

Automated Murmurs: The Social Mobile Tourist Application

Jill Freyne*, Adam J. Brennan[†], Barry Smyth[†], Daragh Byrne[‡], Alan F. Smeaton[‡], Gareth J.F. Jones[§]

*CSIRO ICT Center

GPO Box 1538, Hobart 7001

TAS, Australia

Email: jill.freyne@CSIRO.au

[†]CLARITY: Centre for Sensor Web Technologies

University College Dublin

Belfield, Dublin 4, Ireland

Email: adam.brennan@ucdconnect.ie, barry.smyth@ucd.ie

[‡]CLARITY: Centre for Sensor Web Technologies

Dublin City University

Dublin, Ireland

Email: daragh.byrne@computing.dcu.ie, alan.smeaton@dcu.ie

[§]Centre for Digital Video Processing

Dublin City University

Dublin, Ireland

gareth.jones@computing.dcu.ie

Abstract—The popularity of mobile devices and their increased computing power has given rise to surge in mobile computing technologies. Users are increasingly turning to mobile devices for information relating to their activities and location while on the move. Independent of this, the world has seen a huge uptake in the social web, which has fueled the production of applications where users are the sole providers of valuable information. In this work we present a mobile platform which leverages the popularity of mobile and social computing to produce a location sensitive messaging system which delivers user generated content to the public in the context of their physical location.

I. INTRODUCTION

The now familiar *social web* reflects an important change in the nature of the web and its content. The arrival of blogs in 1999, as a simple way for users to express their views and opinions, ushered in this new era of *user-generated content* (UGC). Many sites quickly began to offer a whole host of UGC alternatives including the ability to leave comments, write reviews, as well as the ability to rate or vote on the comments/opinions of others. The result has been an increased emphasis on people rather than content and, in combination with social networking services, this has precipitated the growth of the *social web* as a platform for communication, sharing, and collaboration.

There is now evidence of the emergence of a new so-called *sensor web*, enabled by the arrival of cheap and reliable sensor technologies, which provide new opportunities for users and applications to interact with the physical world. In short, the sensor web provides a way to bridge the *digital-physical divide* and so enable new types of information services in the physical

world, services which are capable of offering features that were previously only available in the digital domain.

For example, location-based sensing technologies are now becoming commonplace, making it possible to develop new types of location-aware information services. For instance, *digital graffiti* type applications allow users to leave digital messages (text, picture, video, audio) at a physical location, where they can be accessed by nearby users [3]. This type of service provides a new form of mobile messaging and acts as a platform for a range of innovative location-aware services.

In this paper we present one such mobile, location-aware messaging platform, suitable for existing mobile technology, and which has capitalized on the willingness of users to seek out new ways to express themselves, share information, and engage in communities. We detail our general architecture for the Context Aware Mobile Messaging (Camm) platform, an indoor/outdoor capable architecture that targets content to users based on location, time, and social features. This platform permits users to leave a variety of message types at arbitrary locations, allows users to define *context triggers* including location, presence and time to control the activation of these messages in the presence of others. As such Camm allows users to interact with their physical space in new and interesting ways. We present a real world implementation of the Camm architecture in the context of a specific a tourist application and outline a number of novel social communication features that it supports.

The rest of this paper is organised as follows. In the next section we present some of the background technologies relevant to our work both in the mobile domain and the social

web. We then describe our context-sensitive messaging system focussing in particular on a single implementation *Automated Murmurs* a context aware, social, tourist information system.

II. BACKGROUND WORK

Context awareness or sensitivity is defined as a knowledge of situational information. Dey and Aboud, [1], define context as any information that can be used to characterise the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and the application themselves. In context aware systems, situational information is used in order to adapt an application or system using event-based or environmental triggers detected through the use of sensors. The “Forget-me-not” system [11] presented ideas for a handheld memory prosthesis device where memory notes or reminders would appear on a handheld device based on the user’s location, and co-location with others. The technology used to realise location-awareness was the “Active Badge System” [17] which was deployed in a single office building. This operated by the user wearing a battery-powered device that emitted a unique radio signal every 15 seconds which was picked up by an array of sensors and from that the wearer’s location could be pinpointed. Other examples of location-aware services based on active badges were automatic re-routing of phone calls to different handsets, which allowed wearers to move around a building and have their phone calls routed to them, or diverted to voicemail if they happened to be co-located with several others inferring that a meeting was likely to be taking place. The “Shopping Guide” [2] is a PDA-based system that would give a shopper directions through a shopping mall on the basis of their shopping preferences and purchases to date, and also based on their location in the shopping mall. “PlaceIt” and the hospital specific message system [13] both rely on location awareness to drive their messaging service. PlaceIt [16] relies on simple location detection through GSM technology to trigger reminder messages and Munoz et al [13] developed a hospital messaging system which takes into consideration time, staff roles, artifact states and location to trigger its messages. “Just-For-Us” [7] is a location aware social application project proposed in 2005. Researchers recognised the potential of the the growing popularity of online social networking and developed a system which allowed existing friends to communicate through a PDAs when meeting in a familiar environment.

Many of the systems detailed here were built as prototypes with expensive PDA’s required for interfacing due to technology limitations at the time of conception. A major stumbling block for these systems was the mobile devices available and their limited computing power. Twenty-five years on from the “Forget me not” proposal, technology in the mobile computing domain has caught up with the demands placed on it by users and applications [5]. Mobile phones are now mini-computers which sense the environment through which they travel and they gather and generate data which can be easily harnessed and exploited yielding more intelligent

aware systems. With these technological advances came more sophisticated location-aware proposals. The popularity of the iPhone, with built in location and bluetooth sensors, fast reliable internet access for communication and the availability of an API for independent development has seen a surge in the number of location aware applications on the market.

The content delivered by the majority of the systems described until now is pre-specified and static. Users who encounter the physical spaces days, months or even years apart will have the same experiences. Newer mobile applications are leveraging the popularity of the Social Web revolution and in particular systems which encourage users to generate and share content online to enhance the static location aware proposals of the past to produce enhanced social mobile applications.

Google Latitude and Centr1 implement what is known as *Geosocial Networking applications* that allow users to pinpoint the location of friends on a map interface. Yelp, the restaurant review Web site, which relies on users to contribute reviews on eateries has enhanced its mobile interfacing by automating the process of selecting a location for which the user wishes to receive recommendations by communicating with the GPS sensors on the iPhone and consulting with its repository of restaurant addresses to determine the relevance of restaurants to a users location. Other systems have been similarly integrated with location sensitivity to provide ease of use in the mobile domain.

FindByClick.com is an interesting example of where users come together to contribute information about a geographical location. Users are encouraged to share tags, stories about encounters, restaurants reviews, retail information etc in any location that they wish. Contributing is online and achieved through a clickable map interface.

The increasingly popular microblogging platform, Twitter [6], [10], enables users to broadcast content quickly and easily to their social network. Twitter’s open API has fostered development of countless applications designed to support a multitude of platforms and devices and as a consequence Twitter is available to user through variety of devices and supports in-situ and on-demand messaging on-the-move. While the protocol by default supports only 140 characters, services such as TwitPic ¹ allow a user to connect photo-based media to a message. Further to this, the platform supports the geocoding or addition of a geographic reference to each message created and this enables content to be attributed to a particular location within space. The ability to correlate the content with particular location has given rise to a range of new applications for this social data, from simple visualisation of recent messages from around the globe in twittearth ², to modeling, predicting or recording movement of people around the globe as in ‘Just Landed’ a project which mined Twitter content for tweets containing the words ‘just landed’ to detect instances

¹<http://www.twitpic.com>

²<http://www.twittearth.com/>

of air travel³ or social weather mapping⁴ which leverages location tagged Tweets which contain mentions of weather conditions. The notion of geographically placed content is gaining traction and note. Recently Microsoft have announced Vine <http://www.vine.net/> a communication platform designed to take this notion of geo-tagged content one step further. It is designed specifically to allow you to receive and send alerts about places and people of interest to you.

While there is an emphasis in attributing digital and user-contributed content to particular places, there is additional increasing interest in exposing and creating awareness about people, their presence in and movement through real world spaces. Loopt⁵ is an online social network which provides a user interface to that network which can run on a mobile device. Through this mobile application, the user is positioned in the real world and their relative position to their contacts is exposed on a map-based interface and as such it affords the greater visibility and awareness of their social network. Other applications such as the Nokia Sensor [14] and Social Serendipity [4] have also sought to enable messaging between co-located physical people through mobile and digital technologies. More recently, the CityWare project conducted at the University of Bath has looked to leverage the movements of people through physical spaces sensed through Bluetooth as a way to augment online digital social networks with real world representations of those social encounters [9], [8]. In this work we detail a system which brings together many elements of the systems described above. We present a platform which delivers messages to users based on their presence in a physical space or co-location of people. Messages take the form of video, audio or text and can be contributed by a system administrator or by members of the public. The system is primarily mobile phone based but also has a Web interface.

III. CONTEXT AWARE MOBILE MESSAGING

The Context Aware Mobile Messaging System, or CAMM, is a context-triggered messaging platform. In this system, messages (SMS, MMS and/or digital signage messages) are delivered based to their target recipient(s) when the geographic location and/or co-location of the recipient(s) with other users satisfied the contextual preconditions specified by the message sender at the time of creation [5]. CAMM affords message senders an enhanced level of control over their messages by allowing them to specify requirements for a message's delivery including a time frame, location or co-located group. CAMM relies on up-to-the-minute mobile and location-based technologies to provide a user-friendly, as well as a highly customisable, adaptable and accessible framework for communication. To achieve this, CAMM exploits a range of mobile technologies including SMS, MMS and Bluetooth presence as well as Ubisense and GPS for indoor and outdoor positioning respectively.

³<http://blog.blprnt.com/blog/blprnt/just-landed-processing-twitter-metacarta-hidden-data>

⁴<http://smalltalkapp.com/#all>

⁵<http://www.loopt.com/>

Comparisons can be drawn between CAMM and Twitter [6], [10]. Twitter encourages users to tell the world “what you are doing” and broadcasts this information to subscriber lists and the general public for consumption. While both offer socio-digital messaging platforms the goals of the two and the means in which they maintain a relationship between the users and their content are wholly different. Even though ‘tweets’ can be geocoded, there is no formal relationship between Twitter’s messages and the real world spaces in which they are consumed. As such Twitter offers messaging within a virtual space while CAMM’s messaging is inherently based on movements and encounters within real-world spaces. Content is connected to real people and places and can only be accessed in-situ when the user’s context matches that of the target message. While CAMM supports microblogging it broadcasts posts in the context of a physical location or the co-location of people in order to deliver the information to more targeted audiences.

These social features are the key factors which distinguishes CAMM from other location-aware systems (a detailed outline of such systems is provided in Section II). As the messages can be ‘attached’ to locations or to people within that environment, CAMM allows users to contribute digital media which relate to encounters with their physical environment. Photo-, video, audio- and text-based content can be shared easily with others through an application running on a mobile device which in turn communicates with a central server. The range of media supported by the platform as well as the ability to garner and exploit contextual information of both the users and the message content results in a context-aware platform for communication which can be applied in multiple domains and settings. Within this paper we present one such setting — that of the social mobile tourist — though the next section outlines the infrastructure of CAMM. Here we outline the architecture of the platform and how it leverages existing mobile device technology to achieve a context-aware messaging system and its extensibility, before moving on to detail a the tourist implementation of the CAMM architecture.

A. System architecture

At the core of the CAMM platform is a Java-based messaging server. The server accepts new messages (along with contextual triggers) from an application deployed on mobile devices or input through a web interface and it marshals the delivery of messages as the context of the systems’ users updates. This server is tightly coupled with a Python Series60 Symbian OS mobile phone application from which it receives regular feedback from GPS location sensors (Unisense in the indoor version) and Bluetooth co-location sensors to detect the location and co-location of its users. As it receives this feedback, it can then determine if there is a message appropriate to the current contextual setting of the user. CAMM’s server maintains an Oracle database, exploits an SMS gateway to send messages and communicates with the python application via XML data sent over GPRS.




Fig. 1. System Architecture

The suitability of location sensitive technology for context-aware systems is determined by the accuracy of the location information required. CAMM was designed to cover both indoor and outdoor locations in order to broaden its applicability. Two versions of the system exist but here we will focus our attention on the outdoor implementation which utilises GPS for localisation of the users. The indoor application was integrated with Ubisense location technology, however other alternatives exist including WiFi based sensing, more detail on the indoor implementation is available in [5].

The CAMM architecture requires the user to carry a mobile device which runs the CAMM Python application. This application facilitates the creation and delivery of messages as the user moves through real-world spaces and is responsible for gathering the location and co-location data for that user. This serves two core purposes. First, the gathered context can be used in the creation of new messages. When a user contributes to the message database i.e sends a message, the context of the message is required in order to configure its delivery triggers. Secondly, the gathered context is passed periodically to the server, which determines if the current environs of the user have changed and if the delivery of a message should be triggered.

Identifying the context, i.e. the location and co-location, of users at any time involves the java application on the central server and the python application on the individual handsets. Co-location detection uses Bluetooth presence technology. The CAMM application installed on the users phone “sniffs” out other visible Bluetooth devices in its proximity at 1 minute intervals and compiles a list of co-located devices CL_i for that user. CAMM exploits the familiarity of other devices in order to decipher which users are important to individual users [12].

With the popularity of GPS enabled devices in particular the high-end cell phones, such as the Nokia N95 and the iPhone, GPS was the natural choice for the location sensing required

```

<CSM>
<GUPDATE>
<OWNER>00:1C:D4:47:1B:2D</OWNER>
<CO-LOCATION>
<ADD><MAC>00:17:E3:E6:42:03</MAC></ADD>
<REMOVE></REMOVE>
</CO-LOCATION>
<GPS>
<LAT>53.3090</LAT>
<LON>-6.2230</LON>
</GPS>
</GUPDATE>
<TIME>15:04:2008 18:52:50 GMT</TIME>
</CSM>

```

Fig. 2. Update Process

for the CAMM outdoor implementation. The GPS sensors in the phone determine the longitude and latitude of the phone’s location and are polled at the same 1 minute intervals to report the location information, L_i to the CAMM python application. The location and co-location information is sent via GPRS to the CAMM Java server for processing as a simple XML packet. A sample context packet can be seen in Figure 2. 2

B. Message Creation & Delivery

CAMM is designed to exploit any combination of location, co-location of users, and time as its parameters for the creation and delivery of messages to and from users. It is also designed to support the creation of several forms of user contributed content including text, video and audio all of which can be contributed via a cell phone or other mobile device or through a Web interface. Messages are capped to the length of a standard SMS. Video and audio can be captured on the users phone and sent via MMS to the central server or uploaded onto the website however the size of the file is restricted to facilitate timely delivery over GPRS or via SMS/MMS to mobile users.

The CAMM engine receives messages from users via the web or mobile phone and stores them according to contextual

triggers in the database. In addition to the message content, the server must receive location and/or co-location information, and an expiry time for a message in order to store it in the database. A recipient list is provided optionally and if no recipient list is received, the message is assumed to be a public message. This represents the minimum set of information required for each message. Further information could be captured about the message to increase its utility and how targeted its delivery might be. For example, a message type could be used to categorise a message and determine its relevance to particular users. However, requiring this information at time of creation increases the workload for and burden on users, and particularly when using a mobile device for entry this is undesirable given that text entry and form filling can often be a cumbersome process. Rather than explicitly capturing this information, we aim to automate this process by using personalization techniques. As with other mobile personalization technology [15] monitoring user habits and preferences would sid us greatly in this task. Using simple profiling techniques the system could determine the most likely candidate or candidate combinations for recipient selection and location determination.

The CAMM messaging solution offers a versatile platform for communications and opens up a range of new possibilities for messaging scenarios and situations in which traditional text messaging might not be appropriate or suitable. We now elaborate on some of these messaging scenarios. In the case of advertising, if posting an advertisement of a sale on an public information implementation of CAMM in order to broadcast an advertisement to passers by the message would specify which physical location the advert is relevant in (outside the business, at the bottom of the street close to the business etc) who is to receive the advert (either a list of users or a public message) and the time in which it is relevant (i.e. when the advert should first be delivered and when it should cease to be delivered). In the case of a member of the public who is meeting two friends who do not know each and is running late, they can specify a text to be delivered when the two users are both located at the designated meeting point if they arrive before their common friend. The location, co-location and time parameters are all required in this instance.

CAMM's message delivery process is triggered by updates from the python phone application. Co-location and location information is passed to the messaging server periodically and is used to determine if the context of any messages have been satisfied. So for a given a user, i , whose last know location is L_i and who is currently co-located with a set of users CL_i , CAMM checks to see if any messages MS have been set which are constrained by the current time, CL_i and/or L_i . If so these messages are sent via an SMS gateway and marked as read in the database.

While there are currently some limitations to widespread use of our CAMM system on a mobile device, these obstacles are far from insurmountable. Internet access costs on mobile devices are currently expensive and as the system regularly polls for collocated devices and receives message requests via

the internet, the resulting cost of CAMM use on a mobile device can be high. However, there are a factors which help to overcome this. Data bandwidth can be minimised by storing messages for collocated delivery on the individual handset thereby requiring less frequent queries to the server and by performing familiarity calculations on the handset and in effect compressing the raw encounter data. Also there is a trend of mobile data access edging toward cheap high speed mobile internet access and which will additionally make our system more cost-effective in real-world scenarios.

Another limitation of the current implementation is that message creation on a mobile phone assumes that the location for a messages delivery is the location of the user at the time of creation. Messaging to other locations is currently only achievable when creating a message through the web site. While this feature has its benefits, there are many scenarios where having the ability to deliver message to locations other than the one you are in would be very useful e.g. when in the office you could send yourself a message to deliver when you got home. In future versions of the software, the integration of dropdown lists of 'favourited' locations or mapping tools, such as Google Maps, to enable the selection of GPS coordinates by moving across a map, allow such functionality to be added to CAMM.

IV. AUTOMATED MURMURS

The motivation for our work in the area of context awareness was sparked by the developments in the area of mobile devices and the huge popularity of the Social Web. By combining these two factors we have realized a platform which combines context-aware computing with social contribution and collaboration to form an environment where by users can contribute to the experiences of others as they pass through a physical space. We will illustrate the use of this platform in a very real world example often used in location aware computing, a tourist application, and highlight the collaborative process in which a rich collection of varying data relevant to tourists can be gathered through a social computing application.

[Murmur] is a documentary oral history project that records stories and memories set in specific Dublin locations, as told by Dubliners themselves. They collect and make accessible people's personal histories and anecdotes about the places in their neighbourhoods that are important to them. The Bealtaine Festival 2007 partnered with the Dublin Docklands Development Authority to co-commission [murmur] Dublin Docklands, in which Dubliners' stories of the docklands would be heard in situ via a mobile phone connection⁶. It was one of several implementations of the worldwide project. Other cities that [murmur] was run in include Toronto, Montreal and Edinburgh.

The current implementation of [murmur] is that at each physical location mentioned in a story a [murmur] Dublin Docklands sign (large green ear attached to a street sign, see 4), with telephone number and location code, marks where

⁶<http://murmurdublindocklands.info/>

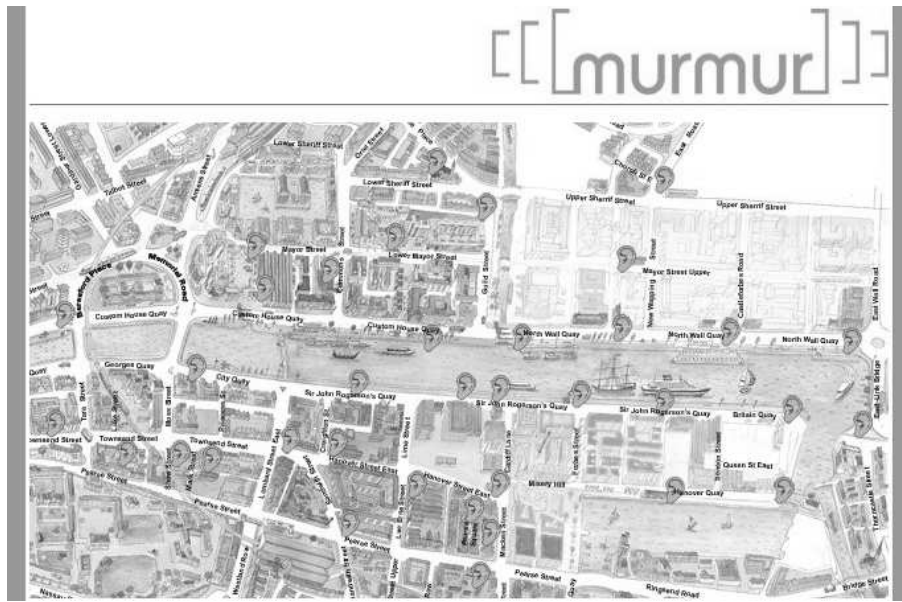


Fig. 3. Murmur Dublin



Fig. 4. Murmur Signage

stories are available. With a mobile phone, users can listen to the story while there experiencing the place itself.

The goal of our project was to automate this service and exploit the trend of user-contributed data, in this case through a mobile device to enhance the data already in existence with other information which is relevant to users. In the [murmur] system users have to be aware of where they may encounter [murmur] signs (see 3) and locate them in what can be a cluttered environment. Then when they find a sign (see 4) they must ring the number on it to hear a story about their location. This method of message retrieval puts the workload on the user to locate the information and retrieve the content and in looking for the [murmur] signs a person may not fully see the city around them. The experience with the non-digital [murmur] installation is also a single user experience, for every

member of a group to take part each must make the phone call and listen to the audio independently. Other restrictions include the noise levels of the city making hearing the audio difficult, missing signs and thus content and when the audio snippet refers to buildings the listener has very few clues as to which building is being discussed. The stories were only offered in audio format and the addition of supplementary material in the form of text, video or image content could have enhanced the richness of the experience by offered new modalities through which the story could have been supported. A secondary goal of the project was to enhance the experience of [murmur] by providing a solution which dealt with many of these limitations.

We have build and initial prototype to demonstrate the CAMM architecture in a system named Automated Murmurs. The application was preloaded with geocoded the [murmur] data relating to Dublin city. Automated Murmurs brings new social functionality to the [murmur] project. The primary social aspect of the system is the opportunity for users to contribute their own content to the message repository and for this content to be shared with individuals, groups or to the public. Users can upload photographs or videos of their experiences at locations. They can share a little about themselves and their thoughts about the tour and exhibits. They can also add content which isn't directly related to the tour itself but interesting to others for example restaurant or hotel reviews, shops or services of interest, suggested routes etc. Each message serves to build up a social layer on top of the existing tour data which harnesses the power of the tourists to create a more dynamic and informed system. The second social feature exploits the co-location information and informs users who are co-located of the presence of others. This tool can be tweaked to inform users who appear to be taking the

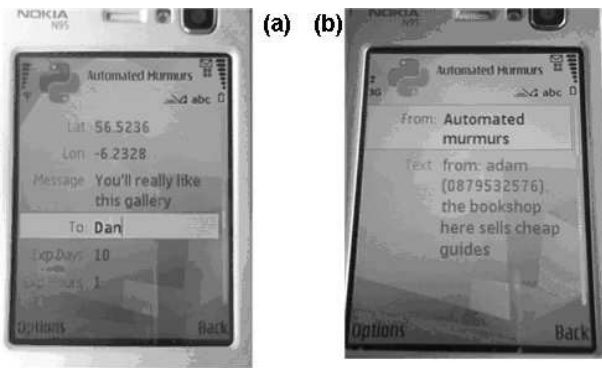


Fig. 5. Mobile phone interface

same route and could target users who are traveling alone and encourage them to communicate or even continue the tour together.

A. User Interface Options

As mentioned previously, CAMM was designed with both phone and Web based interfaces. Primarily the functionality of both interfaces are the same, users may send and view content relating to a physical location. The phone interface is modeled on a traditional SMS or MMS interface with users specifying recipients for private messages, including content of some form. However, in the case of CAMM messages the additional fields for specifying an expiry time, a location and co-location parameters are required. Recipients and the people who are required to be co-located (if any) are elected using phone numbers and by default the user's last known location is used as a the delivery location. Users can opt to select locations from the list of CAMM specified locations but in its current implementation selecting additional locations is not facilitated. With an upcoming implementation of CAMM for the iPhone and a simple integration with Google Maps, the ability for users to leave messages at any geographical location will be supported. Messages delivered from CAMM are received via SMS or MMS, the messages can contain a combination of media and text (see Figure 5(a)&(b)). Users can store this media on their phone and revisit it again at a later date or pass around a handset for more than one user to benefit from a single message's delivery.

The Web site provides users with both pre- and post-visit information about the region that they visit. All of the messages both from the system and contributed by users are available for consumption. The content is overlaid on maps for ease of route planning and areas of high activity in terms of user contribution are easily visible highlighting popular or controversial parts of the physical space. Messages from the community as a whole are identifiable from a particular users contribution through colour coding. Users can listen to the audio, read text, view photographs (see Figure 6) and leave messages for others before their trip or even if they are not planning a trip. Once a user has visited an area their route is visible on their personalized map as shown in Figure 7. From

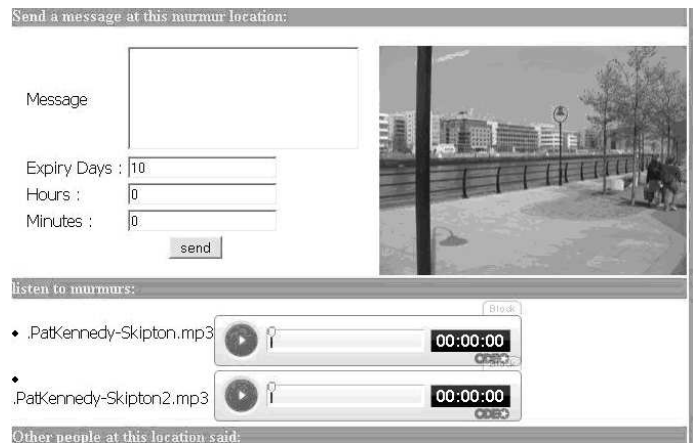


Fig. 6. Creating a message at an existing location with the Web interface

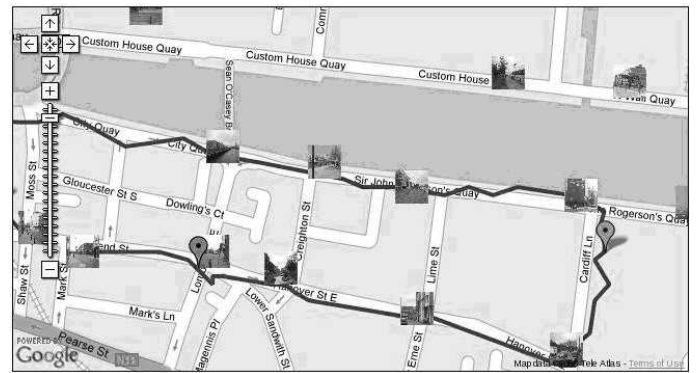


Fig. 7. Personalized map showing the path and contributions of an individual user.

this map users can see what content they did not consume or where to return to the following day.

Depending on the environment in which CAMM is operating, the delivery and display of messages can also take the form of digital signage. For instance in the city environment, several large in-situ screens, bus shelters or billboards etc. can be used by CAMM to display public messages relating to the area.

V. FUTURE WORK

The work described here represents an important step on the road to the development of a new generation of interactive mobile information systems, capable of responding to the context and preferences of individuals or groups of users in the real-world while harnessing the power of the social Web to produce next generation, lively, evolving social applications.

User generated content while rich and varying can result in the familiar problem of information overload. Users of systems such as those described in this work will tire easily of relentlessly receiving messages as they move through a physical environment. Placing expiry dates on messages will reduce the information overload to a degree but necessity then calls for the presence of personalized content delivery

mechanisms in order to successfully determine the relevance of messages to users. Personalization in this domain can be difficult, user preferences can be difficult to anticipate especially when interaction with the system is over a short period of time, multimedia messages as well as short text based messages can be difficult to classify. The success of CAMM based and other UGC systems is inherent on delivering the right content to the right subset of users and this issue is of prime importance to the authors.

ACKNOWLEDGMENTS

This work was supported by Science Foundation Ireland through Grant No. 07/CE/I1147.

REFERENCES

- [1] G.D. Abowd, A.K. Dey, P.J. Brown, N. Davies, M. Smith, and P. Steggle. Towards a better understanding of context and context-awareness. *Lecture notes in computer science*, pages 304–307, 1999.
- [2] T. Bohnenberger, A. Jameson, A. Krüger, and A. Butz. User acceptance of a decision-theoretic location-aware shopping guide. In *Proceedings of the 7th international conference on Intelligent user interfaces*, pages 178–179. ACM New York, NY, USA, 2002.
- [3] Scott Carter, Elizabeth Churchill, Laurent Denoue, Jonathan Helfman, and Les Nelson. Digital graffiti: public annotation of multimedia content. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1207–1210, New York, NY, USA, 2004. ACM.
- [4] Nathan Eagle and Alex Pentland. Social serendipity: Mobilizing social software. *IEEE Pervasive Computing*, 4(2):p28 – 34, 2005.
- [5] J. Freyne, E. Varga, D. Byrne, A.F. Smeaton, B. Smyth, and G.J.F. Jones. Realising context-sensitive mobile messaging. *Lecture Notes in Computer Science*, 4805:407, 2007.
- [6] Akshay Java, Xiaodan Song, Tim Finin, and Belle Tseng. Why we twitter: understanding microblogging usage and communities. In *WebKDD/SNA-KDD '07: Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 56–65, New York, NY, USA, 2007. ACM.
- [7] J. Kjeldskov and J. Paay. Just-for-us: a context-aware mobile information system facilitating sociality. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, pages 23–30. ACM New York, NY, USA, 2005.
- [8] Vassilis Kostakos. Social networking 2.0. In *CHI '08: CHI '08 extended abstracts on Human factors in computing systems*, pages 3381–3386, New York, NY, USA, 2008. ACM.
- [9] Vassilis Kostakos and Eamonn O'Neill. Capturing and visualising bluetooth encounters. In *adjunct proceedings of the conference on Human factors in computing systems (CHI 2008)*, Florence, Italy, April 2008.
- [10] Balachander Krishnamurthy, Phillipa Gill, and Martin Arlitt. A few chirps about twitter. In *WOSP '08: Proceedings of the first workshop on Online social networks*, pages 19–24, New York, NY, USA, 2008. ACM.
- [11] M. Lamming and M. Flynn. Forget-me-not: Intimate computing in support of human memory. In *Proceedings of FRIEND21*, volume 94, pages 2–4, 1994.
- [12] B. Lavelle, D. Byrne, C. Gurrin, A.F. Smeaton, and G.J.F. Jones. Bluetooth familiarity: Methods of calculation, applications and limitations. In *MIRW 2007 - Mobile Interaction with the Real World, Workshop at the MobileHCI07: 9th International Conference on Human Computer Interaction with Mobile Devices and Services*, 2007.
- [13] MA Munoz, M. Rodriguez, J. Favela, AI Martinez-Garcia, and VM Gonzalez. Context-aware mobile communication in hospitals. *Computer*, 36(9):38–46, 2003.
- [14] Per Persson and Younghee Jung. Nokia sensor: from research to product. In *DUX '05: Proceedings of the 2005 conference on Designing for User eXperience*, page 53, New York, NY, USA, 2005. AIGA: American Institute of Graphic Arts.
- [15] B. Smyth and P. Cotter. Intelligent navigation for mobile internet portals. In *IJCAI Workshop on AI Moves to IA: Workshop on Artificial Intelligence, Information Access, and Mobile Computing. The 18th International Joint Conference on Artificial Intelligence*. Citeseer.
- [16] T. Sohn, K.A. Li, G. Lee, I. Smith, J. Scott, and W.G. Griswold. Place-its: A study of location-based reminders on mobile phones. *Lecture notes in computer science*, 3660:232, 2005.
- [17] R. Want, A. Hopper, and V.F.J. GIBBONS. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, 1992.