

Instrumenting the city: developing methods for observing and understanding the digital cityscape

Eamonn O'Neill¹, Vassilis Kostakos¹, Tim Kindberg², Ava Fatah gen. Schiek³, Alan Penn³, Danaë Stanton Fraser⁴ and Tim Jones⁴

¹ Department of Computer Science, University of Bath, UK, BA2 7AY
{eamonn, vk}@cs.bath.ac.uk
<http://www.cityware.org.uk>

² Hewlett-Packard Laboratories, Bristol, UK, BS34 8QZ
timothy@hpl.hp.com

³ The Bartlett, University College London, UK, WC1E 6BT
{ucftajf, a.penn}@ucl.ac.uk

⁴ Department of Psychology, University of Bath, UK, BA2 7AY
{D.StantonFraser, t.jones}@bath.ac.uk

Abstract. We approach the design of ubiquitous computing systems in the urban environment as integral to urban design. To understand the city as a system encompassing physical and digital forms and their relationships with people's behaviours, we are developing, applying and refining methods of observing, recording, modelling and analysing the city, physically, digitally and socially. We draw on established methods used in the space syntax approach to urban design. Here we describe how we have combined scanning for discoverable Bluetooth devices with two such methods, gatecounts and static snapshots. We report our experiences in developing, field testing and refining these augmented methods. We present initial findings on the Bluetooth landscape in a city in terms of patterns of Bluetooth presence and Bluetooth naming practices.

1 Introduction

Building ubiquitous computing systems in our cities requires new ways of thinking about the design and use of technologies and how they interweave with the built environment. We propose a holistic approach to designing the urban environment as an integrated system of architecture and ubiquitous technologies. Our goal is to understand the city as a system, encompassing both its physical and its digital forms and their relationships with people's behaviours in the city. Achieving this goal will require the development, application and refinement of methods of observing, recording, modelling and analysing the city – physically, digitally and socially.

Researchers in the fields of ubiquitous and mobile computing have tended to investigate interventions or experiences in specific urban situations, rather than studying systemic properties of the city as a whole. For example, the Mobile Bristol project [Reid et al. 2005] created situated experiences, such as an “interactive play” in Bristol’s Queen Square. The Equator Citywide project [Benford et al. 2003] developed games played in urban settings. The goal of those activities was to *overlay* an experience on city spaces, by giving mobile devices with GPS receivers to users, rather than by placing technology into the settings themselves – which is often not possible for practical or regulatory reasons. By contrast, several projects have *embedded* technological artefacts into urban situations. For example, Urban Atmospheres’ Jetsam project [Paulos & Jenkins 2005] augmented a city rubbish bin, so that it projects on to the pavement a representation of the activities in which it is involved, which are otherwise latent.

All of these projects are city probes: a technological perturbation of the city experience made with the aim of understanding more about how to design experiences in the city. Other work has attempted to understand existing city behaviours, principally as a resource for designing new applications. Paulos and Goodman [2004] studied the phenomenon of familiar strangers – people we become accustomed to seeing in urban settings but do not communicate with – by asking subjects in Berkeley to record the people they recognised. This became the basis for tools designed, for example, to augment the user’s sense of social relationship to different parts of a city. Others have looked principally at physical behaviours in cities, which themselves are often rooted in social behaviours. Höflich [2005] studied the movements and body language of people in the Piazza Matteotti in Udine as they made mobile telephone calls, relating them to the architectural features of that square and the different types of engagement people have with their interlocutors versus their surroundings. He identified signature patterns and paths of movement, which the work of Mobile Bristol and Urban Atmospheres, cited above, also identified in their particular settings, reflecting a common interest in how technologies affect paths through space. But those patterns are informally described. Moreover, no attempt has been made to generalise them so that we could, for example, compare a class of behaviours in different settings. Mainwaring et al. [2005] studied “urbanites” in three major cities as somehow representative and thus a basis for comparing them. While this informal study gives a flavour of aspects of city life in different places, we aim to develop a basis for more systematic comparisons.

Space syntax [Hillier & Hanson, 1984; Hillier et al., 1993] provides us with a systemic approach to understanding and designing the city and a range of methods and modelling tools that have been extensively tried and tested in both analytical and design practice [Stonor, 1997]. Our previous space syntax research has revealed how, through its structuring of space, urban design plays a critical role in the construction of society and social behaviours [Hillier et al., 1987]. In this paper we report our ongoing efforts to develop, apply and refine methods for understanding ubiquitous systems as an integral facet of the city. In section 2, we provide a brief introduction to space syntax, describing some of its main features and methods. In section 3, we describe how we are extending some of the methods used by space syntax to take account of the digital, as well as physical, form of the city. We illustrate this with examples from a study in the city of Bath, UK, in which we are developing and apply-

ing “digitally augmented” versions of two key space syntax observation methods: gatecounts and static snapshots. In section 4, we present some initial findings from our field trials with these methods, describing patterns of Bluetooth presence and Bluetooth naming practices. We conclude by summarising our ongoing work on developing methods as part of a systemic approach and toolkit for analysing and understanding ubiquitous computing systems as integral facets of the urban environment.

2 Space syntax

Space syntax analyses cities as systems of space created by the physical artefacts of architecture and urban design in order to understand how the spatial structure of the city is related to aspects of its social function. Space syntax is distinct from other forms of spatial analysis in that it characterises spatial elements (rooms, street segments, squares etc) first and foremost in terms of their relations (e.g. their graph distance) to other elements in the system. The resulting patterns of values for elements have been shown to be related with many functional phenomena, such as pedestrian and vehicular movement and land use patterns. These findings suggest how it is that the city considered as a pattern of connected spaces takes on social meaning by constructing patterns of copresence between people in space. In the Cityware project this allows us to integrate a range of social, cultural and economic factors within a single study methodology.

A fundamental concept of space syntax is that a city can be represented as a graph of nodes and links. The graph is constructed from a map of the city by first making an “axial map” of the streets. In an axial map, the longest lines passing down streets are considered as nodes and their intersections as links in the graph. This graph can then be analysed in terms of its properties such as the depth between the nodes. This is characterised in space syntax as the level of integration of a node, i.e. the deepness or shallowness of a node in relation to the other nodes in the graph. Shallow streets are essentially fewer changes of direction from the other streets of a city, while deep streets are relatively isolated from other streets. Space syntax has found a consistent correlation between the shallowness of streets (in terms of integration) and pedestrian flows in the city.

The ability to interrogate the spatial structure of a city plan, and to investigate what factors lead to the presence or absence of people on the street is one of the central contributions of space syntax methodology. Designers can manipulate a city’s map (e.g. when considering the construction of a new bridge), and can make predictions about the resulting impact on people’s movements in the city. The first step in studies of this sort is to gain an understanding of how people use and move through urban space. Empirical observations record the fact of people’s presence in the environment essentially *without* taking account of their intentions. If we ask an individual in the city about her pattern of movement, she is likely to respond in terms of purposes of journeys. However, the collective activity of a whole population gives rise to a pattern of use and movement that in itself forms an important aspect of the social context of the environment [Hillier & Hanson, 1984]. Through empirical observation, we can

retrieve something of the objective properties of the built environment through its influence on people's behaviour at this aggregate level. Space syntax draws on an extensive range of empirical observation methods, many of them also used more widely in urban studies and ethnography. Here we note two of the methods most commonly used by space syntax: gatecounts and static snapshots.

2.1 Gatecounts

Gatecounts are used to establish the flows of people at sampled locations within the city over the course of a day. A gate is a conceptual line across a street, and gatecounts entail counting the number of people crossing that line. The observer stands on the street and counts the number of people crossing the gate in either direction. To demonstrate how gatecounts are traditionally performed, here we describe a set of observations in Bath. We established 96 gates throughout the city, and counted the number of people crossing them. Our observers took 5-minute samples from each gate in 5 cycles throughout the day, from 8:30am to 4:00pm over 2 days.

Observed flows of people ranged from high flows of 2750-4000 people per hour to low flows of 250 people per hour or less. The correlation between the predicted and actual flows of people is shown in Figure 1. In this figure we see a scattergram correlating the observed flows (vertical axis) with the integration value of the street (horizontal axis). The degree of correlation in Bath is low in comparison to that found in other cities. This indicates that patterns of movement are likely to be heavily influenced by a range of other factors – location of tourist “attractors” for example [Fatah gen Schieck et al., 2005; 2006].




Figure 1. Correlation between pedestrian movement (y-axis) and integration (x-axis) for 40 streets in the centre of Bath ($R^2 = 0.47$).

2.2 Static snapshots

A second important method used by space syntax is *static snapshots*. Here, open spaces of the city are considered in detail. These spaces may be external, such as a plaza, or internal, such as a café. The method can be used for recording both stationary and moving activities, and is useful when a direct comparison is being made between the two types of space use. This method makes apparent the relationships between different types of space use in an urban area. For each open space under consideration, the observers record the movements in and out of the space, as well as the type of activity taking place in the space. This gives us an understanding of how people appropriate and make use of a particular space, and how these patterns of use bring people into contact with each other. For example, we may observe that a seating area in a park is actually not used for seating but for playing by children. A common observation is the use of certain spaces by people making calls on their mobile phones or using their laptop computers, and the way that these people then locate themselves with respect to the surrounding urban fabric and other people.

The space syntax methods described in this section provide a well established tool-set for understanding the architectural landscape of a city, for considering design changes to be made (e.g. adding a new bridge, blocking a street), and for observing and evaluating the results of design decisions. In the next section we describe our first steps in incorporating space syntax methods into the understanding, design and evaluation of ubiquitous computing systems as integral facets of the urban landscape.

3 Studying architectural spaces and interaction spaces in the city

In previous work, we have compared architecturally defined spaces with the “interaction spaces” that are created by artefacts or devices such as computer displays [Kostakos, 2005]. These interaction spaces define the boundaries within which the artefact is usable [O’Neill et al., 1999]. An essential feature of interaction spaces is that they are defined both by the characteristics of the device and by the architectural space in which they are situated. For example, within a public architectural space, a large display can create a public interaction space. In addition to visual interaction spaces, interaction spaces may also be auditory. For example, open-air broadcast of audio content via loudspeakers can create a public interaction space while a headset can create a private interaction space that includes just the person wearing the headset [O’Neill et al., 2004].

Interaction spaces may also be wireless. For example, the wireless interaction spaces generated by 802.11, GPRS or 3G access points define spaces within which certain devices (such as phones and PDAs) and services (email, browsing etc) are usable. We differentiate between such fixed wireless interaction spaces and more mobile wireless interaction spaces. The former tend to be static in relation to the physical environment and their location within the city (although, for example, 3G cell coverage can expand and contract quite dramatically with network load). On the other hand, the wireless interaction spaces created by technologies such as Bluetooth are often mobile, and move around as users carry their devices through the city. As

they move, they may come into contact with various other features of the digital landscape: services beaming out of an interactive poster, Bluetooth phones belonging to friends, colleagues and strangers, as well as various Bluetooth devices such as headsets, keyboards and mice.

Together, the concepts of architectural space and interaction space aid us in mapping from urban location to the technological artefacts that are available to us and the forms of interaction we wish to support. A challenge we face is recording, representing and understanding the patterns of presence and use of the diverse forms of interaction spaces that are emerging in our cities through the use of ubiquitous technologies. In this section we report our first steps in adapting and extending space syntax methods to meet this challenge, extending space syntax's consideration of the architectural spaces created by the built environment to include the wireless interaction spaces created by Bluetooth devices.

Bluetooth technology has a characteristic that renders it appropriate for study by methods derived from those of space syntax. The vast majority of Bluetooth interaction spaces are created by small, personal devices such as mobile phones. Thus, in contrast to the interaction spaces created by typically static WiFi access points, the wireless interaction spaces created by Bluetooth devices map very closely to the movements of people around the city, which in turn are a primary concern of space syntax. In our work so far, we have extended both the gatecount and the static snapshot methods to include the observation and recording of Bluetooth interaction spaces and their relationship with people's movements in the city.

3.1 Extending the gatecount method to include Bluetooth interaction spaces

Our basic extension to the gatecount method was to record the movement of discoverable Bluetooth devices past a gate. Even this simple extension allows us to correlate pedestrian movements with Bluetooth device movements, providing baseline data about the penetration of Bluetooth into city life. Beyond simply counting the appearance of Bluetooth devices, we have been using this method to uncover interesting data on patterns of presence of Bluetooth devices, and Bluetooth device names.

We face several technical challenges in developing a Bluetooth scanning method that will effectively augment the standard gatecount method. The main difficulty is that a space syntax gate is a distinct (conceptual) line across the street through which people are counted, whereas Bluetooth devices are counted when the wireless interaction space of the device intersects the wireless interaction space of our scanner. Even in a narrow street with no junctions, this intersection is not entirely predictable due to the variability in device characteristics and environmental influences on Bluetooth propagation. In our urban scans, we have the added complication of picking up devices moving along adjacent streets and in buildings.

A directional Bluetooth antenna pointing directly across the street from our scanner would seem to suggest itself as a means of creating a Bluetooth gate. However, the Bluetooth technology itself mitigates against this. Bluetooth discovery takes up to several seconds. Even with multiple-dongle scanning, described below, this would mean failing to record devices that moved through a narrow linear gate. Hence, we

adopted the approach of recording with standard Bluetooth devices and performing post-analysis to identify the patterns in the data. We also used Bluetooth scanners with a range of only 10 metres to mitigate the effect of discovering distant devices.

3.1.1 Pilot Bluetooth gatecounts

We ran a series of pilot Bluetooth gatecounts in various locations around the city, involving one of the authors standing on the side of a street with a notebook computer that performed Bluetooth scanning. Drawing on the results of the gatecounts reported in Section 2.1, we selected locations with low, medium and high pedestrian flows. Conventional space syntax gatecounts involve an observer standing at each gate for a relatively brief period (typically as little as 5 minutes) and iterating rapidly around multiple gates, repeating each gate several times. We could not predict in advance how many people would be carrying discoverable Bluetooth devices but expected it to be nowhere near every passer-by. Hence, one concern was that this conventional pattern of observation might under-record passing Bluetooth devices. We therefore varied the time period for these gatecounts in order to assess what would be a suitable period. In Table 1 we show the confidence intervals obtained by varying the period of scanning. From these trials we identified 30 minutes as a suitable period for Bluetooth gatecounts at both high and medium traffic locations, as this is enough to generate small errors ($\pm 3\%$). For low traffic locations a period of roughly 2 hours would achieve such small margins.

	Pedestrians	Bluetooth	%	Error
Location 1 (low traffic)				
10 Minutes	21	4	19%	$\pm 16.8\%$
20 Minutes	57	6	10.5%	$\pm 8\%$
30 Minutes	77	13	16.9%	$\pm 8.4\%$
Location 2 (medium traffic)				
10 Minutes	103	13	12.6%	$\pm 6.4\%$
20 Minutes	210	20	9.5%	$\pm 3.9\%$
30 Minutes	331	31	9.3%	$\pm 3.1\%$
Location 3 (high traffic)				
5 Minutes	384	39	10.2%	$\pm 3\%$
10 Minutes	746	69	9.2%	$\pm 2.1\%$

Table 1. Sample confidence intervals derived from varying the period of our mobile Bluetooth gatecounts ($p=0.05$) for locations with low, medium and high pedestrian flows.

We learned several other lessons from these pilot trials and refined our techniques accordingly. First, a number of issues relating to the accuracy of Bluetooth scanning became apparent. We used a single Bluetooth dongle, which meant that if many Bluetooth devices passed the gate simultaneously, less information (or in some cases no information) was recorded. We found that if 20 people carrying discoverable Bluetooth devices crossed the gate simultaneously at walking pace, then approximately 60% of these devices would be recorded. This is due to the sequential nature of our

scanning dongle's communication with the passing devices and the short time that they are in range. We subsequently mitigated this problem by using 3 dongles simultaneously. Because of the random frequency hopping employed by the Bluetooth discovery mechanism, using more dongles increases the chance of discovering a device as it moves past our scanner. 3 was the maximum number of dongles that our notebook computers could power for long enough.

We were also able to increase the discovery rate of our mobile Bluetooth scanning by reducing the amount of information we recorded for each device. By recording only the unique Bluetooth addresses of the discovered devices, we reduced the time our dongles spent communicating with each discovered device, allowing our dongles to contact more devices in a given period. The trade-off is that we lose much of the richness of the recorded Bluetooth data, such as Bluetooth names and service identifiers, reducing the scan's value as an augmentation to the standard gatecount method.

3.1.2 Bluetooth-augmented space syntax gatecounts

In the next stage of our study we carried out 10 gatecounts throughout the city of Bath. Once again, the locations were selected based on the results of the gatecounts reported in Section 2.1 to cover low, medium and high pedestrian flows, but also to cover various types of spaces from open spaces to long narrow streets with no nearby junctions. These gatecounts closely resembled the conventional gatecount method but involved a pair of observers working together at each gate. One observer performed the manual pedestrian count while the other performed the Bluetooth count using our mobile scanner. The observers iterated around the gates throughout the city, recording the flow of people and Bluetooth activity at each gate over the course of two days. We applied the lessons learned from our pilot trials. Here we used 3 dongles to perform Bluetooth scanning, we recorded for 30 minutes at each location, and we recorded only the unique Bluetooth addresses of the discovered devices.

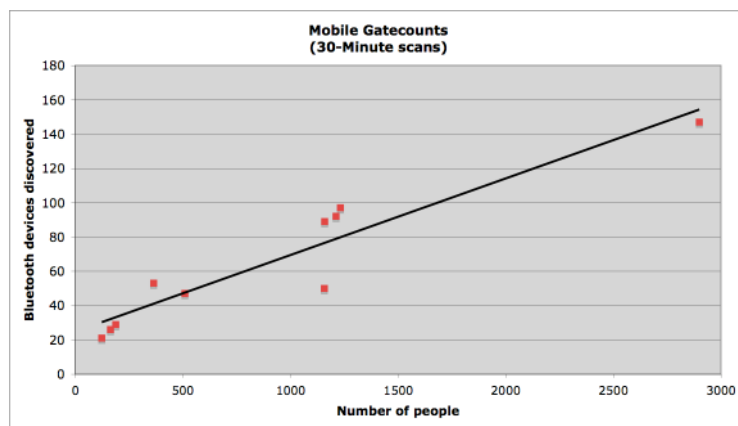


Figure 2. Correlation between number of people and number of discoverable Bluetooth devices across different locations in Bath ($R^2 = 0.8855$).

Figure 2 shows a high correlation between discovered Bluetooth devices and observed pedestrians. Overall, our scans suggested that 7% of observed pedestrians had discoverable Bluetooth devices.

From our Bluetooth-augmented gatecounts we learned further lessons for the refinement of our methods. At some of our sites, adjacent offices had Bluetooth-enabled computers, and these devices generated high Bluetooth activity. In this study, we were not interested in total Bluetooth activity, but rather the presence of Bluetooth devices. To measure this, we counted the number of distinct Bluetooth addresses in the recorded data as opposed to the raw Bluetooth activity. We also applied a filter to differentiate between persistent devices and those devices that appeared to be transient, reflecting the distinction between static and mobile interaction spaces. From our gatecount data and experimentation with known static and transient devices, we established empirically that a transient device typically appears for up to 90 seconds while it crosses a gate.

3.1.3 Fixed long-term gatecounts

In addition to the Bluetooth-augmented gatecounts described above, we ran 2 fixed, long-term Bluetooth gatecounts: 1 at a site on the University of Bath campus (Figure 3), the other on a street in the centre of Bath. At these locations we installed scanners that continuously search for Bluetooth devices that are set to be discoverable. For each device we attempted to record the following data: unique Bluetooth address, Bluetooth name, date, time, class of device (e.g. phone, laptop), and the services offered by the device (e.g. OBEX push, modem, fax, etc).




Figure 3. A Bluetooth gate with a Bluetooth scanner placed on the inside of a window.

Having the scanners continuously recording over a long period (at the time of writing, 3 months) avoids the problem that the conventional pattern of short observation periods at each gate might under-record passing Bluetooth devices. It also provides

very clear data on the cyclical nature of passing Bluetooth traffic. This illustrates the new opportunities offered by our combined methods, providing continuity of data that is effectively impossible with human observers and can be complemented by manually collected data at desired intervals at the same gate.

3.2 Extending static snapshots to include Bluetooth interaction spaces

We also extended the static snapshot method with Bluetooth scanning, drawing on the lessons learned in developing and refining our augmented gatecount method. To trial our Bluetooth-augmented static snapshot method, we installed a long-term scanner in a city centre pub, and another in a café. These scanners were technologically identical to the ones we used for our fixed long-term gatecounts. The difference was in the siting of the scanner. Both our fixed university gate and fixed city street gate employed scanners facing directly out a window. Although some Bluetooth activity was picked up from inside the building, devices crossing our gate generated most of the transient activity. In the pub we placed our scanner near the bar area, while in the café our scanner was placed near the seating area. In each case, we recorded Bluetooth activity generated mainly by devices in these areas.

In addition to the constant Bluetooth scanning, we carried out two 30-minute observation sessions in each of the two static snapshot locations. The purpose of these observations was to verify and correlate the recorded Bluetooth activity with the human activity in the area. Our human observers recorded people's positions, behaviours and movements through space, as well as the precise time of these activities. We subsequently compared these observations with the data recorded by our Bluetooth scanners, generating aggregate data, reported in Section 4.1, which was unavailable using conventional space syntax methods.

4 Field trials of our augmented methods

In this section we present our field trials using our newly developed methods. Although we are still refining our methods, these findings offer concrete examples of the results our methods can provide. We present our findings in terms of patterns of presence of Bluetooth activity in the city and an initial analysis of the Bluetooth names we recorded.

4.1 Patterns of Bluetooth presence

Simply looking at the raw Bluetooth activity can be misleading. From Figure 4 it appears that the university gate was much busier than the city centre gate. However, filtering out multiple records per device and persistent devices (indicating nearby static Bluetooth devices), we can identify the transient Bluetooth devices, shown in Figure 5.




Figure 4: Raw Bluetooth activity data for our campus and city centre fixed gates.

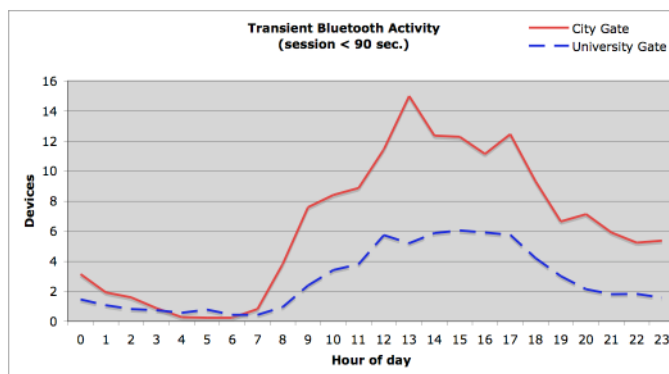


Figure 5: Transient Bluetooth devices for our campus and city centre fixed gates.

So, the city centre gate peaks at 15 unique transient devices per hour, while the campus gate peaks at 6 devices per hour. Intriguingly, the 2 graphs have a very similar profile despite recording Bluetooth traffic at very different sites. Another feature of note is the temporal pattern, with Bluetooth traffic at both sites peaking in mid-afternoon. The peak of 15 devices per hour for the city centre gate refers to the period 1pm to 2pm. This corresponds to 7.8% of the pedestrian traffic for that location, as recorded by our gatecount study reported in Section 2.1. This closely matches the 7% found in our Bluetooth-augmented gatecounts reported in Section 3.1.2, giving us a useful confidence check on our results.

The raw data recorded by our scanners are in the form of timestamped events. A single discovered device typically generates multiple events while it is within range of our scanners. In analysing the data, we developed a timeline visualisation, illustrated in Figure 6 where the top part of the figure is a magnified version of the area marked by the rectangle in the bottom part of the figure.

As a new device is discovered, we assign it to a new timeline. It then generates timestamps, indicated by circles on the timelines. In Figure 6, we see that device 12 was discovered at 15.5 minutes from the start of the observation, was visible for about 40 seconds, and was never seen again. Device 11, on the other hand was discovered at 14.5 minutes, and then again at 21.5 minutes. Devices 14 and 15 were discovered only once, which indicates that they may be at the periphery of our scanning range, and probably did not pass through our gate. As noted in Section 3.1.2, devices which pass through a gate are typically seen for up to 90 seconds. We use this 90-second period as a threshold for identifying highly transient devices.




Figure 6: A timeline visualization of our Bluetooth gatecounts. Each device is given its own timeline (dashed lines in top half) and each discovery event is plotted as a circle on the timeline.

Our timeline visualisation creates the cumulative effect of a diagonal line from bottom left to top right. Any activity recorded below this main diagonal is attributable to persistent devices. In Figure 7 we show data from 3 Bluetooth gatecounts which took place at different locations and reflect contrasting patterns of Bluetooth presence.

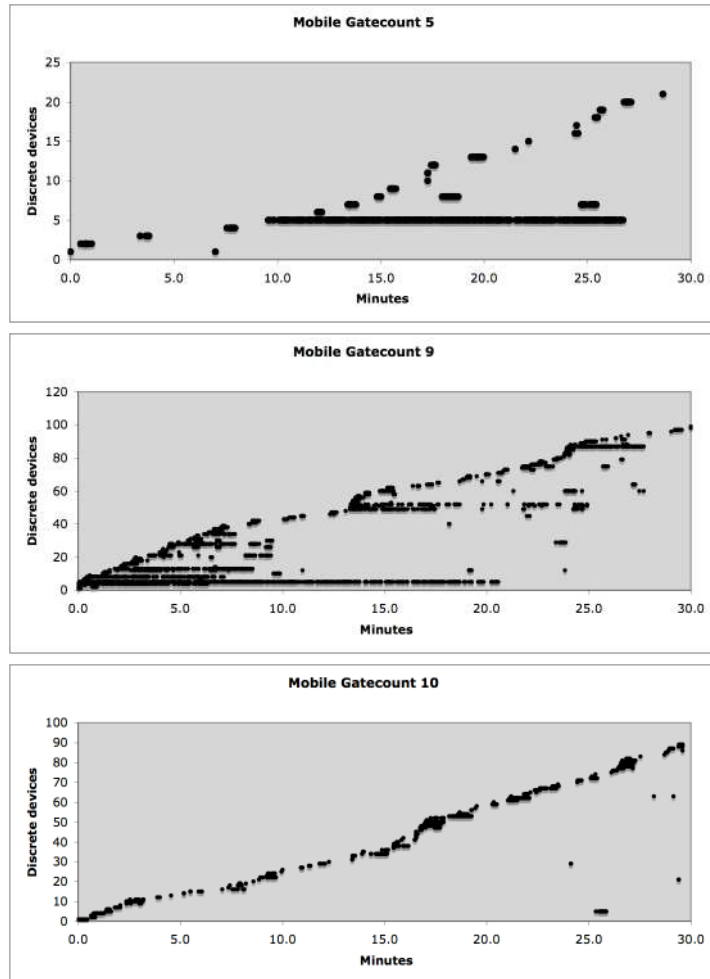


Figure 7: Visualizing Bluetooth gatecount records. Activity below the main diagonal indicates persistent devices.

Gatecount 9 has a relatively high level of persistent devices, while gatecounts 5 and 10 recorded mostly transient devices. In gatecount 9, we can identify bursts of Bluetooth activity recorded at 0, 5, 13 and 23 minutes. Finally, in gatecount 5, we observe the continuous presence of a device for approximately 17 minutes.

We would expect major differences in the patterns of presence between our gatecount and static snapshot data, since one records primarily Bluetooth traffic passing along a street while the other records primarily the Bluetooth devices of people in a café or pub. Figure 8 shows the duration of Bluetooth sessions recorded at the university gate and the pub.

In Figure 8, the bars indicate frequency (measured on the left y-axis), while the line indicates cumulative percentage (measured on the right y-axis). The duration of the

sessions is shown on the x-axis. The difference in patterns of presence between the two locations is clear. The university gate data showed 85% of sessions under 90 seconds, a threshold indicating highly transient devices. In the pub only 35% of sessions were below 90 seconds. The data from the pub also shows much more diversity in session length, with sessions recorded of up to 10 hours – generated by the pub manager’s phone.

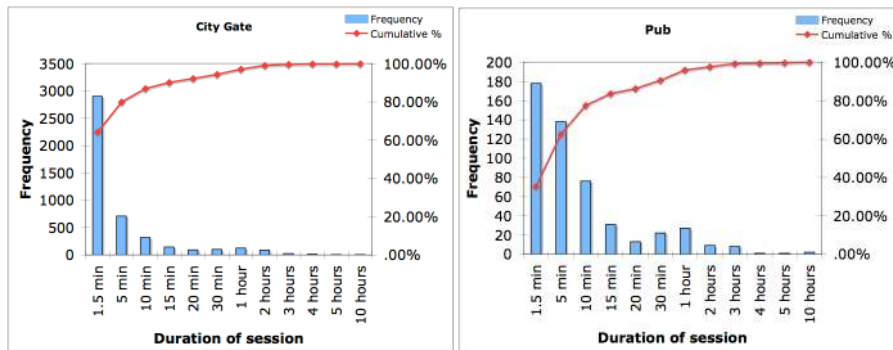


Figure 8. Duration of presence of Bluetooth devices.

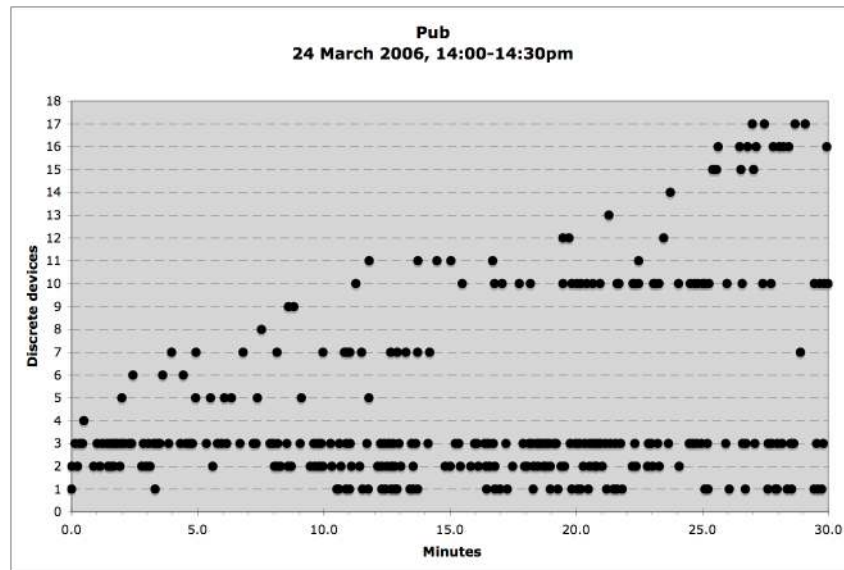


Figure 9. Bluetooth activity recorded during one of our observation sessions in a pub.

Our combination of conventional static snapshot recording by human observers and Bluetooth scanning enables further interpretation of our timeline visualizations. For example, Figure 9 presents the Bluetooth activity in the pub during one of our static snapshot observation sessions. Correlating these data with our observational

data allowed us to identify devices 2 and 3 as belonging to members of staff. Device 7 disappeared when we observed two men leave the area after having lunch. Device 1 was our observer, who was carrying out Bluetooth discoveries with his phone. (During this operation, the Bluetooth device cannot itself be discovered.) When three women entered the area, devices 15, 16 and 17 appeared. These devices had feminine names: “Jen”, “Cass” and “Han”. Finally, we were able to attribute devices 9 and 10 to the same person. This person was briefly present in our observation area, and the two devices were named “Sam K” and “Sam Karoot”. In addition to patterns of Bluetooth presence, our methods are beginning to uncover patterns in user’s Bluetooth naming practices. These are described in the next section.

4.2 Bluetooth names

We collected Bluetooth name data at three of our scanning sites: the university campus, a street in the city centre, and inside a pub in the city centre. Discovering a device’s Bluetooth name takes significantly longer than simply discovering its Bluetooth address, which meant that the names of many recorded transient devices went unrecorded. Nonetheless, we collected 1703 Bluetooth names: 771 from the city centre street, 307 from the pub and 625 from the campus. The great majority of these were names of mobile phones.

Devices are usually configured by default with names that state the make and model, such as “Nokia 6680” or “TomTom GO 300” (a satellite navigation unit). It is straightforward to distinguish variations from those defaults. Our logs revealed that 58% of discoverable devices had user-defined names in the city centre street; for the campus, the figure was 76%; and for the pub it was 88%. Interpreting these figures demands contextual data about user practices and intentions if we are to draw meaningful conclusions. But the figures are at least indicative of a thriving culture of Bluetooth naming. With factory default names, we simply cannot tell whether the user is unaware of the potential for setting – and making discoverable – a Bluetooth name, or if she is aware of Bluetooth and its properties and has chosen to retain the default name. However, from the user-defined names we may infer with some confidence that the user has an awareness of Bluetooth and its properties.

The choices of names themselves may suggest different forms of awareness of Bluetooth’s properties, and variations in the degree and type of effect that users wanted to have upon other Bluetooth users when they were discovered. Some names show an acute awareness of Bluetooth and the possibility of being discovered: “Clear off!!!”, “Pick me pick me”, “Send me porn”, “U Found Meee...”.

Being “Bluetooth aware” serves as a precondition for choosing names that may be characterised on a *self-others* spectrum, running from one extreme of simple presentation of self to another extreme of seeking an effect on other people. Many people chose simply to identify themselves. Some, such as “Nokia 6280 Wayne” did not remove the default name altogether, possibly because the type of phone had some significance. Others removed the default name entirely and provided a name in either a short (“Annie”) or long (“John K. Taylor”) form. Still others, such as “Snagglepuss” and “Crown Jools”, gave themselves aliases. Here, we begin to see a

movement along the spectrum from simply stating one's name to projecting an identity. This becomes even more apparent with others who identified themselves not so much by name as by what (or whom) they wished to associate themselves with. "Beer boy" is on the one hand an alias, but also associates this individual with drinking beer. "M.C.F.C OK!" refers to a football club; and "Pezza's girl" even associates herself with a (boy)friend. As we move towards the other extreme of this spectrum, the term *name* is less obviously apt for what people entered on their phones. Many provided texts that one might find in public places, such as graffiti and T-shirts: "4 a gay time call 077...", "\$LiK JiM", "LesbiansAre4Girls", "Ima kettle".

Many accounts are possible of the foregoing examples. Our account here represents first thoughts on dimensions of variation in the data we are seeing. We are inferring patterns in what the names themselves suggest, rather than what the users intended. The latter requires further research and we are engaged in investigations to understand naming as a set of practices in particular contexts. A fuller understanding of the culture and practice of Bluetooth naming requires not just scanning passing devices but gathering data from users about their Bluetooth naming practices.

What these names substantiate, however, is that, for the majority of users who are Bluetooth aware, Bluetooth on mobile phones gives rise to a de facto rather than merely potential interaction space. There is little point in altering the device's name from the default unless there is an intention for either functional or social interaction (or both) through this new electronic medium. Through her choice of name, the user defines the "feel" of that interaction space. The fusion of such wireless interaction spaces with the spaces defined by the features of our built environment has the potential strongly to influence people's relationships with those urban spaces. For example, consider the different potential effects on your behaviour if you became aware that the café in which you had always felt most comfortable had amongst its customers a preponderance of discoverable Bluetooth names that you found either upsetting or reassuring.

5 Conclusions and ongoing work

Whilst we have already generated a large corpus of interesting data, the primary purpose of the work reported in this paper was to develop and refine the methods themselves, and we have made considerable progress in this respect. Our extended methods do not simply consist in our Bluetooth scanning. Rather, they consist in our Bluetooth scanning *in combination with* the conventional observational techniques.

Our data on the use of Bluetooth names suggests the appropriation of the technology to project identity and to engage with others whilst enjoying a cloak of relative anonymity. Other data shows people visiting a café for extended periods not to eat or drink, or even to engage socially with others in the café, but to use the WiFi service. Each of these examples was thrown up by our field studies but drew on only one side of our extended methods, conventional or electronic. The Bluetooth names were recorded by our electronic scanning and did not rely on the accompanying human observations, while the observations of non-eating, non-drinking WiFi users were made by standard static snapshot techniques and did not rely on the accompanying

electronic scanning. These findings confirm the utility of our extended methods in capturing a wide range of data that neither approach alone would encompass.

In addition, our extended methods capture data that require the *combination* of both electronic and human observation. For example, a characteristic of conventional gatecounts and static snapshots that we are not able to reproduce with Bluetooth scans alone is the ability to classify pedestrian flow depending on people's characteristics, such as men, women, locals, tourists, children or adults. On the other hand, our Bluetooth scanning allows us to record and classify data that is simply not accessible to conventional gatecount and static snapshot methods, including characteristics of both the devices and the people carrying them. Thus, in addition to recording the presence and flows of Bluetooth interaction spaces, our extended methods gather data about the devices generating them, such as whether they are mobile phones or notebook computers, while our recording of Bluetooth names provides a rich dataset that allows us to classify people in terms other than the characteristics recorded by conventional observational methods.

We will continue to refine our methods, and indeed develop other methods, for both the capture and analysis of data. For example, we have not yet overcome the problem of discovering Bluetooth devices on adjacent streets or in adjacent buildings. We can in principle mitigate the adjacent streets problem by careful choice of scanner location. The ideal site for a gatecount scanner is on a long, narrow street with no nearby junctions, where we can be more confident that discovered devices have actually crossed our conceptual gate line. However, it is unlikely in general that sites where we can place fixed long-term scanners will coincide neatly with this ideal spatial morphology. We have much more freedom in choosing the location of our mobile scanners but this will not always help because our choice of site may be constrained by, for example, our research questions. Thus, we may be interested in studying Bluetooth flow in a particular urban space regardless of how well that space lends itself to our scanning.

The adjacent buildings problem is mitigated by our developing data analysis methods. While our analysis differentiates between static and transient Bluetooth devices, rather than between activity on the street and in buildings, we can usually infer location in a building for highly persistent devices. We have found empirically that fixed devices in buildings account for the majority of highly persistent devices recorded by gatecounts. While our analysis is less useful at distinguishing between less persistent devices in buildings and passing devices on the street, mobile devices within buildings are often more persistent than the transient devices of passers-by in the street. In this way, we have demonstrated that Bluetooth scanning enables short-term capture of data that lends itself well to the study of wireless interaction spaces in the city. We hope to extend this approach to inspect changes in Bluetooth activities over longer time scales and thereby to monitor the patterned effects of our future technological interventions.

The field of ubiquitous computing in general lacks concrete methods grounded in the urban context. The methods we are exploring and extending, conventionally applied in analysing and understanding the traditional architectural features of the urban environment, can help us to analyse and understand ubiquitous computing features as integral aspects of that environment. Our ongoing research continues to develop and refine these and other methods, while gathering further interesting data as we progress

towards our goal of understanding the city as a system encompassing both the built environment and ubiquitous technologies.

Acknowledgements

This research is funded by the UK Engineering and Physical Sciences Research Council grant EP/C547683/1 (Cityware: urban design and pervasive systems).

References

- Benford, S., Anastasi, R., Flintham, M., Drozd, A., Crabtree, A., Greenhalgh, C., Tanda-Vanitj, N., Adams, M. and Row-Farr, J. (2003). Coping with uncertainty in a location-based game. *IEEE Pervasive Computing*, 2(3): 34-41.
- Fatah gen Schieck, A., Lopez de Vallejo, I. and Penn, A., (2005). Urban Space and Pervasive Systems. In *proceedings UbiComp 2005*, Tokyo, Japan, (poster).
- Fatah gen. Schieck, Penn, A., Kostakos, V., O'Neill, E., Kindberg, T., Stanton Fraser, D., and Jones, T. (2006). Design Tools for Pervasive Computing in Urban Environment. In *proceedings 8th International Conference on Design & Decision Support Systems in Architecture and Urban Planning*, Springer, in press.
- Hillier B, Penn A, Hanson J, Grajewski T & Xu J. (1993). Natural movement; or, configuration and attraction in urban space use, *Environment and Planning B*, 1993.
- Hillier, B., & Hanson, J. (1984). *The social logic of space*. Cambridge: Cambridge University Press.
- Hillier, B., Burdett, R., Peponis J., and Penn, A. (1987), Creating Life: Or Does Architecture Determine Anything? In *Architecture and Behaviour*, 3(3): 233-250.
- Höfllich, J. R (2005). A certain sense of place. In *A sense of place – the global and the local in mobile communication*. K. Nyíri (ed.), Passagen Verlag, pp. 159-168.
- Kostakos, V. (2005). *A design framework for pervasive computing systems*. PhD Thesis, University of Bath, UK. Technical Report CSBU2005-02, ISSN 1740-9497.
- Mainwaring, S.D., Anderson, K., and Chang, M.F. (2005). Living for the global city: mobile kits, urban interfaces, and ubicomp. In *proceedings UbiComp 2005*, Tokyo, pp. 269-286.
- O'Neill, E., Johnson, P., & Johnson, H. (1999). Representations and User-Developer Interaction in Cooperative Analysis and Design. *Human Computer Interaction*, 14(1/2): 43-92.
- O'Neill, E., Woodgate, D., & Kostakos, V. (2004). Easing the Wait in the Emergency Room: Building a Theory of Public Information Systems. In *proceedings Designing Interactive Systems (DIS 2004)*, Cambridge, Massachusetts, pp. 17-25.
- Paulos, E., and Goodman, E. (2004). The familiar stranger: anxiety, comfort, and play in public places. In *proceedings CHI 2004*, ACM, pp. 24-29.
- Paulos, E., and Jenkins, T. (2005). Urban probes: encountering our emerging urban atmospheres. In *proceedings CHI 2005*, ACM, pp. 341-350.
- Reid, J., Hull, R., Cater, K., and Fleuriot, C. (2005). Magic moments in situated mediascapes. In *proceedings ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE 2005)*, Spain, ACM.
- Stonor, T. (1997). Making space: the use of space syntax in design. In *proceedings First International Symposium on Space Syntax*, 16-18 April 1997, UCL, London.