (Figure 10 b). These enable the user to capture images from the camera, browse web pages, and create and annotate images, notes and drawings in the topic book.

Figure 10. The main HandLeR screen, showing a) avatar and interaction buttons, b) annotating an image, c) idea map, d) communication link with another user.

Tapping on the avatar’s “brain” opens the idea map to organise, share and reflect on the captured learning objects (Figure 10 c). The screen shows (on the right) a history list of items created in the topic book and web pages browsed. The idea map was designed to be easy to operate on a tablet computer. Although the interaction is more constrained than other idea map tools such as Inspiration (Inspiration Software Inc., 2001), it only requires two basic operations (drag and tap) to build and explore the map. Dragging an item from the history list attaches it to the central box in the idea map. To move through the map, the user repeatedly taps on one of the outer boxes which brings it to the centre and shows the items linked to it. Tapping on the central box opens its associated learning object (web page or topic book item).

A tap on the icon in the lower right of the screen opens a dialogue box to connect to another person. Selecting from a list brings up the other user's avatar (Figure 10 d), which can be tapped to select the user's profile (the heart), or a direct voice connection (the mouth). An aim for a future version of the software is to show and capture a view from the other person's camera (by clicking on the avatar's eyes) and to share items from the idea map (by dragging them over the user's icon).

The main interface and idea map for the prototype system were implemented in Java, with calls to standard software applications to provide the main tools. The camera is controlled through the PictureWorks Live software and the topic book is Microsoft FrontPage Express. Each of these tools has its own interface which creates inconsistency. A more developed system would require custom-designed tools with consistent appearance and interaction.

The hardware for the prototype system consisted of:

- a Fujitsu Stylistic LT Pen Tablet computer containing a 8.4 inch 800 x 600 SVGA display with a touch sensitive screen and a 233MHz Intel Pentium processor running the Windows 98 operating system;
- a 3Com Home Connect camera in a custom-built bracket connected by a USB lead to the computer;
- a Lucent Technologies IEEE 802.11 standard PCMCIA card, providing wireless connection to a local area network with data rates of up to 11Mbs;
- a PCMCIA Nokia CardPhone, enabling direct voice communication from the computer to a mobile phone or to another computer, and data connection at 9.6Kbs.
-

Figure 11. Hardware for HandLeR prototype

Figure 11 shows the computer with camera and wireless LAN card. The Wireless LAN can either link the computer to a fixed network through a base station, or it can be configured in an ad hoc network that enables a group of people with handheld computers to exchange data at high speed up to a range of about 100 metres.

9. User Testing and Evaluation

The development method of socio-cognitive engineering is intended to promote testing throughout the system lifecycle. The General Requirements provide a set of constraints on the design process and criteria against which to validate the finished system. To produce the Task Model requires successful integration of the Theory of Use and Field Studies. If these conflict (for example if the theory of learning predicts learner behaviour that does not occur in practice) then further work is needed to reconcile them. The Design Concept can be validated by checking that it supports those activities and cognitive processes that were identified as being important or problematic in the Task Model (such as support of learning projects across multiple contexts). The OVID method specifies a series of tests at each step in the design and implementation sequence. Lastly, the Task Model provides a reference point to compare the way tasks are performed with existing technology against how it can be carried out with the new tools. The methods that were used to study and assess the original tasks can be used to compare with the transformed activities.

The prototype system largely, but not entirely, satisfies the General Requirements. The hardware is *portable*, in a single package that can be carried in one hand. It is easy to direct the camera and capture an image while watching the computer screen, though the display can be difficult to see in bright sunlight. The pen tablet can be operated by cradling the machine on one hand and tapping or writing with the other (Figure 12), but the weight of the device, 1.5 Kg, meant that some children sat or squatted and balanced it on a knee. Hardware based on a PDA such as the Compaq Ipaq would be more portable, but would sacrifice the 800x600 SVGA screen for a 240 x 320 display. Further work is needed to determine whether a display of that size can support useful learning, but digital cameras and idea map tools (Soloway et al., 2001) are already available for PDA devices.

Figure 12. A child interacting with the HandLeR device

As an *individual* learning resource, the HandLeR software is intended to support a wide range of learning activities and abilities, rather than adapt to specific learners. It can support children with differing approaches to learning, for example by providing pre-prepared resources and idea maps or enabling learners to construct their own.

The system is relatively *unobtrusive* and can be used in settings (such as field trips or informal conversations) where a laptop computer and camera would intrude on the activity. A particular advantage is being able to capture an image without having to hold the device in front of the face.

Communication is *available* anywhere within range of a cellular telephone system and the interface enables a call to be made by a single tap on the "mouth" of the other party's avatar. It has proved difficult to adjust the audio settings to enable good quality voice conversation, but an acceptable conversation in relatively quiet surroundings can be carried out with the built-in microphone and speaker. Data communication by mobile phone at 9600 baud is too slow for internet access or transfer of images. With the devices in peer communication through the wireless LAN cards images and data can be transferred rapidly up to approximately 100 metres line of site. Indoors or obstructed by buildings, the effective range of communication drops to about 10 metres.

The current version is not *adaptable* to the learner's changing needs and abilities, nor is it designed to be *persistent*. We have research in progress to develop mobile technology for long-term learning projects (Vavoula & Sharples, 2001)

The prototype system was evaluated for *ease of use* and *usefulness* through a combination of methods. These included: videotaped observations of three 11-year-old children using an early prototype; usability questionnaires with 29 children aged 10; and a day trial of the system, using an earlier hardware configuration based on a larger Fujitsu Stylistic 2300 tablet computer, with six 11-year-old children and their teacher carrying out a guided activity to explore canals in central Birmingham (this was videotaped by the BBC for a Tomorrow's World science programme).

The children in the observation study successfully completed the tasks set for them such as creating a topic book, making a drawing, capturing an image and a movie, and setting up a phone call from the HandLeR to a mobile phone. Although they had suggestions to improve the interface, they found the layout and main functions easy to operate.

For the questionnaire study, the children were given printouts of the main HandLeR interface and asked to identify the functions associated with parts of the avatar. Each function was correctly identified by more than 50% of the children, apart from the feet (for internet access) and the watch (for a diary). In general, the children said that the design of the avatar was attractive and colourful but some of them commented that the cartoon rabbit was too "babyish". The children suggested a variety of alternatives, including animals, aliens, robots and cyborgs. We can conclude that if an avatar is used as the means of representing the user on the screen, then the user should be able to choose from a suite of images, designed in collaboration with children and appropriate for different ages. A more fundamental problem is that although the avatar metaphor could be extended to provide animated help and guidance, it could easily become contrived and overloaded with tools. Alternatives should be considered, such as the Pad interface based on a camcorder metaphor of pan and zoom (Perlin & Fox, 1993).

For the day field trial, the children were divided into two groups of three and presented with a “mission” on the HandLeR screen, which was to explore the canals in central Birmingham, answer two questions, and return with evidence to support their answers. For one team, the questions were “What were canal boat used for in the 1850s?” and “How were the boats powered in the 1850s?” The other team were asked similar questions about modern-day canal boats. Visual evidence that they could collect from exploring their surroundings included disused warehouses and canal bridges with notches caused by tow ropes, as well as working boats and water bus signs. They could also refer to pre-cached web pages that were linked to the HandLeR idea map. The groups were encouraged to converse via the mobile phone link and share information.

The field trial was successful, in that the children accomplished their tasks, despite having to perform for the BBC camera crew. They were able to navigate the topic map to find background information, make notes, capture still and moving images as evidence, and hold voice conversations between the devices. The main difficulty was in handwriting recognition. The software, PenX V1.66 by Communication Intelligence Corporation, did not recognise cursive script, and the children found it impossible to balance the tablet on one hand and write with the other. By sitting down and resting the pad on their lap, they were able to make short notes. Other problems with the prototype hardware included its weight and short battery life.

10.Conclusions

A main conclusion from the trials is that further evaluation of the usefulness of mobile technology for learning should be delayed until technology has been developed that is easy and intuitive to use. The prime purpose of the system described in this paper is to augment the quality and scope of personal learning, not to deliver instruction, so it cannot simply be assessed by measuring prescribed learning gains. As with personal management software such as diaries and meeting organisers, the benefits come from enabling people to manage their lives more effectively. A successful learning organiser should fit into the daily activities of informal learning. It should allow people to capture and recall an object or event that they would otherwise forget, integrate disparate sources of information into coherent schemas, assist in performing experiments and solving problems in the everyday world, and augment conversations by providing a way to exchange and share relevant information. We have projects underway to investigate these aspects of informal learning and develop a generic systems architecture and interface for personal learning organisers.

Another finding from the current project is the importance and complexity of context. Mobile learning is more strongly mediated by its context than classroom instruction. Context involves the familiar dimensions of time and location, but also includes the learner’s trajectory, goals and motivation, the surrounding resources, co-learners, and other available conversants. Consider, for example, a visitor to a museum or historic site. Their current learning goals and activities will depend, *inter alia*, on where they are located and the time of day (the museum may offer scheduled events), the route they have taken and what they have already seen on the way, their interests and motivation to learn, whether the museum objects are accompanied by captions or teaching material, and whether a more knowledgeable person is near and available. A mobile learning resource should, ideally, fit seamlessly into this complex pattern of learning opportunities and resources.

The project described in this paper has developed a proof of concept for mobile learning, informed by a relevant theory of learning and field tested with 11-year-old children. It has shown that a handheld or wearable device, with appropriate learning tools and resources, an intuitive interface and high-speed communication could offer a new generation of portable “learning organisers” for people of all ages.

References

Abowd, G. (1999). Classroom2000: an experiment with the instrumentation of a living educational environment. *IBM Systems Journal, 38,* 508-530.

Brown, A., & Campione, J. (1996). Psychological theory and design of innovative learning environments: On procedures, principles, and systems. In L. Schauble & R. Glaser (Eds.), *Innovations in Learning: New Environments for Education* (pp. 289-325). Mahwah, N.J.: Erlbaum.

Buzan, T. (1989). *Use Your Head* (Revised ed.). London: BBC Books.

Corlett, D. (2000). Innovating with OVID. *Interactions, 7,* 19-26.

Druin, A., Stewart, J., Proft, D., Bederson, B. B., & Holland, J. D. (1997). KidPad: A Design Collaboration Between Children, Technologists, and Educators, *Proceedings of CHI '97* (pp. 463-470). Atlanta, Georgia: ACM/Addison-Wesley.

Engeström, Y. (1987). *Learning by expanding: An activity-theoretical approach to developmental research*. Helsinki: Orienta-Konsultit.

Eraut, M. (1994). *Developing Professional Knowledge and Competence*. London: Falmer Press.

Fischer, G., & Scharff, E. (1998). Learning Technologies in Support of Self-Directed Learning. *Journal of Interactive Media in Education, 98*(4).

Inkpen, K. M. (2000). Designing Handheld Technologies for Kids. *Personal Technologies, 3*(1&2), 81-89.

Inspiration Software Inc. (2001). *Inspiration: visual thinking and learning software, http://www.inspiration.com/.* Available: http://www.inspiration.com/.

Kolb, D. (1984). *Experiential Learning*. Englewood Cliffs, New Jersey: Prentice Hall.

LTSC. (2001). *IEEE Learning Technology Standards Committee website*. Available: http://ltsc.ieee.org.

Norman, D. A. (1986). Cognitive Engineering. In D. A. Norman & S. W. Draper (Eds.), *User Centred System Design* (pp. 31-61). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Novak, J. D., Gowin, D. B., & Johansen, C. T. (1983). The use of concept mapping and knowledge vee mapping with junior high school teachers. *Science Education*(67), 625 – 645.

Oosterholt, R., Kusano, M., & de vries, G. (1996). Interaction design and human factors support in the development of a personal communicator for children, *Proceedings of CHI '96* (pp. 450-457): ACM/ Addison Wesley.

Pask, A. G. S. (1976). *Conversation Theory: Applications in Education and Epistemology*. Amsterdam and New York: Elsevier.

Perlin, K., & Fox, D. (1993). Pad: An Alternative Approach to the Computer Interface, *Proceedings of SIGGRAPH '93* (Vol. 93, pp. 57-64). New York, NY.

Roberts, D., Berry, D., Isensee, S., & Mullaly, J. (1998). *Designing for the User with Ovid: Bridging User Interface Design and Software Engineering*. London: Macmillan Technical Publishing.

Roth, W.-M. (2000). Learning environments research, lifeworld analysis, and solidarity in practice. *Learning Environments Research, 2*(3), 225-247.

Ryan, N. S., Morse, D. R., & Pascoe, J. (1999). *FieldNote: a Handheld Information System for the Field.* Paper presented at the TeleGeo'99, 1st International Workshop on TeleGeoProcessing, Lyon.

Sharples, M. (2000). The Design of Personal Mobile Technologies for Lifelong Learning. *Computers and Education, 34*(177-193).

Sharples, M., Goodlet, J., & Pemberton, L. (1992). Developing a Writer's Assistant. In J. Hartley (Ed.), *Technology and Writing: Readings in the Psychology of Written Communication* (pp. 209–220). London: Jessica Kingsley.

Sharples, M., Jeffery, N., du Boulay, J. B. H., Teather, D., Teather, B., & du Boulay, G. H. (2000a). Socio-cognitive engineering: a methodology for the design of human-centred technology. *European Journal of Operational Research, In press*.

Sharples, M., Jeffery, N., Teather, D., Teather, B., & du Boulay, G. (1997). A Socio-Cognitive Engineering Approach to the Development of a Knowledge-based Training System for Neuroradiology, *World Conference on Artificial Intelligence in Education (AI-ED '97)* (pp. 402 – 409). Kobe, Japan.

Sharples, M., Jeffery, N. P., du Boulay, B., Teather, B. A., Teather, D., & du Boulay, G. H. (2000b). Structured Computer-based Training and Decision Support in the Interpretation of Neuroradiological Images. *International Journal of Medical Informatics, 60.*(30), 263-228.

Soloway, E., Norris, C., Blumenfield, P., Fishman, B., Krajcik, J., & Marx, R. (2001). Handheld Devices are Ready at Hand. *Communications of the ACM, 44*(6), 15-20.

Star, S. L. (1989). The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In L. Gasser & M. N. Huhns (Eds.), *Distributed Artificial Intelligence* (pp. 37-54). London: Pitman.

Vavoula, G. N., & Sharples, M. (2001). Studying the Learning Practice: Implications for the Design of a Lifelong Learning Support System. In T. Okamoto & R. Hartley & Kinshuk & J. P. Klus (Eds.), *Proceedings of IEEE International Conference on Advanced Learning Technologies (ICALT 2001)* (pp. 379-380). Madison, USA: IEEE Computer Society.

Walker, L., Rockman, S., & Chessler, M. (2000). *A More Complex Picture: Laptop Use and Impact in the Context of Changing Home and School Access* (Third in a Series of Research Studies on Microsoft's Anytime Anywhere Learning Program). San Francisco: Rockman et al.