Observations on Pervasive Information Systems Design

Mark J. Weal[1], Danius T. Michaelides[1], David E. Millard[1], David C. De Roure[1] and Geraldine Fitzpatrick[2]

[1] Pervasive Systems Centre, University of Southampton, UK
[2] Interact Lab, Dept. of Informatics, University of Sussex, UK

Abstract. This paper seeks to lay out a number of observations on Pervasive information systems design that have been made during the Chawton House project, an augmented field trip for school children. After laying out the Chawton House case study, the information infrastructure is described and arising issues are examined concerning design for persistence, empowerment of non technical users, and re-use.

1 Introduction

Many pervasive applications are very much about information—collecting it, using it, delivering it and interacting with it. In this paper we acknowledge this by adopting an information perspective on pervasive application design. The potential benefit of this approach is that we can maximise the re-use of information in all phases and from application to application.Instead of a pervasive experience being a one-shot endeavour, we envisage a world with multiple, ongoing pervasive deployments, where there is an accumulation of information from multiple user engagements, multiple deployments and multiple designs.

For our case study we have been working with different groups (primarily the curators of a country house and a group of teachers) to build location based pervasive experiences on top of a common infrastructure and information model. This enables us to look at how experience builders could construct their own pervasive experiences with less intervention from system experts as well as exploring how multiple pervasive experiences can co-exist on a single architecture. In doing this we have been able to use Semantic Web technologies to facilitate a wide variety of re-use at all phases of the information lifecycle.

2 The Chawton House Case Study

In order to move towards a generic infrastructure to support a wide range of activities, we have been working with the curators of Chawton House Library, an Elizabethan manor house that is now owned by a charitable foundation that operates it as a study centre of early English women's writing.

We have worked with the curators to produce a visitor system in the tradition of locative information systems [2, 1]. In addition, we have also been working

with teachers from a local school in Hampshire to create an augmented field trip for a group of Year 5 students (aged 10-11) using the same underlying infrastructure and information model. Other learning experiences for children using locative, context-based technologies have focused on scientific discovery [10], learning about nature [3], learning about history, etc. Here, the main emphasis of the school experience has been on using the landscape as a writing aide.

The literacy experience applied the infrastructure framework to the creation of a field trip for a group of children with the aim of steering them through the gathering of material for them to later use when writing a fictional story. The teachers devised a scenario where at various locations around the grounds the children would be given a range of different activities to carry out. Some of these activities required them to write in their exercise books while others involved the recording of text and audio on the PDAs.

The children gathered information through listening to audio prepared by the curators, following instructions on the PDA set out by the teachers and observing, describing and speculating about the landscape, flora and fauna around them. Having completed the first phase the children were gathered together to briefly share their findings with each other before moving on to the second phase, where they went back to two locations of their choice and received further instructions on more specific activities to help in the creation of their stories.

3 The Information Infrastructure

RDF was created as a framework for metadata to provide interoperability across applications that exchange machine-understandable information on the Web. It has a very simple relational model which accommodates structured and semi-structured data, and in fact can be seen as a universal format for data on the Web, providing greater interoperability and re-use than XML alone.

RDF was created as a framework for metadata to provide interoperability across applications that exchange machine-understandable information primarily on the Web. RDF is also very useful within pervasive infrastructures. The museum experience described in [6] is a very good example of the use of semantic descriptions in a real (and real-time) application. It uses inference rules alongside representation of user models and content descriptions, and involves several ontologies. The 'Semantic Smart Laboratory' work [7] also uses RDF from the very first stage of capturing the activities of chemists working in a laboratory, as well as a sensor network to capture laboratory environmental conditions. This is used as the first stage in establishing a complete provenance trail through to scholarly output, enabling researchers to chase back to the original data.

A variety of notions of context may be expressed, including location and user tasks [9, 11]. Ontologies can also be used to describe device capabilities, for example to facilitate content delivery to devices with diverse characteristics [8]. In this project we use geographical location to link together semantic annotations and to describe content, and we have developed a simple ontology for record and re-use which provides a framework to facilitate the capture of pervasive comput-

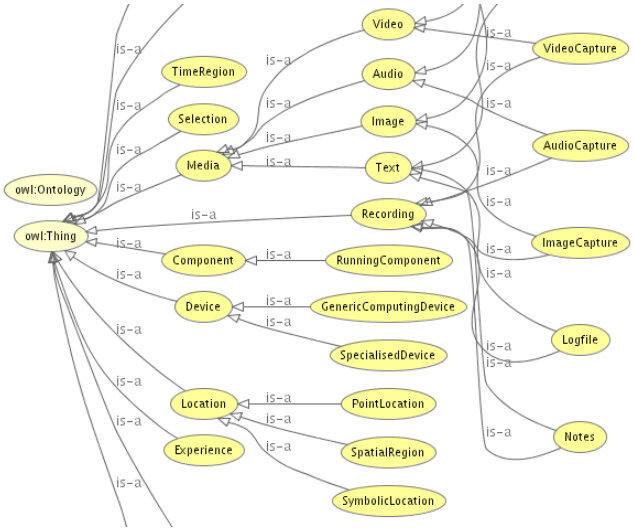

Fig. 1. Part of the Record and Re-use ontology graph.

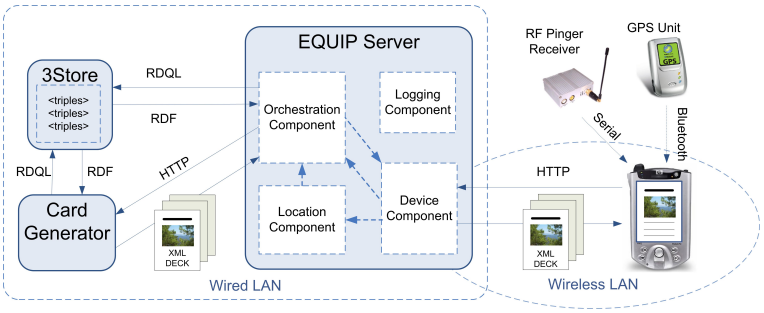

Fig. 2. The system architecture.

ing experiences and experiments, Figure 1 shows part of the graph. The base ontology defines a number of classes of objects that represent 'things' used in an experience — people, locations, devices, artefacts, software components, content, recordings, annotations — with associated relationships between them. The intention is that a base ontology can be used in conjunction with an experience specific ontology to give more appropriate information model.

Figure 2 shows a simplified overview of the architecture. The pervasive infrastructure comprised at the base level: wireless networking, location sensing technology (GPS and RF pingers) as well as PDAs carried by the children. On top of this infrastructure sat the generic information system, an orchestration tool (constructed using EQUIP, the EQUator Infrastructure Platform) alongside a triple store (the AKT 3store [5]), an RDF storage and query engine developed to support the handling of large RDF knowledge bases.

```
<SpatialRegion rdf:about="http://www.chawton.org/locations#graveyard">
    <rdf:type rdf:resource="http://equator.ecs.soton.ac.uk/rnr/ontology#SymbolicLocation"/>
    <frame rdf:resource="http://nurl.org/0/geography/SRSCatalog/wgs84"/>
    <has-description>The graveyard of the church</has-description>
    <has-pretty-name>The Graveyard</has-pretty-name>
    <points>51.076730,-0.594024 51.076730,-0.594024 51.076791,-0.593866
            51.076861,-0.593833 51.076870,-0.593830 51.076877,-0.593833
            %51.076881,-0.593839 51.076883,-0.593849 51.076883,-0.593849
            %51.076881,-0.593876 51.076869,-0.593970 51.076869,-0.59397
            51.076868,-0.594015 51.076866,-0.594022 51.076858,-0.594035
            51.076845,-0.594045 51.076845,-0.594045 51.076845,-0.594045
            51.076836,-0.594044</points>
</SpatialRegion>

<SymbolicLocation rdf:about="http://www.chawton.org/locations#libraryTerrace">
    <has-description>The terrace outside the library.</has-description>
    <has-pretty-name>Library Terrace</has-pretty-name>
</SymbolicLocation>
```

Fig. 3. RDF Representations of a spatial and symbolic location.

3.1 Location and context

The base ontology defines three types of location — symbolic, point and spatial region. These form a rudimentary method of conversing about location in terms of named resources. Figure 3 shows the RDF representations of a spatial region and of a symbolic location. Pinger locations were tied to symbolic locations with the spatial regions used to map GPS positions to named lcoations.

Many spatial annotation efforts are emerging. For example, accumulation of annotations in a spatial region is the basis of the OpenGuides "WIKI" city guides (http://openguides.org/). The Basic Geo vocabulary is used in Locative Packets for spatial annotation (locative.net) and the Open Geospatial Consortium (OGC) pursues standards for a number of geospatial and location based services (http://www.opengeospatial.org/).

3.2 Media content, annotations and metadata

All media in the system is marked up in the base ontology which describes the different types of media files. Media appears in two forms in an experience; as content and as recordings of the experience. Recordings in one experience can easily be re-used as content in another because they share the common media concepts, whilst maintaining the context and annotations associated with the original capture of the material.

A card metaphor was used to construct the user's main experience of the Chawton system. This model was used for orchestration of the pervasive experience by maintaining a user context, or system state, and then deciding which of the available *cards* held by the system fit with the current context. In the case of the Chawton project trials the context included the participants location as well as a global clock which controlled the overall experience. When a card is experienced by the participant, the user context can change as a result of the assertions held on the card. This allows us to build in things such as pre-requisites

(when a participant has to experience certain cards before they are able to access others). Different types of cards had different functionality associated with them, for instance a simple *information card* might have a piece of audio attached to it to be played automatically, or a *capture card* might ask the children to record a short description. Cards can then be arranged into sequences and decks for particular activities [13].

Annotations provide the key information representation used in analysis and post-experience re-use. The annotation structure enables statements to be constructed about other 'things' in the experience. Annotations maybe derived from log data or maybe authored manually by experience runners or analysts. For example, network activity from log data may indicate areas or periods of poor network activity. An experience runner may have made a note of when the system did not respond at all, or an analyst viewing a video of the experience may notice users not being delivered the correct information in a timely manner.

The final element of the information model is experience metadata that captures the running of an experience, such as names, locations, timing, and scope, as well as the more complex relations between objects that make up the configuration and capture of the experience. The experience metadata ties all the resources together allowing the experience to be packaged up. This can forms a complete record for archival purposes or it can be re-used, for example an identical experience could be rerun by discarding the recorded information and reinstating the system configuration and content.

4 Designing for persistence

One of the greatest challenges for pervasive systems is to have an extended lifetime, beyond a single trial, when the system is unlikely to be supported by experts and must be maintained by users and traditional support staff. At Chawton our intention was to create a system that could grow with use, and would be simple enough for curators and teachers to use for repeatable events.

A significant issue was making the system resilient to hardware and software failures. For example, in the Chawton school trip experience a number of the information cards involved the playback of audio files, but the audio drivers on the older PDAs that we were using were unstable, and in certain conditions could halt the machine. During our trial events we were forced to restart the application when this happened, although the teachers quickly learnt the process for themselves, however the disruption was minimised as the loss of state was minimal; upon restarting the application would obtain the cards relevant to its location and restart any sequences. Because most of the users' history (state) was held on the server, the only loss of information was the children's position in any sequences, and because sequences were short the disruption was minimal.

A persistent infrastructure must also support evolving applications and uses. This means that the information in the system should grow with its use, and infrastructure should be extensible. In Chawton the initial experiences were constructed by the Curators, later the Teachers took the tours that the Curators

had made (and listened to the resources created) and were inspired to build their own experience. Not only did the material motivate them to reuse the application and the landscape, but they could directly draw on some of the resources for their own experience, incorporating stories about the different areas into their tasks and roleplays. In this way a multi-layered experience is built up, with many experiences existing in the same physical and information space.

In summary there were three principles that we felt facilitated and encouraged persistence: coping with failure (enabling restarts), supporting evolving information, and allowing multiple experiences based on the same information.

5 Designing for empowerment

Related to persistence are the issues of control, ownership and empowerment. If a pervasive system is to move from the control of the technical experts who created it to the users and domain experts who are using it, then they must be allowed to take ownership. This requires familiarity with the way in which information is translated into an experience. In Chawton we used a card-based metaphor to accomplish this. Cards are an excellent way of emphasising the non-linear nature of a location-based system, the atomic requirement for content, as well as the constraints of a small display. To cope with prerequisites we introduced the notion of card sequences, but these didn't effect the other characteristics (for example, just because a card is in a sequence doesn't mean it ceases to be atomic - sensible and reusable in its own right).

A good metaphor also allows users to design an experience independently of the implementation, there is no need for them to understand either the underlying information representation or the infrastructure. The metaphor allows a clean separation of the information that *constitutes* the experience, from the information that *enables* the experience. Not only is the former more easily grasped by non-technical users, but it is the part most in need of their domain expertise, and thus the part that they are most likely to take ownership of.

Metaphors only go so far towards empowering users, tools that allow the metaphorical elements to be authored and maintained are also important. For Chawton we developed a number of authoring tools that further emphasised the content of the experience, for example we developed in-situ authoring tools that placed the author in the context of the eventual user [12]. Thus activities could be inspired directly by the landscape and the Teachers could discuss and refine their ideas unhindered by unfamiliar technology.

In summary, three principles helped to empower our Chawton users: the use of metaphors, the separation of the information layer, and the development of lightweight (in-situ) authoring tools.

6 Designing for re-use

For a pervasive system to become naturalised it is not only necessary for it to be persistent and familiar, it must also grow and change with the people and

the space in which it is deployed. Re-use is the building block to enable this; if experiences and resources are re-usable, then it becomes possible to build on past efforts and over time create more complex and sophisticated experiences. In the Chawton infrastructure we have enabled this through the use of Semantic Web based models and our record and re-use ontology

Re-use within the Chawton system takes a number of forms. During the authoring phase the teachers were able to make use of existing locations created for the visitor system by the curators. They could also appropriate anecdotes produced by the curators and use them to provide information to the children appropriate to the various locations around the house and grounds. For the first trials the experience was created as part of a co-design activity [4]. By the second trials the teachers were re-using and updating the content from the first trials using the newly developed authoring tools.

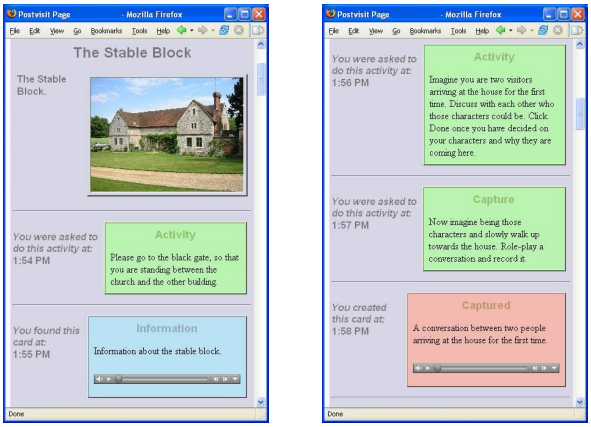

Fig. 4. The replay interface used by the children in the school

The semantic logging allowed the information gathered during the trials to be used in a number of ways after the day of the trials. Two of these were the use by the children back at the school in their writing activities and the use by the researchers in analysing the data gathered during the trials.

When the children returned to the school they were able to use the material they had collected to help them create their fictional stories. The logs of their activities were parsed and used to generate personal journal web pages that presented them with where they had been, what audio clips they had heard and their own personal recordings. Figure 4 show excerpts from one of these journals.

The journal approach, used in earlier projects [14] is made easier by the well structured activity logs and the ease of parsing of the semantic structures into simple web pages. By using the journals the children were able to revisit their activity and refresh their knowledge. As one of the teachers put it, 'They all grabbed little bits of the House yesterday.'

As well as the short term re-use post trial, a number of long term re-use activities are envisaged. Projects that 'instrument' museums and other public spaces are often heavyweight research efforts that are one-offs, relying on a team of skilled developers to carry out any maintenance or change to the content. Again, we see persistence as a crucial issue; there need to be ways that technology can remain in situ, at least partly maintained or changed by its users. To this end, we are continuing to develop specific tools to allow the curators to perform in-situ authoring to continue to increase the body of content available as part of both the visitor experience and for repurposing for other activities.

By archiving detailed logs in a well structured way we are able to make these logs available to other researchers who may wish to carry out analysis on the activities from different perspectives. For instance someone might specifically be interested in how the GPS reception changed according to the speed of movement of the children. This information, although not utilised in this particular experience, is available and easily accessible through the heavily annotated logs.

In summary we believe that any pervasive system should facilitate reuse for authoring, reflection, analysis and archiving.

7 Conclusions

The landscape of pervasive computing research includes many projects which have focused on specific devices or specific deployments. There is also a good tradition of focusing on the end-user of the devices for user requirements, trials and evaluation. In our work we endeavour to take a more holistic view, looking at all users engaged with it, and looking towards a world with multiple, ongoing pervasive deployments.

Our trials have enabled us to look in more detail at the specific requirements of an information infrastructure designed to support a wide range of experiences in a persistent fashion. We have been able construct an ontology upon which to build information systems that can be re-used both during and post experience by participants and researchers alike. Having consolidated our information models we have been able to place the tools to create such experiences into the hands of the domain experts, in our case, both the curators and teachers involved with the project. In this way we can move towards a persistent infrastructure to support the creation of multiple, potentially overlapping experiences that draw on the same annotated information repository.

Acknowledgments

The authors wish to acknowledge their partners: Chawton House Library, Whiteley Primary School and the University of Bristol. EPSRC IRC project "EQUATOR" GR/N15986/01.

References

1. G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton. Cyberguide: a mobile context-aware tour guide. *Wireless Networks*, 3(5):421–433, Oct. 1997.
2. K. Cheverst, N. Davies, K. Mitchell, and A. Friday. Experiences of developing and deploying a context-aware tourist guide: the guide project. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking, Boston, Massachusetts, United States*, pages 20–31. ACM Press, 2000.
3. W. Grant. Wireless coyote: A computer-supported field trip. *Communications of the ACM*, 36(2):57–59, Feb. 1993.
4. J. Halloran, E. Hornecker, G. Fitzpatrick, D. E. Millard, and M. J. Weal. The chawton house experience - augmenting the grounds of a historic manor house. In *Workshop - Re-Thinking Technology in Museums: Towards a New Understanding of People's Experience in Museums, Limerick (Ireland)*, pages 54–65, June 2005.
5. S. Harris and N. Gibbins. 3store: Efficient bulk rdf storage. In *Proceedings of the 1st International Workshop on Practical and Scalable Semantic Systems (PSSS'03), Sanibel Island, Florida, USA*, pages 1–15, 2003.
6. M. Hatala, R. Wakkary, and L. Kalantari. Rules and ontologies in support of real-time ubiquitous application. *Journal of Web Semantics*, 3, 2005.
7. G. Hughes, H. Mills, D. De Roure, J. Frey, L. Moreau, m. c. schraefel, G. Smith, and E. Zaluska. The semantic smart laboratory: a system for supporting the chemical escientist. *Organic and Biomolecular Chemistry*, pages 1–10, 2004. DOI: 10.1039/b410075a.
8. O. Lassila and M. Adler. Semantic gadgets: Ubiquitous computing meets the semantic web. In D. Fensel, J. A. Hendler, H. Lieberman, and W. Wahlster, editors, *Spinning the Semantic Web*, pages 363–376. MIT Press, 2003.
9. I. Millard, D. D. Roure, and N. Shadbolt. Contextually aware information delivery in pervasive computing environments. In *Lecture Notes in Computer Science, Volume 3479*, pages 189 – 197, Jan. 2005.
10. Y. Rogers, S. Price, D. S. Fraser, C. Randell, M. J. Weal, and G. Fitzpatrick. Ubi-learning Integrates indoor and outdoor experiences. *Communications of the ACM*, 48(1):55–59, Jan. 2005.
11. T. Strang and C. Linnhoff-Popien. A context modeling survey. In *First International Workshop on Advanced Context Modelling, Reasoning And Management, held in conjunction with Ubicomp 2004*, Sept. 2004.
12. M. J. Weal, D. T. Michaelides, D. E. Millard, , D. C. D. Roure, E. Hornecker, J. Halloran, and G. Fitzpatrick. Requirements for in-situ authoring of location based experiences. In *Proceedings of 8th ACM International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI 06)*, pages 121–128, 2006.
13. M. J. Weal, D. T. Michaelides, D. E. Millard, , D. C. D. Roure, K. Howland, and G. Fitzpatrick. Card based metaphor for organising pervasive educational experiences. In *Proceedings of the 3rd IEEE International Workshop on PervasivE Learning.*, pages 165–170, 2007.
14. M. J. Weal, D. T. Michaelides, M. K. Thompson, and D. C. D. Roure. The ambient wood journals - replaying the experience. In *Proceedings of Hypertext'03, The fourteenth conference on Hypertext and Hypermedia.*, pages 20–27, 2003.