Journal of Mobile Multimedia, Vol. 1, No.1 (2005) 047-061

ROAD TALK: A ROADSIDE LOCATION-DEPENDENT AUDIO MESSAGE SYSTEM FOR CAR DRIVERS

MATTIAS ÖSTERGREN
Interactive Institute, Stockholm
mattias.ostergren@tii.se

OSKAR JUHLIN
Interactive Institute, Stockholm
oskarj@tii.se

Received February 25, 2005

Abstract. We present the Road Talk application which is a mobile message system for car drivers. It is designed to allow messaging while driving, for instance, sending notes to surrounding drivers to tell them about hazards. The user annotates the road with audio recordings, typically voice memos. This annotation is location-based, which means that messages are automatically coupled to the location where the audio was recorded. When another driver coincides with this location the audio file is automatically played. In the background, messages are also shared with fellow car drivers. This distribution draws on wireless mobile ad hoc networking (MANET) and the physical movement of vehicles. We designed messages to spread 'epidemically', since MANETs between vehicles are often disconnected. Messages hop from one car to the other when there is coverage and then passively travel some physical distance. Then they hop again when coverage reappears. The design of Road Talk is motivated and informed by a statistical analysis of commercially available systems for distributing warnings of speed trap events. These events occur where police monitors conformance to speed limits and although our intention is not to improve these services the analysis gave insights in, for example, temporal and spatial occurrences of roadside location-dependent messages.

Key words: roadside location-dependent, mobile message system, epidemic distribution

1 Introduction

We present the Road Talk application, which is a message system for car drivers enabling them to annotate the roadside with voice messages. These messages are shared with fellow drivers that they meet. The message is played when these drivers coincide with the location for an annotation. The Road Talk prototype is an application for PDAs that uses the built-in microphone, speaker, Bluetooth GPS receiver and 802.11 wireless networking interface of the device to record, play, and share memos.

Road Talk is motivated by drivers' desire to know as much as possible about the state of the upcoming road before getting there. During any trip the driver moves forward into situations that he or she are unaware of prior to reaching them. Take for instance a bend in the road where the section immediately after is obscured. Most of the times the driver would safely travel beyond this bend assuming that nothing unexpected will appear. This is mostly a correct behaviour, as the state of the stretch will be similar to what was before it. But this assumption is wrong if the oncoming section has undergone a drastic change. The driver could be heading for a traffic queue, or a road construction work, instead of free way. A hazards could also be less dramatic e.g. when children are playing dangerously close to road. All of these situations can have a negative impact on traffic safety and driver can more easily deal with the situation if he or she is forewarned. Road Talk is designed to distribute such knowledge. With Road Talk an obstruction would only surprise the first driver to the scene. This driver could record a warning, which would then be automatically spread to other drivers that he or she meets.

Here we account for the design and initial evaluations of the Road Talk prototype. Our work draws on a study of commercially available message systems as well as our extensive experiences with other prototypes for drivers. The purpose of the commercial available systems is to let subscribers warn each other for *speed traps* with short text messages. Speed traps are applied when police monitors conformance to speed limits manually. The police stand by the road side measuring the speed of the passing cars with for example a laser instrument. Our statistical study of such messages shows that people are interested in sharing information about this particular type of roadside-location events using mobile technologies but also concerning other topics than speed traps. We also found that messages are most frequent during the day in urban areas. This influenced the message distribution design of Road Talk. We use the physical movement of cars, where the units are mounted, to share memos over wireless ad hoc networks of limited range when the cars meet. Copies of memos ‘hop’ from one car to the next when they come within wireless reach. This form of sharing is best categorized as an *epidemic information diffusion* process. It gains from high density of cars, which is most prevalent during the day in urban areas.

The work presented here contributes to research on *mobile message systems* and especially *location-dependent message systems* of the context-aware computing domain. Mobile message systems, such as SMS, are used by a vast number of people. They allow sending text-messages practically anywhere and anytime. On the other hand, many messages that are prevalent in our everyday life, such as post-its stickers and road signs, are not available as such. Yet they are powerful because they are put somewhere or attached to something to highlight an aspect of that context. The Road Talk prototype and the analysis of speed trap services we present here contributes to the growing body of research that investigate the benefit of context and particularly location to messaging.

The paper is structured as followed: In section three we present the Road Talk concept, a user scenario and a brief overview of the prototype. In section four we account for the motivation to the design of Road Talk. In this section an overview of the study of speed trap message services and a discussion on the viability of epidemic message distribution can be found. In section five we discuss implementation details of the Road Talk prototype and finally in section six we present an initial performance study.

2 Related Work

The work presented here is related to research on *message systems* i.e. research on technologies that enable people to send and receive messages *asynchronously*. A message is a piece of information of some particular form, for example, text, graphics or sound or combination of all [3]. We refer to the action of sending, receiving and otherwise engaging with such systems with the term ‘messaging.’

Messaging occurs in many shapes and contexts from stationary to mobile based on a wide set of technologies. The most widely spread messaging applications today are based on Internet technology. Undeniably e-mail, instant messaging [15, 20], blogs (e.g. [29]), chats [4, 21] have contributed to the popularity of it. Particularly, the work here is related to mobile message systems, which today also is a widespread phenomenon. In Europe sending messages with the Short Message Service (SMS) available though GSM mobile phones is a huge phenomenon [12]. Similarly, in Japan, using mobile phones to do e-mail messaging is a widespread practice [19].

Mobile message systems allow messaging anytime and anywhere away from the desktop and the always Internet connected PC there. However, many mundane messages concern a particular spatial context and as such the place gives additional meaning to them. For example a sign outside a grocery shop stating “buy strawberries ‘here’!” obviously indicates that ‘here’ should be interpreted as that grocery shop and nothing else. There are several research projects in the context-aware computing domain that aims at bringing such context information into mobile messaging [2, 10, 13]. Typically, these systems explore an extension to the notion of message and couple a reference to a physical location to a text body. This, in turn, lets the user interact with the messages of the systems in mobile fashion literally. When users roam the physical world and coincide with the location for a message, it becomes available. Typically the message ‘pops up’ in the GUI or becomes ‘available’ in some other way. This is still an emerging field and the benefits are still questionable. For instance, in the evaluations of Egraffiti [2] and GeoNotes [24] it was found that the users often mistook the systems for chats entering messages not particularly or concerning actual places. This could be due to many reasons; one might be that although these systems are location-centric they were still practically too similar to many chats and blogs. This is most obvious in GeoNotes where positioning was based on wireless base station access. Such positioning could not discriminate between locations within the base station range and the users had to specify the particular location with so called “place labels” i.e. a field for entering a textual description of it. Messages in chats often have structural information fields too but for other purposes such as ‘subject,’ ‘topic’ or ‘group’ etc. Consequently the place label field was used to indicate a topic not necessarily concerning a place which indicates GeoNotes was basically approached as a chat.

Furthermore, mobile location-dependent messengers of the context-aware domain are primarily designed for the pedestrian. These systems are intended for people walking around, whereas our intention is to design for car driving. Walking and driving is related, in both cases you have a plan ahead, make sure you remain on the decided course, avoid bumping in to others, etc. However, they differ greatly in the means and the speed in which movement occurs. In both cases a person is active, but whereas a walking person uses his or her legs, a driver uses his or her feet and hands to interact with the pedals, steering wheel, gear stick etc of the vehicle. As such, driver has to pay more attention to how and where he or she is travelling than a walking person. But practically there is also room for doing other things while driving, such as listening to music [1] and talking in mobile phone [7, 18].

Finally, common for all of the above mentioned systems are that the users always have access to the messages and they rely on the assumption that opportunities for networking and distributing messages are normally present. No connection is basically considered an error condition. However, there are many other research projects, similar to Road Talk, which investigates situations where communication opportunities are expected to be intermittent for various reasons. Pollen Networking [11] is an example of a mobile messaging where distribution is achieved by people’s movement, typically in an office, and messages are transferred at physical contact or at very short, centimetre-range between devices. In Delay Tolerant Networking [5] the aim is to achieve an interplanetary network, using satellites to send messages across space. Here communication requires alignment of transmitters and receivers. This is not always achievable since satellites orbit around planets.

3 The Road Talk concept

The Road Talk application allows sharing roadside location dependent messages between drivers that meet in traffic. A message in Road Talk is constituted by audio recordings, or voice memos, together with a location reference. When a memo is recorded, it is simultaneously associated with a GPS position. The Road Talk application also tracks the current GPS position and whenever coinciding with the location for a memo, that particular memo is played. Furthermore, in the background Road Talk shares the memos and as such, a user will be able to hear what other users have recorded.

3.1 Usage scenario

We envision Road Talk to be used like in the following scenario (Figure 1). Sue-Ellen heads out on the road and sees the neighbour's boy playing in a water puddle dangerously close to the street. She decides to record a memo to warn others and she pushes the record button while speaking out loud: 'oops kids.' She drives on and in an intersection she meets many other cars and drivers. Some of them are using Road Talk and in the background memos are exchanged. Bobby, an occasional driver in the area, receives Sue-Ellen's warning as he is heading in the opposite direction. Furthermore, he is speeding, as he is late to a meeting. Fortunately he hears Sue-Ellen's voice: 'oops…kids' through his PDA running RoadTalk and slows down in advance of the boy.

Figure 1. Usage scenario. Sue-Ellen sees boy and formulates message (a). She meets Bobby and the memo is shared (b). Bobby reaches initial position and hears Sue-Ellen's message (c).

3.2 The RoadTalk prototype

The Road Talk prototype (Figure 2) is an application for PDAs with GPS positioning and wireless ad hoc networking implementing the concept outlined above. It uses standard GPS positioning to track current location and associate memos to physical locations. Further, it uses wireless ad hoc network to synchronize memos with other Road Talk units. As such it is an application for PDAs running the PocketPC 2003 operating system and it is implemented in C++. Positioning is achieved with Bluetooth GPS receivers and sharing of memos is accomplished with popular standards such as XML, and HTTP over IEEE 802.11b wireless networks in ad hoc mode. These wireless networks are automatically established between the PDAs mounted in cars that move at the speed of traffic.

Figure 2. An HP Ipaq 4150 PDA and a Bluetooth GPS Receiver (a). Screen shots of Road Talk. (b) Map show zoom level 1 (b) and level 2 (c) centred on the same coordinate. The squares indicate locations of memos.

4 Motivation

Three things motivate the particular ideas we have pursued in designing the Road Talk prototype. First, there are *user needs for communicating roadside-location dependent events*, such as hazards. We will ground this argument in a statistical analysis of commercially available speed trap message systems for drivers. Second, this analysis yields that messages roughly occur when and where traffic is dense and therefore the Road Talk concept is a good candidate for using wireless Mobile Ad Hoc Networking and *epidemic diffusion* to distribute the memos. The applicability of the technologies we have chosen to accomplish this is motivated by previous research on distributing web pages in traffic encounters between motorcyclists (the Hocman prototype). Finally, speed trap systems are not well adapted for in-vehicle use and the design of Road Talk message system must particularly *accommodate driving* in order to be properly useable. The design approach we have taken is motivated by the principles that governed the design of infrastructure management support for driving road inspectors (the Placememo prototype).

4.1 User needs for communicating roadside-location events

Car drivers already communicate with various means to achieve smooth coordination [16]. The most obvious is the use of turn signals. Occasionally they also communicate with means for some other original intention, e.g. by flickering the headlights to warn oncoming drivers for a hazard. Still it is not evident that drivers would use mobile technologies to help each other in more elaborate ways. Nor is it obvious that they would require any new technology to do it. The statistical analysis of existing speed trap message systems gives an initial understanding of the needs of drivers to communicate roadside-location events using mobile technologies. It also gives a rough impression of when and where these messages occur, which in return tells us where roadside-location events worthy of communication happen.

A speed trap is an activity taking place by the road side, where the police monitors conformance to speed limits with manual means. They stand by the road and measure the speed of passing vehicles with radar or laser equipment. The purpose of the speed trap message systems, which are available in many countries, is to allow drivers to warn each other for their occurrence. Drivers can report through various means e.g. SMS, wap, the web or phoning to a call centre. A subscriber of such service receives these warnings to his GSM mobile phone as SMS text messages.

Our concern is not to improve speed trap systems per se. In a sense these particular services are counter-productive to police work and therefore immoral. Instead we are interested in these systems as a particular, and peculiar, example of services aimed for drivers to help other drivers. They are of interest since they are message systems, allowing people to communicate knowledge in roadside-location events. Besides, warning for speed traps is not always counter productive to traffic *safety*. Occasionally, knowledge of a speed trap in advance, complies with the intentions of the police. They reveal the time and location of the speed trap e.g. at their website, and on local radio. Finally, it must be stressed that it is not illegal in any of the countries to report or subscribe to speed trap messages.

In the following, we present a brief overview of four speed trap services and the messages they contain. The statistics provide us with an understanding of the *quantity*, *location*, *temporal distribution, structure* and *content* of the messages. We collected messages from the public web pages of the service operators over a period of roughly five months, from the 10th of October 2003 to the 3rd of Mars 2004. In total we downloaded the *quantity* of 6140 messages. This data does not accurately tell us how many that actually received the messages. Yet we believe that the speed trap message systems were actually used for receiving messages in this period. For instance, the number of subscribers of one particular service was about 1500 according to one operator's website.

A user subscribe to messages about speed traps that occur in a specific region. We have found that the messages were *located* in no less than 644 such regions of Denmark, Norway and Sweden. The average number of messages per region is nine and the standard deviation is 20. The large deviation is due to the messages are unevenly distributed and the top twenty regions have about 31% of all the messages. In the very top we find urban regions corresponding of larger cities in Scandinavia: Bergen / Hordaland (198), Göteborg (186), København (170), Stockholm (154) and Trondheim / Sør-Trøndelag (141). It may seem surprising not find Oslo in the top five. After all it is the largest city in Norway. The reason is that it was divided into smaller regions. Accordingly, there are six smaller regions that together cover Oslo and collectively they had 173 messages.

The messages are *temporally distributed* during 24 hours as shown in Figure 3. The analysis yield a pattern where the most frequent activity occurs between nine a.m. to around six p.m. There is a distinguished peeks at noon. The message frequency seems to decrease from about six p.m. and vanish altogether at four in the morning. Per day, the average number of messages is 42 and the standard deviation is 23. The message rate varies with the weekdays and holidays, when the number drops.

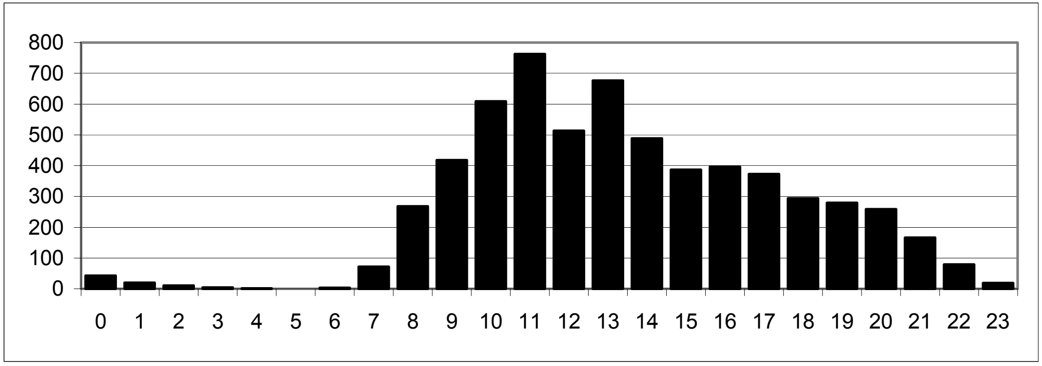

Figure 3. Temporal distribution of messages during 24 hours.

We have also examined the wording of the first 416 messages closer (Table 1). The messages had a common three-part *structure*. The first part is a brief description of the activity e.g. stating the occurrence of a police raid. At times, the activity description includes a description of the type of police control e.g. whether the police controls attendance to speed limits, safety belt prescriptions, or allowed vehicle weight. The type is sometimes implicit by the reference to the particular instrument the police are using, such as 'laser', or 'photo-vehicle', which are used for controlling speed limits. Second, the message provides a reference to the location where this activity occurs, such as 'Roskilde street in Vipperød beside gas station 100-150 meter within city sign.' Finally, the message occasionally includes the impact of the activity, e.g. informing that a particular road is closed-off. Further, we found that as much as 12% of the messages are reports not concerning speed traps at all, but accidents, weather conditions etc. This is a surprisingly high rate seeing that members pay to receive speed trap warnings and are probably reluctant to messages concerning other topics. But still people post such messages and they get through the system. We believe this indicates that users also need to share information about other topics than speed trap warnings.

Orignal message	English translation
HOLBÆK FOTOVOGN PÅ ROSKILDEVEJ I VIPPERØD LIGE VED TANKSTATION 100-150 METER INDENFOR BYSKILT	Holbaek photo-vehicle at Roskilde street in Vipperød beside gas station 100-150 meter within city sign
Laser på E4 i höjd med Bölesjön	Laser at E4 around Lake Böle
KONTROLL PÅ E39 CA 2 KM ETTER SØYLANDSKIOSKEN RETNING STAVANGER.	Inspection at road E39, about 2 kilometers after the kiosk at Søyland heading Stavanger

Table 1. Examples of speed trap warning messages

The messages are most frequent during the daylight hours on weekdays and occur mostly in urban regions. Essentially the messages coincide with the time and place when traffic is most intense. However, a closer look at the distribution over the day reveals that messages are not very frequent during morning and afternoon rush hour traffic, when people commute to and back from work. Thus it seams that the users of the speed trap message systems are not commuters, but people who do work by driving e.g. delivery work, and taxi.

4.2 Epidemic message diffusion

The speed trap analysis indicated that messages are most frequent where and when traffic is intense. In Road Talk we investigate an approach that uses this fact for message distribution. We use the physical movement of cars and wireless networking of limited range performing transfers when cars meet to distribute data. As such messages move in traffic with the cars, and 'hop' from car to car when they meet. The Hocman prototype for motorcyclists showed that it is technically feasible to transfer HTML files with so called wireless Mobile Ad Hoc Networks (MANETS) between nodes moving at the speed of traffic [9]. Using such MANET technology is appealing for distributing messages in Road Talk, as it requires very little infrastructure to operate. In fact, only the PDAs and the networking interfaces they contain would be needed to form a fast and cheap distribution method.

Message distribution in Road Talk is classified as an *epidemic information diffusion* process [17]. Such process shares some similarities with how diseases like influenza spread among people. Virus spread when somebody carrying it encounters somebody who is not yet infected. The spreading becomes epidemic when a large number of people are infected. Epidemic processes are well studied and describe many phenomena besides diseases, for example the adoption of high-technology products [21] and information diffusion in MANETs [17]. What is particularly interesting with epidemic diffusion in MANETs is that there exists an optimal node density yielding an optimal diffusion speed. Khelil et al. found through a simulation of a distribution scheme that resembles Road Talk an optimal node density of 620 nodes per square kilometres [17]. Their findings are not directly applicable to Road Talk as they are based on 'artificial' simulations that only vaguely resemble traffic flow and wireless communication on streets. For instance, the nodes were confined to a square area and moving randomly and the communication range was fixed to 75 meters. Nevertheless, this metric gives a crude estimate of the density of Road Talk applications to achieve a satisfactory spreading of memos and considering actual statistics on the number of cars in an urban region, the number is surprisingly low. For example, in 2003 there was about 278 500 cars registered in the municipality of Stockholm [26] and the municipality is responsible for roughly 16 square kilometres of streets [25]. As such there would be about 17 400 cars per square kilometre, if the cars were evenly spread out on the streets. Hence only about 3.6 % or (or about 10000) of the total number of cars would need to have RoadTalk, to achieve an optimal node density of 620.

4.3 Design that accommodates driving

The speed trap message systems have been somewhat successful with people that do work by driving. Still, there is room for considerable improvement, especially concerning user interaction. The Placememo prototype for infrastructure managers working along the roads [6] show that it is possible to accomplish a user interface that allows messaging while driving. The Placememo reminder application was developed for road inspectors who travel along the roads to keep doing maintenance work. The application was developed to help them better remember the defects they spot along the roads. This application couples an audio recording, typically a voice memo, to the location of the defect. The Placememo reminds the inspector of a particular defect by playing the memo when approaching the location of it. Accurate GPS positioning was chosen not only to make the road inspectors find the spot again, but also to support stopping by the spot in a safe manner. It was found that road inspectors often found it hard to stop at the location of defects. At the time they could see the defect, it was often too dangerous to break hard. As the inspectors operate in dense traffic, breaking hard would often risk a rear end collision with the cars behind. The Placememo prototype plays the memo some hundred meters in advance, depending on their speed, to allow bringing his or her vehicle to a stop safely.

4.4 Design requirements for Road Talk

The arguments above is summarised as a set of requirements for an improved service for messaging roadside-location events between drivers in traffic.

- **Sharing messages concerning roadside-location events:** The quantity of speed trap messages indicates that there are roadside location-dependent topics interesting for drivers to share with the use of mobile technologies.
- **Allowing a broad topic**: We interpret the fact that other types of information occurred as much as 12 % of the examined messages, as an indication that there are many other roadside location topics of interest for a driver. An improved service should allow as many types of messages as possible in a flexible manner.
- **Epidemic message distribution**: Speed trap messages use SMS and GSM mobile phones. SMS communication is costly and does not provide any service guarantees. However, based on mobile ad hoc networking technology we may achieve epidemic distribution that would be free and potentially very efficient given a certain level of deployment.
- **Accommodating driving**: Speed trap message systems are text based. Reading and writing text messages conflicts with driving. It is thus paradoxical that a service that distributes information about the roads is practically inaccessible while driving. The Placememo prototype is an example

of a design that allows interacting while driving. Essentially it showed that audio-based interaction is clearly feasible when driving and much preferred over text-based.

- **Accurate positioning:** Text-based messages are appropriate for planning a trip ahead, but can hardly be used to warn accurately when the driver is approaching a place for where there is a roadside event occurring. A design similar to Placememo would also be applicable for Road Talk. Still giving a user an opportunity to plan ahead and avoid hazards altogether is important aspect and should be reflected.

5 Implementation

The RoadTalk prototype is implemented in C++ for the Pocket PC 2003 operating system. The application has two major collections of software. The first collection concerns the GUI (Figure 3) and the other, the *Monger library*, deals with recording, playing and sharing of memos. As such, the latter implements the core concept of the system. The software of this library is modelled around a scheduler class to accommodate intricacies of positioning and networking as efficiently as possible. The main aim with the scheduler is to interface communication primitives of WLAN networking and GPS positioning for as much time as possible and as such quickly react on such events (Figure 4). For example, if there is data waiting on a particular networking socket, dealing with HTTP, the scheduler signals the class associated with this socket through a call back mechanism. This particular class is then responsible for interpreting the data and taking action in this case according to the HTTP specification. The called upon class must make sure that it does not block indefinitely as that would halt operation of the scheduler and break functionality of the application. In a similar non-blocking fashion the scheduler also checks network interface events (i.e. concerning addresses); serial port communication (reading Bluetooth GPS data); thread messages (mainly audio recording and playback events); and delayed operations (e.g. communication time-outs).

```
void Scheduler::runWhile( bool * test, DWORD inThread ) {
        Scheduler::inThread = inThread;
        while( (*test) ) {
                pollNetworkEvents();
                pollSockets();
                pollComPorts();
                threadMessages();
                pollSockets();
                expire();
                Sleep( 0 );  //yield to other processes
        }
        threadMessages();
}
```

Figure 4. The core of the Monger library and the scheduler class. Networking handler is scheduled twice per iteration as socket communication is deemed particularly vulnerable for delays.

5.1 Memo recording and playback

The Monger library does two things practically simultaneously to accomplish recording memos and associating them with locations. First, it records audio data encoded in the GSM10 format and continuously captures the position formatted in NMEA GPRMC [28] from the Bluetooth GPS. Second, when the user stops the recording, it is associated with the position where the user was located halfway through that activity. When approaching a location for where there is a memo recorded, the Monger library estimates the time to intersect the position. In this estimation Monger considers travelling speed and the length of the memo and times playback so that half of it will be played when coinciding with this position.

5.2 Memo sharing

The Road Talk application accomplishes memo sharing over wireless single-hop ad hoc networks using popular standards such as XML and HTTP. It is carefully designed to react on fleeting connection opportunities and transient links, caused by the extreme movement and the limited range of wireless radio transmitters. The application relies on a peer-to-peer system architecture, which practically means that each node is equally capable of exchanging and processing information. Furthermore, the system is fully distributed, that is, it requires no external infrastructure to operate. This approach allows a peer to establish network connectivity autonomously, which fits the brief constellation of nodes of the wireless ad hoc network. More so, the links over which these information exchange transfers are performed are not expected to be persistent or predictable. Nodes are not expected to stay in range of each other's transmitters for long. However, sometimes they do, which would allow for extended networking opportunities.

The RoadTalk implementation of memo sharing has two key components i.e. the Rapid Mutual Peer Discovery (RMPD) and the Monger Diffusion Scheme. The RMPD module is responsible for monitoring presence of other nodes in wireless reach and is designed to rapidly discover changes. Furthermore the aim is to provide mutual discovery. In other protocols, such as Jini, UPnP and SLP, discovery is strictly one-sided. These protocols are designed for a client discovering some resources, but not the other way around. The resources are never assumed to be interested in discovering the clients. The RMPD is based on repeated pings and whenever hearing a new ping a node echoes a 'pong.' A detailed description and analysis of RMPD can be found in [8, 23]. The Road Talk prototype includes an extension with the aim to ensure a mutual understanding in the event of asymmetrical wireless connectivity. This is important in wireless networking as two nodes with unequal transmitter range may cause asymmetrical connectivity i.e. one part may reach the the other, but not the other way around. The extension works as follows. First, assume two peers, X and Y, which experience asymmetrical connectivity. In this case peer X is able to hear pings from Y, but peer Y cannot receive pings from X. In this case the extension will have peer X *ignore* the pings from Y until peer X is certain that Y is able to hear its pings again. Peer X can determine this when X hears a 'pong' message from Y.

5.3 Monger Diffusion Scheme

The Monger Diffusion Scheme synchronises memos among peers in a wireless ad hoc network as a two-step process. First, whenever RMPD discovers a peer it downloads a particular XML file from it. The download is performed with HTTP. Consequently each peer contains a HTTP client to query for files, as well as a HTTP server to offer them. The client and the server implement a limited version of HTTP/1.0 tailored for the unreliable nature of wireless ad hoc networking. The downloaded XML file is essentially a database of the memos stored on that peer. However, the database does not contain the actual audio data that make up the memos. They are stored separately as GSM10 encoded '.wav' files in a server directory. When the download operation completes the Monger platform loads and examines the XML file. This load operation is performed with our port of the publicly available EXPAT XML parsing library [27]. The loading operation concludes the first phase of the synchronization. Second, the monger downloads each memo it does not currently store. Whenever a download of such memo completes, a reference to it is included in the local memo database and the file is copied to the server root. As such these files are ready to be shared to a "third party."

The Monger Diffusion Scheme bears some resemblance to SPIN-1, which is an algorithm for information diffusion in sensor networks [14]. However, SPIN-1 uses a push/pull approach. In SPIN-1, whenever it acquires a new entity, a node pushes the meta-information describing the information entities it has. The entity may come from either the sensor or other nodes. All nodes hearing this advertisement then "pulls" the entities they are missing. As such new information is propagated throughout the sensor network by "hopping" from node to node. The SPIN-1 algorithm is not appropriate for Road Talk. It is very likely that new data is produced out of range of any other node. A push operation under such circumstances would then be wasted, as any other would not receive it.

6 Initial prototype performance evaluations

It is not obvious that Road Talk is technically feasible. After all, the hardware we used here is intended for indoor office settings where people at most walk about. Therefore, we decided to examine the prototype performance in terms of sharing memos in mobile situations. First we recorded 22 memos (the locations of the memos are depicted in Figure 3) together with the memo database constituting about 92 kilobytes[a] of data. We kept one PDA with Road Talk stationary and drove by in various speeds with another installed in a car. The PDAs were used "as is" and *no* extra equipment, such as range extending antennas, was used. In the first ten drive-bys we kept speed to about 50-60 km/h and cleared the stationary PDA of downloaded memos for every second run. Altogether we had five complete downloads and one failure. This also means for four drive-by runs there was no download operation simply because the synchronization was complete. We then increased the speed to about 80 km/h and cleared the stationary Road Talk of memos for each complete download. At this speed the download immediately failed once. At the time we believed this was due to some signal path problems caused by an elevation in the road. When we moved the stationary PDA about 100 meters to the crest of it, downloading worked for three straight runs at this velocity. We then increased speed to about 100 km/h and performed a successful download.

We have also logged some RMPD events, such as node discovery and connection loss. We found two cases of RMPD behaviour. First, in eight test runs there were two RMPD discovery and loss event pairs for a single drive-by. This means that discovery was established and then connection was lost and after a few seconds this repeated (Figure 5a). Second, in five drive-bys the there was an asymmetrical number of events. This means that one peer experienced two pairs of RMPD discovery and loss events, but the other experienced only one, (Figure 5b). When we moved the stationary PDA, the phenomenon of patchy and asymmetrical sessions did not occur (Figure 5c).

When we examined various parts of the Monger library we realized that some of them performed poorly. Synchronizing memos includes parsing an XML file and then iterating over the memos it indexes. This operation constituted a processing bottleneck, because when performing it, networking tasks that let the other PDA to complete its synchronization were not attended. We believe that the patchy and asymmetrical behaviour of RMPD in the test could be due to signal path propagation obstructions *in conjunction with* timing of parsing and iteration. Moving the stationary PDA essentially made the timing different. Monger dealt with RMPD messages and performed HTTP downloads when

---

[a] The Memos contained about 0.5 to 3 seconds of audio data encoded in the GSM10 format. This format is very storage efficient, for example 3 seconds of audio takes about 7 kilobytes.

there was actual network coverage. Therefore, we inserted calls to code that attend networking also from within parsing and iteration over memos.

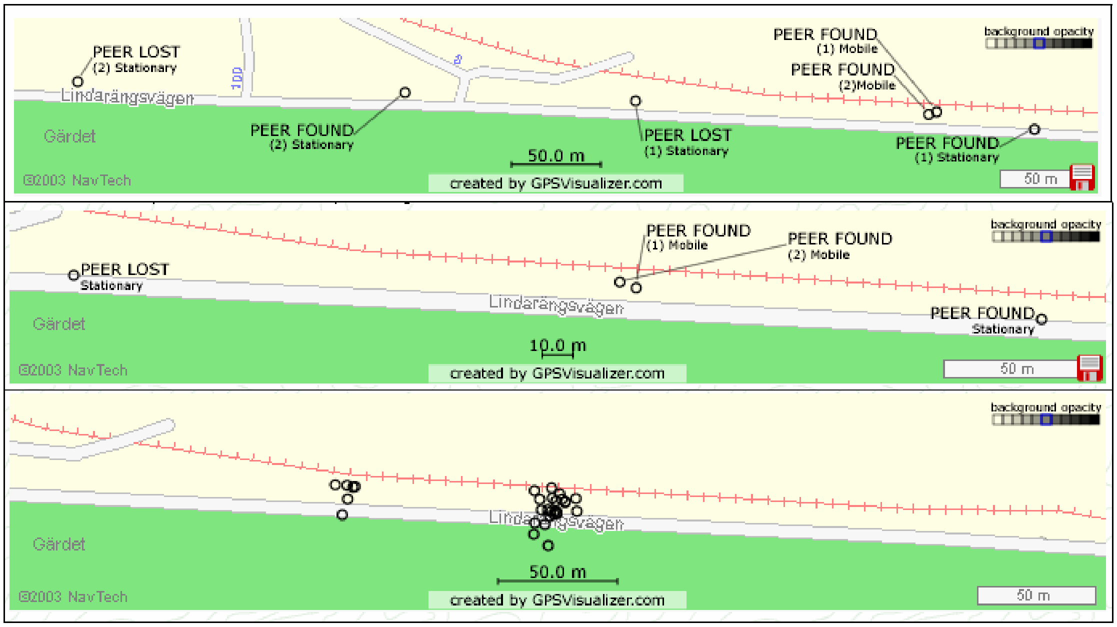

Figure 5. .Patchy RMPD behaviour (a). Asymmetric discovery (b). Approximate location of stationary PDA (c).

The logged RMPD events were also coupled to the time and position where they occurred. We could not obtain an accurate GPS position fix for all the test runs and therefore the analysis presented here is based on the 13 runs where we did. On average the time elapsed between RMPD discovery and connection loss session is about 17 seconds. The minimum distance between discovery and connection loss is about 120 meters occurring at a velocity of 53 km/h. The maximum distance of about 520 meters occurs at 90 km/h. We also noted that this range seemed to increase with speed, which is an effect of that RMPD uses a near constant time-out to determine when connection has dropped. Naturally, during a constant period of time a peer travels farther at higher velocity. An improved discovery protocol should somehow take the relative speed of nodes into account when determine this time-out.

7 Discussion

Although the prototype development is still in an initial phase and need further refinement, both in terms of interaction design and technical performance, we argue for that the current Road Talk prototype is well grounded in the issues raised in the motivation. First, obviously Road Talk is designed to allow sharing of messages concerning roadside location events. The initial prototype performance evaluation indicated that sharing is possible during a wide range of speeds. Second, memos consist of voice recordings and distributed in their original form. There is no discrimination of topic and a user may mediate any issue, such as warning for hazards, to the fellow driver. Third and concerning epidemic distribution: the application supports exchange of memos immediately from one car with Road Talk to the next and also indirectly, from one car to the next, via a third car. As such the memos would spread epidemically following the physical movement of vehicles and spread quickly wherever and whenever traffic would be dense. Fourth, we argue that Road Talk accommodates interacting while driving. Drawing on the design of Placememo, the application plays memos automatically when approaching the location associated with it and requires no other intervention besides driving. Finally, Road Talk also makes use of accurate GPS positioning to capture the location of a memo and continuously track the position to determine when approaching one. Besides, Road Talk supports planning ahead as the application also features a zoom-able map with all memos (stored on the local PDA) are plotted as squares. A user may click on the squares to hear the corresponding memo.

8 Future work

The prototype performance evaluation showed the current Road Talk prototype is stable enough to get additional feedback through field experiments with users. The RMPD proved to feature some limitations and message sharing suffered somewhat from problems of patchy and asymmetric discovery patterns. These problems are due to signal path problems and timing of network operations, but can be approached with optimized software performance and using range-extending antennas (mounted on the outside of the car). We are working on inductive antennas as most PDA with built-in wireless networking interface lack ports for external ones. We are also working on staging the application on mobile phones, which would make the concept more appealing for a larger audience and thus more manageable in a large-scale user evaluation. Finally, we are also further investigating the speed trap services to learn more about how these messages are produced and how successful they are at warning other drivers.

9 Conclusion

The state of the upcoming road is desirable to know in advance and a system to deliver information about such state would aid traffic safety. Here we have demonstrated Road Talk a mobile message system that let drivers communicate about roadside-location events such as hazards or situations that deserve extra caution. The contribution to mobile message systems is two-fold. First, Road Talk is targeted for a previously unexplored context. It introduces safe messaging among drivers. Second, Road Talk explores epidemic distribution of messages. This is a process where messages spread as cars physically move and coincide at a certain distance with other cars. Given a certain level of deployment this method would be very efficient in spreading messages. We have also confirmed that the Road Talk prototype is able to share messages in a wide range of speeds in a performance study. Finally, the work also contributes to the understanding of context-aware computing by a statistical study of commercially available mobile message systems for distributing warnings about speed traps. As much as 12 % of the messages concern other topics which motivate motivates designing system supporting roadside location-dependent messaging. Furthermore, speed trap messages occur when and where traffic is dense, which fits with epidemic distribution.

Acknowledgements

We would like to thank the members of the Mobility studio at the Interactive Institute for valuable input to the design of Road Talk and writing this paper. Particularly, thank you Mattias Esbjörnsson for your practical contributions to the performance evaluation. The Swedish Foundation for Strategic Research funded this work through the Mobile Life project.

References

[1] Bull, M. (2004). Automobility and the Power of Sound. In Journal of Theory, Culture and Society, Volume 21, Issue 4, SAGE.

[2] Burell, J. and Gay, G. (2002). E-graffiti: evaluating real-world use of a context-aware system. In Interacting with Computers 14, Elsevier Science.

[3] Borghoff, U. and Sclichter, J. (2000). Computer-Supported Cooperative Work: Introduction to Distributed Applications. p 120, Springer.

[4] Bradner, E. Kellog, W. and Erickson, T. (1999). The Adoption and Use of 'Babble': A Field Study of Chat in Workplace. In Proceedings of European Conference on Computer-Supported Cooperative Work, Dordrecht, The Netherlands.

[5] Cerf, V. Burleigh S. Hooke, A. Torgerson, L. Durst, R. Scott, K. Fall, K. and Weiss, H. (2003) Delay-tolerant networking: an approach to interplanetary Internet. In Communications Magazine, IEEE, Volume 41, Issue 6.

[6] Esbjörnsson, M. (2005). From Ethnography on Infrastructure Management to Initial User Feedback on PlaceMemo. To appear in Personal and Ubiquitous Computing: Theme issue on Interactive Mobile Information Access.

[7] Esbjörnsson, M. and Juhlin, O. (2003). Combining Mobile Phone Conversations and Driving - Studying a Mundane Activity in its Naturalistic Setting. In Proceedings of Intelligent Transportation Systems, Madrid, Spain.

[8] Esbjörnsson, M. Juhlin, O. and Östergren, M. (2004). The Hocman Prototype – Fast Motor Bikers and Ad Hoc Networking. In Proceedings of Conference on Mobile and Ubiquitous Multimedia, Oulu, Finland.

[9] Esbjörnsson, M. Juhlin, O. and Östergren, M. (2004). Traffic Encounters and Hocman: Associating motorcycle ethnography with design. In Journal of Personal and Ubiquitous Computing: Special Issue Mobile Human-Computer Interaction, Volume 8, Springer Verlag.

[10] Espinoza, F. Persson, P. Sandin, Anna. Nyström, H. Cacciatore and E. Bylund, M. (2001). Social and Navigational Aspects of Location-Based Information Systems. In Proceedings of Ubicomp, Atlanta, Georgia, USA.

[11] Glance, N. Snowdon, D. and Meunier, J.-L. (2001). Pollen: using people as a communication medium. In Computer Networks: Special Issue Pervasive Computing, Volume 35, Issue 4, Elsevier Science.

[12] Grinter, R. and Eldrige, M. (2001). y do tngrs luv 2 txt. In Proceedings of European Conference on Computer-Supported Cooperative Work, Bonn, Germany.

[13] Griswold, W. Boyer, R. Brown, S. Truong, T. M. Bhasker, E. Jay, G. and Shapiro, B. (2002). Using Mobile Technology to Create Opportunistic Interactions on a University Campus. Presented at Ubicomp 2002 Workshop on Supporting Spontaneous Interaction, Göteborg, Sweden.

[14] Heinzelman, R. Kulik, J. and Balakrishnan, H. (1999). Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In Proceedings of Conference of Mobile Computing and Networking, Seattle, Washington.

[15] Isaacs, E. Walendowski, A. Whittaker, S. Schiano, D. and Kamm, C. (2002). The Character, Function and Styles of Instant Messaging in the Workplace. In proceedings of Computer Supported Cooperative Work, New Orleans, Louisiana, USA.

[16] Juhlin, O. (1999). Traffic Behaviour as Social Interaction - Implications for the Design of Artificial Drivers. In Proceedings of 6th World Congress on Intelligent Transportation Systems, Toronto, Canada.

[17] Khelil, A. Becker, C. Tian, J. And Rothermel, K. (2002). An Epidemic Model for Information Diffusion in MANETs. In Proceedings of Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems, Atlanta, Georgia, USA.

[18] Laurier, E. (2002). Notes on dividing the attention of a car driver. Team Ethno Online.

[19] Masui, T. and Takabayashi, S. (2003). Instant Group Communication with QuickML. In proceedings of GROUP, Sanibel Island, FL, USA.

[20] Nardi, B. Whittaker, S. and Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. In proceedings of Computer Supported Cooperative Work. Philadelphia, PA, USA.

[21] Norton, J. and Bass, F. (1987). A diffusion theory model of adoption and substitution of for successive generations of high-technology products, Journal of Management Science, Vol. 33, No. 9, pp. 1069-1087, USA.

[22] O'Neil, J. and Martin, D. (2003). Text Chat in Action. In proceedings of GROUP, Sanibel Island, FL, USA.

[23] Östergren, M. (2004). Sound Pryer: Adding Value to Traffic Encounters with Streaming Audio. In Proceedings of Entertainment Computing, Eindhoven, The Netherlands.

[24] Persson, P. and Fagerberg, P. (2002). GeoNotes: a real-use study of a public location-aware community system. Swedish Institute of Computer Science Technical Report T2002:27, ISSN 1100-3154.

[25] Redovisning av verksamhetsmått/faktauppgifter och nyckeltal. (2005). http://www.stockholm.se/files/64500-64599/file_64558.pdf

[26] Stockholm 2004 sifferguide. (2005). http://www.stockholm.se/files/71900-71999/file_71939.pdf

[27] The Expat XML Parser. (2005). http://expat.sourceforge.net

[28] The NMEA FAQ. (2005). http://www.ziegler.desaign.de/GPS/nmea.txt.

[29] We make money not art. (2005). http://www.we-make-money-not-art.com