Mobile Pollution Mapping in the City

Paul Rudman, Steve North, Matthew Chalmers
Computing Science, University of Glasgow
matthew@dcs.gla.ac.uk

ABSTRACT

Mobile computing has the potential to allow both experts and the public to collect and understand environmental data such as pollutants in urban areas. We describe an experimental system—eGS—that allows users to explore a city area while collaboratively visualising a common atmospheric pollutant— carbon monoxide—in real-time. Users carry a networked tablet PC. Using GPS and an attached sensor, a map shows pollutant values as a colour-coded trail as the user moves around the city. Users may take photographs of pollution-significant situations that are referenced against their current map location. Pollutant readings and photographs appear on all users' maps as shared information for potential collaboration. We report on lessons learned and design issues arising from the implementation and us of this research prototype. In particular, we question some assumptions regarding the use of map-based representations with transient environmental information.

1. INTRODUCTION

Mobile computing has the potential to allow both experts and the public to collect and understand environmental data such as pollutants in urban areas. In this paper, we describe an experiment that allowed users to explore a city area while collaboratively visualising a common atmospheric pollutant—carbon monoxide—in real-time. We report on lessons learned and design issues arising from the implementation and use of our research prototype.

Carbon monoxide is usually generated when combustion takes place, e.g. cigarettes, gas fuelled cooking and heating systems, portable petrol–driven generators. In the urban environment, however, it is road traffic that produces the bulk of CO [4, pg. 47]. While one motor vehicle with its engine idling may produce several hundred parts per million (ppm) at the exhaust, the gas dissipates rapidly. The current recommended maximum exposure to CO is a mean value of 10ppm over a 10–hour period [4, pg. 35], and this level is therefore of particular interest within environmental science. On a windy day there is relatively little overall exposure to CO caused by traffic in open space. The problem is that CO may build up to above the 10ppm limit given appropriate geography (e.g. narrow roads and high buildings), meteorology (e.g. still air) and traffic flow (e.g. continual queuing and movement at traffic lights). While fatal levels would not be expected in open air, 35ppm plus may be found in some circumstances [9].

Most CO research to date has used a small number of large (e.g. $2m^3$) fixed sensors placed sparingly throughout a city, collecting data and storing it locally for manual collection. Analysis gives general or coarse–grained CO estimates for a large area. This, however, does not necessarily represent any one individual's exposure to the gas in his or her daily life. In order to quantify this, more fine-grained data gathering is necessary. Recently, this has been made possible by the development of small, portable, accurate CO sensing devices. One such sensor is manufactured by City Technology Limited [3]. It provides measurement of carbon monoxide from 0 to 500ppm with a resolution of 0.1ppm. Response time is quoted as less than 40 seconds to register 90% of a change in CO concentration. The sensor itself requires associated electronics, and a complete CO monitoring unit has been developed by Learian [6]. The unit is reasonably compact (11 x 8 x 4 cm), and researchers at University College London [7] have obtained large data sets of geo-referenced CO readings from people, each of whom carried a CO monitor, GPS and data recorder while cycling or walking around London. Post–collection analysis was then used to determine recurring CO features.

The work reported here goes further, by supporting collaborative visualisation of data readings on tablet PCs. Overall, our motivation was to develop an early pilot system to explore how scientists as well as the public might use such technologies to do both 'official' environmental science and the provision of personally (or publicly) relevant pollution information. Given that environmental scientists currently have only relatively coarse–grained data, we were interested in supporting people working together to find and collect pollution data in a fine–grained way. A second motivation relates to greater public understanding of science—or even public *involvement* in science. Given the trend for sensor miniaturisation and for ubiquitous wireless communications via mobile phones, it is possible that far greater public involvement in environmental science could take place. Rather like the valuable side–effect of image labelling created by people playing games described in [8], we suggest that useful collective data might be collected by either explicit or passive recording of geo–referenced pollution data, in the places that people actually go—rather than where scientists expect them to.

The next section describes the system, *eGS,* which we built as an initial exploration of these concepts and aims. Later sections outline the evaluation of the system, our initial findings, and more general discussion of our experience, design lessons and possible future work.

2. THE eGS SYSTEM

The system is intended to allow users to collect and visualise carbon monoxide readings over a city area, either individually or in collaboration with other system users. Each user carries a lightweight tablet PC. This provides him or her with a map of the local area on which all the users' locations are marked—each user's location is taken from a GPS device attached to the tablet. Once per second, a coloured dot appears on the map for each user, indicating the current carbon monoxide level at his or her location. An attached camera allows the user to take photographs, which are shown to all users, both as a temporally ordered 'filmstrip' and as thumbnails on the map. Each thumbnail is placed at the map location corresponding to where the photograph was taken. In some of our trials, voice–over–IP was used to allow continuous audio communication between users.

The software is based upon an existing system, *George Square* [2], developed as part of work investigating collaborative tourism. George Square supported co-visitors, e.g. one in a city square and one in a cafe, to share locations, photographs and audio. In George Square, users also shared recommendations of objects and locations of interest, but in eGS ('eScience George Square') co-investigators share CO readings.

Behind the application software lies the EQUIP middleware [5]. This is a distributed tuple space system that facilitates the transfer of data between peers without the requirement for a central server. Discrete events, such as a GPS reading or a photograph URL, are passed between applications, regardless of the machine on which they reside—provided a suitable network connection is available. Furthermore, EQUIP supports self-discovery of tuple spaces, and so new instances of eGS—being run by new users—are able to receive and share data with others with no explicit setup procedure.

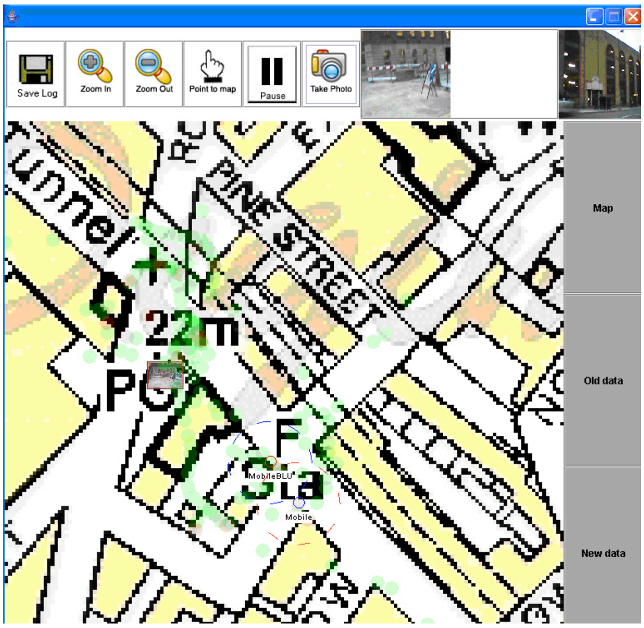

Figure 1. The eGS system's interface centres on a map of the local area that shows geo–referenced sensor readings and users' locations. A set of buttons (top left) lets the user save a log of sensed data, zoom the map in or out, control a telepointer, pause sensor recording and take photographs. Photographs taken by all the users are shared in real–time and shown in a filmstrip (top right) and on the map.

Figure 1 shows an example screenshot. The map takes up the bulk of the screen area. In this example the red and blue circles locating two users are visible (just below the central letter "F"), along with the users' names. One photograph has been taken, which appears on the filmstrip (beside the "take photo" button) and on the map (below the "22m" text). The green dots are CO readings. At the top of the screenshot are buttons allowing the user to save the CO log file (not used during these evaluations), zoom the map in or out (panning was achieved by dragging the map with the stylus), initiate a telepointer (which generates a large coloured dot at a selected location for a few seconds on the other users' maps), pause the taking of CO readings (used at the start of each trial while the GPS 'settled down') and take a photo. The far right of the filmstrip is the camera's viewfinder. Buttons down the right-hand-side were not used during these evaluations. Lastly, as each user walks through the area, a trail of coloured dots is created, one new dot per second, showing a history of CO levels. Thus, users could see and share locations, CO readings and geo-referenced photographs.

Each participant carried an NEC Versa T400 tablet computer, with a SysOn CFplus compact flash GPS, a USB 'stalk' camera, and a Learian CO monitor connected via a serial to USB cable. Since the CO monitor was too large to attach to the tablet, it was either placed in a small rucksack worn by the participant, or held by a second participant. The tablet's internal 802.11 wireless provided connection to a base station. Finally, a microphone and headset could be connected to the tablet's audio sockets.

An 802.11 network was required to allow communication between multiple tablets. Experience from the George Square work suggested that line-of-sight was essential to maintain a continuous connection, and so a wireless network access point (AP) was placed high on a rooftop, reducing signal interruption by obstacles, passing vehicles and so forth. We only required the network as a relay between units (and would not use any Internet bandwidth in this experiment) and so we built a self-contained mobile LAN comprised of a D-Link 802.11 AP, two 12 volt sealed lead-acid batteries (in parallel) and a voltage regulator. The AP could be used with either the original antenna or an external high gain antenna. The batteries last about five hours. This system was used to provide network coverage during the studies.

3. EVALUATIONS

Based on our discussions with experts at University College London, the following potential user exercises were identified:

- Compare participants' existing perceptions of pollution to actual CO in an area
- Find moving CO peaks—it is suggested that accumulated CO may drift through city spaces, but this has yet to be observed
- Identify 'lines' of the same CO level—lines denoting the distance from a CO source where the CO value is consistent. This requires still air to be meaningful.
- Investigate local CO variations—differences between sides of a street, sides of a pavement, around corners, etc.
- Validate places for fixed sensors—investigate the local area to see if a sensor is at a representative location

Note that all these tasks require low wind speeds, since at medium or high wind speeds (e.g. above 10 metres per second (22 mph)) the spatial structure of the CO changes too quickly to be measured accurately, with the gas often being dispersed.

3.1 One tablet, two users

Initial evaluations used a simplified version of the system whereby two participants shared one tablet PC The location used was Euston Road, W1, London, between Euston Square and Warren Street underground stations, and Tottenham Court Road from Warren Street to Goodge Street stations. This area was well known to the participants, who were students with no particular environmental science knowledge. Initially, the participants were given a (paper) map of the area and asked to mark areas they considered to be highly polluted. One participant was then given a tablet PC while the other carried the CO monitor. They were asked to walk through the defined area, observing the CO levels as reported by the system. After the walk, the participants were given another (paper) map of the area and asked again to mark areas they considered to be highly polluted.

A colour scheme for sensor readings was chosen that would show small changes in CO. Colours were chosen to stand out clearly, both when changing between adjacent colours and when seen on the tablet PC's display. (Defining the relationship between individual colours and CO levels was problematic, and is discussed in detail later in this paper.) The colours used were as follows:

CO (ppm)	Colour
<1.0	BLACK
1.0 - 1.1	ORANGE
1.2 - 1.3	GREEN
1.4 - 1.5	BLUE
1.6 - 1.7	BROWN
>1.7	RED

3.2 Two tablets, two users

In this configuration, two participants were each given a tablet PC running eGS, and a small rucksack for the CO sensor. Our wireless network allowed the tablet PCs to be interconnected, allowing CO readings and photographs to be shared immediately and automatically between users. We used two locations for this configuration. The first location was Farringdon Road, WC1, London, between Rosebery Avenue and Clerkenwell Road. Here, pairs of participants were given the task of: comparing sides of the road for differences in CO. A wireless network was created by placing the mobile LAN on the top floor of a multi-story car park (some five stories in height). By placing the external antenna horizontally over Farringdon Road the whole visible area was covered with few interruptions.

The second location was, initially, the square by Clerkenwell Road and St. John Street, London. Pairs of participants were then given 45 minutes to roam as they chose. Since their initial location was on the edge of the system's map, it was suggested to the participants that they may like to begin by walking more into the centre of the mapped area. Their task was to find high and low areas of CO concentration (a variation on the CO lines task). A wireless network was created by placing the mobile LAN in a bag carried by a researcher. This allowed the participants to walk where they chose, either together or apart (but within visible distance of the researcher). A colour scheme for sensor readings was chosen that would show up changes considered to be scientifically relevant. Colours used were as follows:

CO (ppm)	Colour
<5.5	BLUE
5.5 – 9.9	GREEN
>=10	RED

4. FINDINGS

We present a qualitative analysis of the initial evaluations of the system outlined in the previous section, and describe lessons learned for the design of such systems.

4.1 One tablet, two users

Prior to using the system, participants tended to see CO as being directly proportional to traffic. Thus, on the (paper) map they marked large roads and junctions as most polluted and small roads as having low pollution. Use of the system showed that this was too simplistic, as seen in Figure 2. The maps they made showed increases in CO around relatively enclosed junctions. Euston Road, which is at that location broad and open with free-flowing traffic, gave low CO values. Tottenham Court Road, however, is narrower, enclosed and has stationary traffic, generating higher readings.

After using the system, participants showed an improved understanding of the relevance of stationary traffic in (relatively) enclosed spaces as being an important factor in CO build-up. For example, one participant said:

> "Before, we said this area [Euston Road] is the most polluted area because it's a crossroads, but when we tested that area... it's quite a huge area and the wind is quite strong... it disperses quickly, but around the [Warren Street] tube station, because there's big buildings when we stood here the wind was less strong so we got an orange point [higher reading]"

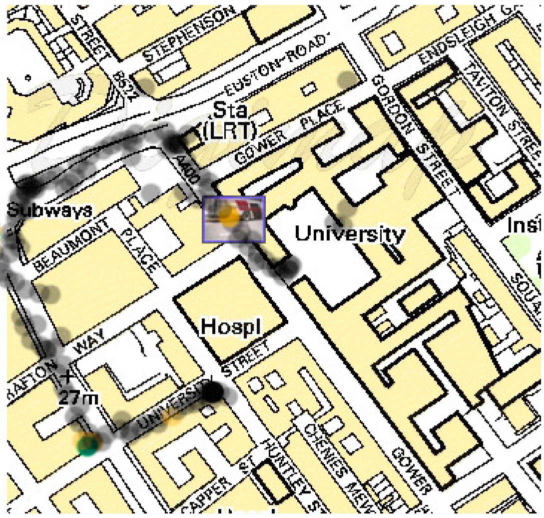

Figure 2. A participant's map after the 'user perceptions' scenario.

4.2 Two tablets, two users

Figure 3 shows participants' CO readings after the first scenario, investigating CO levels on opposite sides of the road going from top left to bottom right in the figure. It does not, however, show two lines of dots (CO readings), one line on each side of the road. There are two problems here. The first is GPS inaccuracy, which has moved the dots randomly up to one road-width in any direction away from the road the participants walked on. This problem is discussed at length in Section 4.3. A second issue was that users' readings (dots) appeared with no clear way to distinguish which individual had recorded them. Participants found this confusing, as they wished to retain information on who had recorded which readings. This was at odds with our expectation of more objective or objectifying scientific–style observation in which the identity would be of little concern.

Since the colour scheme was changed from the single system runs to make the readings more scientifically relevant, the displayed level tended to show the lowest level (blue) most of the time. Unlike with the single system, participants did not notice when higher readings did appear (there are some green dots around the junction with Baker's Row). This suggests that if CO level changes are not frequent, the user needs to be clearly alerted when change does happen—clear colour changes or an alarm, for example.

Figure 3. A participant's map after the “road sides” scenario

In the second scenario, participants were allowed to chhose the direction in which they walked. It was assumed that they would wish to remain in the area for which the system had map details. This was not the case, as can be seen in Figure 4. This was unexpected, since in the George Square tourism study participants used the shared map as an important aspect of their shared experience. This issue is discussed further in Section 4.4.

Figure 4. Participant's map after the “free roaming” scenario, showing significant movement away from the area with map data.

4.3 Practical problems

As with the earlier George Square system, the hardware worked well, apart from a significantly reduced battery life on the tablets (from 90 minutes to barely 60), probably due to the near-zero temperatures during the evaluations. While the CO monitor unit was not bulky as such, it was not small enough to attach to the tablet PC. It is likely that at some time in the future a small CO sensor will become available that can be attached to the PC or on the body, in the same way as the GPS unit has evolved from an external wired box to a CF card or a Bluetooth box. The headsets were regarded by participants as preventing them from being aware of potential dangers (traffic, other pedestrians, muggers, etc.). It was also very difficult to hear speech over the headsets due to the ambient noise. Therefore the headsets were discontinued early on. Participants instead were encouraged to walk apart where the task required this, and to meet regularly to discuss the task in person.

There is an unresolved question over the sensitivity at which the system should be set. With the single system, the sensitivity was set to pick up almost any increase in CO. (The system began showing colour changes at 1.0 ppm, from a background level of approximately 0.5 ppm and subsequent colour changes every 0.2 ppm up to 2.0 ppm.) These colour-value mappings allowed participants to see small increases in CO. While these may be considered genuine readings—as opposed to noise from the sensor for example—the problem is that from an environmental science or public health perspective these values have little significance. 1.0 is very close to the baseline value of 0.5, in comparison to the 10.0 ppm target value for maximum exposure. A value of 1.0 may easily be obtained by a small number of vehicles and is unlikely to represent a significant build-up of CO. Setting the system at a lower sensitivity (colour changes at 5.0 and 10.0 ppm) on the other hand has scientific relevance, but is much less likely to occur. A value of 5.0 ppm requires a build-up of CO that requires appropriate geography, meteorology and traffic flow. This is not easy to predict, and may not occur for several days, making user trials difficult. Probably the best solution is to set the sensitivity according to the task, as has been the case here.

The eGS system is intended to expand the user’s experience of their surroundings by providing real-time information about one type of current pollutant level—carbon monoxide. This assumes that an accurate reading may be obtained once per second, but in these evaluations this was not always the case. We found the situation where a participant walked towards a busy junction with a light wind behind them—no change in CO occurred. As they turned the corner (out of the wind) the CO value began to rise, going through the colour levels over a period of a few seconds. It is possible that the CO level was slowly rising. However, there is a suspicion that this slow rise may be an artefact of the sensor itself. The sensor is rated as reaching 90% of a new CO value in less than 40 seconds. This is very good for most situations. However, the eGS system implies to the user that every second a correct reading is given (not that they should ignore readings until they ‘settle down’). It may therefore be necessary to inform future participants of this factor.

GPS accuracy in the ‘urban canyons’ of London streets was an issue, particularly in the dual-system evaluation. Figure 3 shows the problem. The participants each walked up one side of the road. What is not obvious from the static map is that the dots on one side of the road were not all generated while the person was on that side. The width of Clerkenwell Road at this point ranges from approximately 15 to 19 metres. Thus, even a ‘good’ GPS reading cannot reliably determine which side of the road the person is located on. However, the participants appeared to circumvent GPS accuracy problems. They observed the map as the readings appeared, seeing each dot as representing their current reading, even though it might not be perfectly accurately placed.

Fitting with this view is the participants’ behaviour when given free rein to walk where they chose. As can be seen from Figure 4, the participants in this trial run were not worried about leaving the area covered by the on-screen map. When these participants were later asked why they had not walked so as to remain on the map, they replied that they could still see the colours (CO values) as they appeared, so didn’t need the map. Thus, we are reminded that participants were asked to explore CO readings in the city, and the map was merely an aid to this task. Getting the ‘right dots in the right place’ was not necessarily part of the task, as they could verbally recount what they had discovered. For more persistent records or other forms of sharing and discussion, it may be that self–reporting of location [1] may be a useful means to allow users to correct GPS errors in situations that they consider to be significant.

4.4 The map

Carbon monoxide propagation through the atmosphere is complex and very difficult to model accurately. As stated earlier, traffic is the main source of CO. The obvious feature of traffic is that, overall, it moves. Thus, not only is CO movement difficult to predict but the source is continually moving. Where CO is analysed retrospectively, this movement can be adjusted to give typical CO values in each location [7]. However, using the eGS system, only one value remains per reading location. This generates local effects that would be removed by a longitudinal study.

For example, participants using the single system walked past a stationary car with its engine idling. There was no perceptible wind at that time, allowing CO to build up around the vehicle. Thus the system registered a higher reading at that point. On the map then, a small area of an otherwise low CO road shows a higher reading. (A similar increase has been seen in early system tests when walking past people smoking a cigarette.) This does not mean that that area would generally be expected to have higher CO values, only that it did so on that occasion. Conversely, it is possible to stand at a busy junction and register no increase in CO values. This may be due for example to wind blowing the CO away from the observer, or—given the complexity of air movement—for no apparent reason. Again, a fixed sensor at that point would, over a period of time, register that location differently as it would then be presenting average CO levels that would be higher than that particular low reading.

From an HCI viewpoint, this calls into question the use of a map as visualisation of the CO data. The intention was to allow the user to build up a picture of CO levels in a given area, using the map as a representation of the physical area. By displaying CO levels as colours on the map at the locations they were generated, it was intended that the user could understand of CO in the physical area using the map representation. With the George Square system this worked well, as the participants there were able to use the shared map when discussing features in the physical landscape.

With the system described here, CO values become the main features, even though the map no longer fully or precisely represents that landscape. The reason for this is that the amount of traffic present at the time each location was visited is crucial to the perception and experience of that place. Thus, a location is relevant here not just as a place, but as a place and time with its associated traffic density, something not directly represented by a map. The participants' use of photographs as reminders of the traffic situation at particular locations did ameliorate this problem to some extent. As photographs appear on the map where they are taken, this augments the map with time/traffic information. It would, however, be advantageous if this process could be automated in some way, to provide a consistent set of reference photographs related to the readings. This would assist users in using the dot-trail as a recording of their experience, connecting CO values to place more explicitly. It could possibly further sideline the map. However, this may be task-specific. Where the participants walked off the edge of the map their task was to find high and low areas of CO. Had they have been asked to find lines of equal CO, or moving CO peaks, or to create a permanent veridical representation of CO distribution, they may have considered the map of greater relevance. Further, more extensive studies will be required to answer these questions.

5. CONCLUSION

The pilot system described here has revealed a number of assumptions and issues with regard to the system's design and use. In retrospect, the participants treated the system in a way that suited the tasks and technologies at hand. Especially in the face of GPS errors, they focused on the readings that they knew best—their own—and did not attempt to make a common and precise map of CO coverage. They used the camera provided in the system to help record and recount their experience, and wished for individually identifiable readings that might have helped further that aim.

Furthermore, as we begin to look towards smaller and more numerous mobile devices for recording and presenting pollution, e.g. attached to or integrated with mobile phones for example, we suggest that it may not be appropriate to expect that collective construction of a shared and 'perfect' model of CO across a city will be the aim of every user. We do still consider that contribution to a collective data set as a side–effect of individual system use may afford useful new scientific analysis as well as a variety of uses by the public. However, further research will be needed to find specific or customisable tools for individual members of the public that will motivate them to offer up passively or explicitly recorded data for the 'common good'. Clearly, a vital element of such a system will be the protection of contributors' privacy, although we suggest that maintaining individuals' ability to retrieve or identify their own data may nevertheless be an important design requirement.

6. REFERENCES

1. Benford, S., Seagar, W., et al., The Error of Our Ways: The Experience of Self-Reported Position in a Location-Based Game. in *Proceedings of Ubicomp*, (Tokyo, 2004), 70-87.
2. Brown, B., Chalmers, M., et al., Sharing the square: collaborative visiting in the city streets. in *To appear in: Proceedings of ECSCW 2005*.
3. CTL, A3CO EnviroceL, http://www.citytech.com/PDF-Datasheets/a3co.pdf.
4. DEFRA. The Air Quality Strategy for England, Scotland, Wales and Northern Ireland, Department of the Environment, Transport and the Regions, Norwich, 1999.
5. Greenhalgh, C. Equip: A Software Platform for Distributed Interactive Systems *Equator Technical Report*, University of Nottingham, 2002.
6. Learian, Learian Mini Logger--Specifications, http://www.learian.com/pdfs/icom.pdf.
7. Milton, R. and Steed, A., Correcting GPS Readings from a Tracked Mobile Sensor. in *To appear: International Workshop on Location- and Context-Awareness (LoCA 2005)*, (Munich, Germany, 2005).
8. von-Ahn, L. and Dabbish, L., Labeling Images with a Computer Game. in *Proc. ACM CHI*, (Vienna, 2004), 319-326.
9. WHO, Carbon Monoxide, http://www.inchem.org/documents/ehc/ehc/ehc213.htm.