Social Gravity: A Virtual Elastic Tether for Casual, Privacy-Preserving Pedestrian Rendezvous

John Williamson[1]**, Simon Robinson**[2]**, Craig Stewart**[1]**,**
Roderick Murray-Smith[1]**, Matt Jones**[2]**, Stephen Brewster**[1]

[1] Department of Computing Science
University of Glasgow, G12 8QQ, UK
{ jhw, craig, rod, stephen } @dcs.gla.ac.uk

[2] Future Interaction Technology Lab
Swansea University, SA2 8PP, UK
{ s.n.w.robinson, matt.jones } @swan.ac.uk

ABSTRACT

We describe a virtual "tether" for mobile devices that allows groups to have quick, simple and privacy-preserving meetups. Our design provides cues which allow dynamic coordination of rendezvous without revealing user's positions. Using accelerometers and magnetometers, combined with GPS positioning and non-visual feedback, users can probe and sense a dynamic virtual object representing the nearest meeting point. The *Social Gravity* system makes social bonds tangible in a virtual world which is geographically grounded, using haptic feedback to help users rendezvous. We show dynamic navigation using this physical model-based system to be efficient and robust in significant field trials, even in the presence of low-quality positioning. The use of simulators to build models of mobile geolocated systems for pre-validation purposes is discussed, and results compared with those from our trials. Our results show interesting behaviours in the social coordination task, which lead to guidelines for geosocial interaction design. The *Social Gravity* system proved to be very successful in allowing groups to rendezvous efficiently and simply and can be implemented using only commercially available hardware.

Author Keywords

Mobile, navigation, vibrotactile, GPS, geosocial interaction.

ACM Classification Keywords

H.5.2 User Interfaces: Haptic I/O H.5.2 Prototyping: H.5.1 Artificial, augmented and virtual realities.

General Terms

Human Factors

INTRODUCTION

In this paper, we describe a system which creates a dynamic virtual "tether" structure that draw groups of people together. This structure gives cues as to the "group centre"; users can be brought together without revealing their individual locations – only the location of the group as a whole.

We designed the system to satisfy the following constraints: It should allow pedestrians in a common geographical area to meet up without revealing their current positions, and with very simple, efficient interactions, using non-visual feedback and inertial sensing to give users virtual probes to explore the space around them. The locations of users engaged in rendezvous activity should not be shared directly between participants; they should have no direct perception of the position of other members of the group. The problem of reaching consensus on rendezvous locations (one which scales quadratically with the number of participants) will be avoided by simply guiding all interacting parties to the closest point at which they will meet. Users can combine the information about the group centre with their knowledge of their immediate environment (obstacles, faster routes, etc.). We envision the system as an additional *sense* which uses the Earth-centred nature of the sensing hardware to provide a common environment, and which can be used as a tool in organising meetups.

In summary, the proposed interaction does not share user positions; does not require the following of sequences of directions; avoids time-consuming social negotiation of meet up points; can be operated non-visually and with minimal attention; can cope with substantial positioning and sensing noise and is efficient enough to be practical, even in physically-constrained environments such as built up areas.

Navigation without navigating

The world is abundant with applications for mobile devices for navigating, whether it be for vehicles or for pedestrian way-finding. The broad availability of good quality positioning information from GPS and WiFi/cell tower triangulation has made such functionality near-standard in mobile phones, expanding rapidly into more basic devices. Some prior and existing services, such as Google Latitude or Nokia's FriendView, take advantage of this to share positional information between groups of friends. This use of geographical information for social co-ordination will become increasingly important as positioning technology becomes ubiquitous.

Most navigation systems are concerned with (a) showing the approximate position of a user, and then (b) navigating to

CHI 2010, April 10-15, 2010, Atlanta, Georgia, USA.

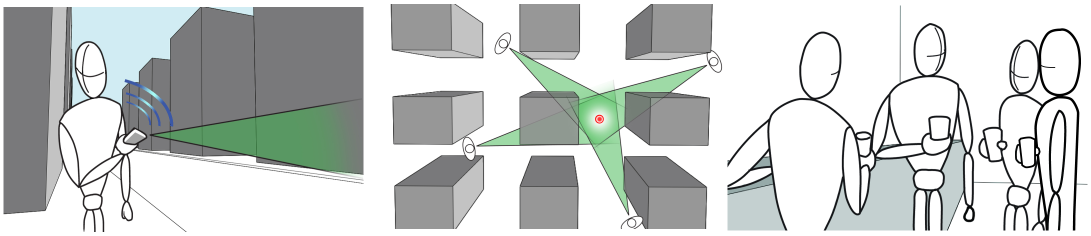

Figure 1. A group of friends have a quick & simple rendezvous after exploring the city. Stragglers can home in on the social centroid.

some desired destination via a series of turn-by-turn directions. Many everyday tasks are organised like this, and these automated systems greatly improve the ease of wayfinding in unfamiliar or difficult situations. But new forms of social interaction are afforded by the continuous use of positional information. Impromptu meetups based on serendipitous proximity are one obvious class of scenarios, and one which location sharing systems such as Latitude and FriendView facilitate. However, these services generally just display the locations of others on a map, and provide no functionality for the co-ordination of groups for joint social activities.

Design issues

These imagined possibilities can only be realised as successful systems if they can deal with the significant constraints affecting geolocated social interactions in real-world settings. The sources of constraints in building a geolocated social system can be divided into two categories: the technical constraints, which govern the precision, complexity and responsiveness of the fused physical-digital world; and the social constraints, which govern the behaviours which people are willing to engage in, and the information they are willing to share.

Social

Where there is sharing of location information, there is always a significant privacy issue. Few groups of individuals are likely to be close enough that they would be willing to always unconditionally share their respective locations, given the enormous disclosure of their private lives this would entail. Systems which can facilitate meaningful geolocated interaction without requiring the complete exposure of personal location have a wider potential of use than those which mindlessly share private details.

People interacting with a fused world are often operating in public, and the behaviours that they will be willing to engage in are therefore constrained by social norms. Clumsy, unsubtle or awkward interactions that might be acceptable in a private setting will not be appropriate. Interactions must be sufficiently socially-acceptable that they can be performed publicly without embarrassment. This includes both the design of socially-appropriate input – no wild movements or attention demanding intricate interfaces – and the creation of viable feedback mechanisms which do not require the use of senses that are otherwise occupied with everyday activity, or broadcast inappropriately into the surrounding social environment.

Technical

The accuracy to which GPS and other positioning technologies are able to fix a location is often relatively poor (between 10-200m depending on situation) for tasks such as pedestrian navigation, where small distances matter. Accuracy varies in different locations and at different times. Earth-relative sensors (such as accelerometers and magnetometers, which sense the effects of the Earth's gravitational and magnetic fields, respectively) are also subject to substantial variations in fidelity, given context-dependent disturbances. It is not obvious to all users that the heading read by an electronic compass will swing around under powerlines or that positioning quality will degrade in the presence of foliage, and many users will not understand why these deviations occur. There is lack of a shared understanding between how people observe the world and the frames of reference they use, and that of the system with which they are interacting.

Scenarios

One useful function that a system aware of its position and orientation relative to the Earth could perform is to tether together groups who are spread across an area. Obviously such groups can currently communicate via mobile phones or other wireless connectivity. But co-ordinating groups of people becomes more demanding as the size of the group grows, and is more difficult when the environment in which the group is attempting to co-ordinate is unfamiliar [4]. In this paper we explore the concept of a virtual tether which exerts "tension" on users to draw them together in a common meet up point. Once a rendezvous activity is initiated, the tether affords sensing of the group's "centre of mass". The members of the group can then navigate to this dynamically-optimised point. This simple function, if available in a robust, discreet and convenient form would be of value in many common situations, an example of which follows (illustrated in Figure 1).

Conference meetup scenario

Andy, Jen, Lisa, Steve and Tom are at a scientific conference in an unfamiliar city which they have just arrived in that morning. They plan to meet up for a drink following the presentations, but Steve and Andy want to have a chance to explore the city first and Lisa has a workshop session running late. At 6 o'clock, Tom decides to see if the others want to meet up, and sends them an invitation to join a rendezvous group. Steve, Andy and Jen accept, but Lisa is still held back in the workshop. Tom and the others in the rendezvous group then briefly scan around with their phones until they feel the distinctive homing vibration. Each starts walking in the direction of the vibration beacon, checking their direction every few minutes with a brief flick of the phone and panning for the goal. After a few detours around intervening buildings, the four meet up without any further communication, and head along to the nearest bar. Later, Lisa uses her phone to join the rendezvous group and homes in on the bar using the tactile feedback.

PREVIOUS NAVIGATION STUDIES

Previous work has evaluated the use of vibrotactile feedback to guide users in navigation tasks. One of the most common approaches to date has been to use a belt to provide waypoint cues. FeelSpace [14], for example, is an electronic vibrotactile compass, which displays the direction of north in real-time using an array of vibration transducers mounted on a belt. Van Erp *et al.* [24], applied this to navigation, creating a similar vibrotactile belt for use in situations where visual feedback is not appropriate. Studying several combinations of haptic pulses, an initial experiment found that distance-coded feedback was able to successfully guide users between waypoints while walking. Further investigation by van Erp [23] showed that well-placed tactors (small vibration motors) could be accurate as directional indicators.

One obvious application of non-visual navigation is for users with visual impairments. Johnson and Higgins [10], for example, studied such a device, creating a tactor belt to help blind users avoid objects in their paths. Two cameras attached to the belt provide a stereo view of the scene, and feedback is mapped to the distance of each object. Part of the motivation behind their belt-based approach was to avoid restricting the user from other activities. Our approach to this problem is to promote opportunistic navigation interactions allowing the user to make casual, infrequent requests for feedback when it is appropriate, rather than being guided constantly toward a target destination.

The use of a handheld device to provide haptic feedback for navigation is less common than belt-based approaches. Sokoler *et al.* [19] describe a low-resolution tactile approach where one of four pegs is raised to give a direction cue. Lin *et al.* [13] used structured tactons to provide rhythmic vibrations as navigation aids (turn left, right or stop). Their prototype achieved high levels of tacton recognition, finding that users were able to pay attention to their environment at the same time as using the system. However, a Wizard of Oz approach was used to guide users to the target location. Our system uses simpler, but realtime navigational feedback, with the user scanning their mobile device to discover the direction they should head in.

Our use of haptic feedback for privacy-preserving group navigation is particularly novel. Previous work has investigated user behaviours when rendezvousing, but these have used visual-based systems. Axup *et al.* [1] examined an early prototype which allowed users to send group text messages to co-ordinate a rendezvous. They found that while users were able to meet up, the method used was "somewhat unpopular" and had many usability problems. In addition, and particularly related to our design, they found that the visual attention required to operate the system forced users to alternate their attention between screen and environment to avoid walking into obstacles. Our design uses a more lightweight, low-attention interface that allows users to concentrate their attention on their surroundings rather than a mobile device.

Olofsson *et al.* [16] studied user needs for meeting during music festivals, and proposed a concept device to display the location of nearby friends overlaid on a map of the festival, based on their findings from a field study. Nicolai *et al.* [15] explored social contexts in location-aware systems, using mobile proximity-awareness of around 10m to give a group of users feedback when members joined or left. The system provided no navigational assistance; instead it gave users a brief overview of the people nearby, and classified them into familiarity groups. Our system focuses instead on groups of individuals that want to meet each other over larger areas, guiding them to a convenient meeting point without the need for visual feedback.

Other authors have looked at how people behave during rendezvous. Colbert [3, 4] extensively explored a diary study of users' behaviour while rendezvousing, looking at the effect of factors such as group size, time pressure and area familiarity. Larger group sizes were found to cause more stress to participants, but they were still were able to rendezvous successfully. Area familiarity also affected users' rendezvous behaviour: rendezvous in unfamiliar locations required more communication and caused problems attributed to the lack of local knowledge. This aspect is directly addressed by our system – with our design no group communication is necessary, and although local knowledge may shorten the time taken to meet up, it is not required.

Dearman *et al.* [6] conducted an exploratory Wizard of Oz field study to investigate mobile location-aware rendezvous behaviours between pairs of participants. Those using a visual location-aware handheld device chose a meeting location that was a middle point in the majority of cases, with only one pair choosing a landmark. Further work by Dearman *et al.* [7] investigated user requirements during rendezvous, finding that participants often maintained continual awareness of partner and meeting locations. Our system allows for this by providing meeting point awareness on demand, but does not provide partner locations in order to preserve their privacy during the meetup process.

SYSTEM DESIGN

The *Social Gravity* system is intended to illustrate a clear example of how shared, geographically-bound virtual environments can be created with low-latency multimodal feedback and inertial sensing. The system described is a mechanism for bringing together geographically distributed groups. It does not solve the tricky design issue of how users join and leave such social groups. Its mechanism, however, is tolerant of users dynamically entering and leaving rendezvous. The *Social Gravity* design was pre-tested with a realistic movement and sensing simulator to establish whether meaningful interaction could occur given the constraints of sensing with the available mobile hardware. This both indicated the feasibility of the system and permitted the interactive tuning of critical parameters in the interface design.

Privacy control

Information about an individual's location is sensitive. Some existing technologies allow partial location sharing, providing granularity controls to limit the detail at which others can observe movements. The *Social Gravity* system is specifically designed to sidestep the issue of sharing positional information, by only sharing summary attributes of a *group's* location and movement. Here we use the centroid as a way of tethering together groups; other functions of group position could provide other senses – average velocity, convergence rate, radius or local density. All of these can extend the users sense of the social presence around them, without revealing the specific positions of other interactors.

Control rather than way points

The conventional approach to navigation is to either present a visual map indicating current location and target, or to provide a series of turn-by-turn directions (where paths are constrained, as in car navigation). In this work we are creating a fused environment in which the combined sensor/feedback unit becomes a tool which responds to objects in the digital world which appear to be attached to physical places. The immediate goal is not to guide someone to a target but to create a convincing illusion of a shared object in a physical space. The *Social Gravity* system is an elementary example of the class of systems where users' senses are extended such that social relations can be directly perceived in a geographically grounded environment. This is quite distinct from the problem of static wayfinding to a target; people might appropriate such augmentation to use the *Social Gravity* system for other ends – for example simply to have an awareness of other family members around them.

Non-visual, gestural interaction and feedback

We are concerned with the construction of an artefact which extends the user's perceptions so that the social connections pervading the local geographical space are tangible. An essential aspect of this is that the movements of the user are immediately and directly tied to the feedback they receive.

The negative impact of relying solely on visual feedback in mobile navigation tasks is obvious; people need to use their eyes to look around when walking. There is little free visual capacity for watching onscreen indicators while walking in unfamiliar environments outdoors. Of the non-visual modalities, audio feedback can be very rich and has a natural spatial component, and has been successfully used for navigation systems (see [8, 11, 20, 25]). However, it either has the side effect of being broadcast to those in the immediate environs, or requires the use of headphones – and will negatively interfere with any music being listened to (though [20] describes using music modulation for navigation as a potential alternative). Vibration feedback is private and inherently spatialized, but current technology limits the variety of sensations that can be presented, and only a few parts of the body are practical for presenting haptic feedback. Vibration lends itself well to motion sensing interfaces; the illusion of physical objects which are responding to movement can be generated. The most common vibrotactile navigation approach in existing literature is using a vest and presenting vibrations to the torso (e.g. [24]). Less common is the use of vibrotactile feedback provided by a mobile device held in the hand, such as [13].

Simulator

Developing and running mobile experiments is costly and time-consuming, and is often difficult to observe and monitor adequately. This problem is compounded when multiple interacting participants are involved. The system under consideration in this paper has numerous parameters that can be adjusted and various external constraints that can vary. The angles at which feedback is produced can be altered, for example, or the GPS positional fix quality can vary as as the environment changes. Testing of the feasibility of the virtual tether – and then optimising the parameters of it – would have been prohibitively expensive had it been done solely via field trials. Using a custom-built Python agent simulator, a "safe" parameter set was refined, and was then tested in a single field trial.

Obviously a simulator cannot capture the subtle complexities of human behaviour, but the simple navigation task involved here can be reasonably modelled with a few assumptions. Importantly, the uncertainties in the system (for example, inaccuracy in bearing sensing, or limited GPS resolution) can be modelled, and the effects on navigation performance observed and quantified. Much of the inspiration for this approach is drawn from the work in evacuation modelling [9], and crowd modelling [2] where simple models of human behaviour have been shown to predict well how people will move under constraints. These systems have been sufficiently accurate that building design and event management decisions have been taken on the basis of simulations.

Agent model

The simulator models the problem as a set of agents with a simple behaviour system. The agents act according to the following rules: (1) walk in the current direction (with some random wander); (2) occasionally stop, scan for the centroid, and then head in that direction; (3) if an obstacle is reached, turn to face the smallest angle away from the centroid where the obstacle is not in the way; (4) stop when within the stopping radius of the centroid. While simple, this model provides reasonable behavioural characteristics. Modelling the uncertainty, particularly GPS noise, the accuracy of the bearing signal, the impact of obstructions and the effects of network delay are what make the simulator powerful. The model can be explored interactively to check if agent behaviour seems plausible. The simulation can also be run in a batch mode, where repeated runs are performed and overall statistics (such as time to converge, average scanning time, etc.) are reported.

This simulator allowed us to establish that, with a realistic model of sensor noise and environmental constraints, participants would be able to meet up efficiently even with a poor

GPS fix and with network delays exceeding 30s. Various angular widths of the centroid were evaluated, and a value of 60° was finally chosen as an optimal trade off between oversensitivity and efficiency of convergence. Figure 2 shows the feedback layout.

Predictions

Using the simulator, predictions of user performance in the field trials – discussed later – can be performed. The exact positions of the starting points of the participants in the field trial were set in the simulator. The simulator was run with both a simplified set of obstacles representing the area of the field trial, and an empty space. For the no obstacles case, the simulator predicted a mean convergence time of 12:06 minutes (std. dev. 2:19), and 17:45 minutes (std. dev. 5:37) for the obstacles case. In both cases the distribution of end times was skewed such that there was a long tail of long meetup times where agents got stuck or oscillated around. This simulation was run with 5 agents walking at 0.8ms^{-1} (± 0.2), Gaussian GPS noise (std. dev. 8.0m), Gaussian angular noise (std. dev. 8°) and a 60° feedback angle. These parameters reflect a realistic model of human walking behaviour and GPS/inertial sensing accuracy.

Simulator conclusions

It is not difficult to create simple models of human behaviour that can illuminate the design of mobile systems. Realistic modelling of the sensing and environmental constraints that will affect user behaviour (particularly location-fixing inexactitude) can be used to highlight areas which may require additional attention. The simulator used in this work demonstrated the feasibility of the meetup system, even under very difficult conditions. It also allowed the iterative refinement of the free design parameters in the system. The simulator made concrete predictions (with appropriate error bounds) about user performance in the real task.

IMPLEMENTATION

The prototype system was constructed using standard Nokia N95 phones, in conjunction with the SHAKE inertial sensor pack [26]. This Bluetooth sensor pack provides accelerometer and magnetometer readings and includes a pager-style vibration motor for vibrotactile feedback. The N95 has built-in GPS for positioning, and its 3G connectivity is used for communications.

For our prototype the SHAKE device was encased in a mobile phone shaped form factor. The N95 was attached to a lanyard worn around the neck and was connected to the sensor pack by Bluetooth. The phone simply provided a 3G connection to the remote *Social Gravity* server. While at the time of our implementation there were no widely available phones that could provide the same reliable sensing functionality as our apparatus, there are now several phones available with all of the necessary hardware built in (for example the Nokia N97 and the Apple iPhone 3GS).

Feedback

The feedback is designed to allow users to sense the centroid in such a way that the centroid seems grounded in physical space. In creating feedback for relaxed, intermittent pedestrian navigation it would be beneficial to operate without any visual attention, and to be implementable with only a very crude pager motor device. Several iterations of the feedback design were tested in a lab setting, including modulations according to angular distance to target. None of the complex schemes offered significant improvement over simple "on-target" vibration, which has the distinct advantage of being easy to understand and requiring only basic vibrotactile output. Note that no presentation of target range is included (unlike, for example, in [27], where range is sensed by targeting a distance via tilt). Here, we are looking at the *ability* of people to meet up, but distance might be important when users decide *whether* to meet up.

The range of angles at which feedback is produced is critical to the success of the system. Because there are a number of noise sources influencing the user's accuracy, including GPS uncertainty, sensor to feedback latency, network delay and heading sensor uncertainty, a very narrow target angle will be frustratingly hard to target [21]. A very wide target angle will be slow to converge and participants will spiral around on to the target. Tests with the simulator established that an "on-target" zone of 60° results in excellent performance, even if users behave pathologically and always follow the worst angle within the feedback zone. Even in perfect (no noise) conditions, the simulated time to acquire a 20m wide target at 500m in an empty space with a 60° target angle is only 6:50 minutes, compared to 5:40 minutes for a 5° target. Figure 2 illustrates the feedback setup. This wide target means that latency between sensing and feedback has low impact and the uncertainty in heading angle (from magnetic disturbances and accelerometer motion) has little effect.

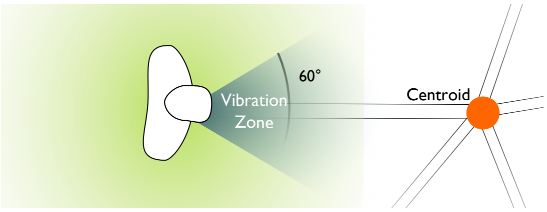

Figure 2. Vibration feedback is produced when the heading from the inertial sensing unit is within 60° of the heading to the centroid.

The vibration feedback was delivered via the pager motor within the SHAKE sensor pack, which has basic intensity control. A pulsing "on target" vibration was used, partly to mitigate sensor disturbance issues as described in the sensing section, and partly because it is less annoying than a continuous vibration. The pulsing nature also masks the discontinuities in the feedback zone transition. Such masking is intended to avoid overcontrolling where users attempt to stay on the boundary, where the *change* in feedback is greatest (behaviours seen in previous GPS navigation studies [21]).

Heading from sensing

The heading of the device can be computed from the direction of the Earth's magnetic field and the gravitational vector. On the implementation platform, these are measured by the accelerometer and magnetometer. The sensed values are converted into yaw (heading), pitch and roll values using a standard algorithm [22]. Although magnetic north can differ by several degrees from true north, these variations are too small to make a significant difference given the coarseness of the heading feedback. The heading computed this way is tilt-compensated and does not vary significantly for reasonable variations in pitch and roll ($< 15^\circ$ for pitch and roll angles $< 45^\circ$). Low-pass filters remove variation due to electrical noise and muscle tremor. However, the system is designed to be used with vibration feedback *produced in the sensor pack*, and the vibrations from the pager motor are sufficient to introduce very large acceleration disturbances. Repeating patterns of a short burst of vibration followed by a longer burst of silence allowed for more accurate sensing. The accelerometer value is "locked" to its current value just before the onset of a vibration interval and "unlocked" immediately afterward; only the magnetometer is used for updates during this period. This assumes the device will not wildly tilt or roll during the vibration pulse. In practice this results in much more reliable heading data. Feedback is automatically disabled when the pitch or roll values exceed $\pm 30^\circ$. This means that the device is passive unless it is held relatively level, and if a user puts the sensor down or has it in a pocket, irritating superfluous feedback will not be produced.

Sprung centroid

The geometric centroid of the positions of a group of people is simply the mean of their positions (for small distances, ignoring the effects of the Earth's spherical geometry). However, because sensed positions are subject to GPS noise, and because position updates may be interrupted when the inherently unreliable network transport fails, simple mean computation leads to jumpy updates in centroid position. Instead, the centroid is considered to be a point under the influence of forces in a pseudo-physical system. These forces are the attractive effects of participants in the *Social Gravity* network. By adjusting the springiness and damping of this simulation, the smoothness of the centroid's movement can be adjusted. The physical modelling approach leads naturally to schemes where users can probe and stimulate the network of springs that bind them together.

The computation of the centroid naturally leads to a client-server model, where each phone sends a position update to a server. The centroid is continuously computed, and the result placed in a common database. Because the the *Social Gravity* system is designed for pedestrian navigation, it does not require extremely rapid updates of centroid position. Updates on the order of 10s or so are acceptable; tests with the simulator showed navigation was not significantly impaired with delays up to 30s. The average velocity of the centroid in the trials was $0.47\ \mathrm{ms}^{-1}$. This is 0.26 degrees/s at 100m and 0.05 degrees/s at 500m. This means the centroid will have swung only 2 degrees or so in a ten second update interval, even when participants are close to the target. Communication was implemented via a web server, using commercial 3G networks. The system can reliably update the centroid on a 4s schedule, and the asynchronous protocol means that the system is tolerant of devices temporarily losing communication.

REAL WORLD EXPERIMENT

We performed a field trial to examine potential usage of the system in a realistic scenario. Our research questions were:

Viability: Can the tool be used by a group of pedestrians to help them meet up?

Impact: What is the effect of using the system on the users' behaviour, compared to walking behaviour without the system's guidance?

Five separate one-hour trials were performed to help understand these questions. 25 participants aged from 18 to 65 were recruited for the trials, in 5 groups of 5 people. 15 participants were male, 10 female; 6 students and 19 members of university staff. None of the participants worked in areas directly related to HCI.

Tasks

Participants completed two primary tasks during the study:

Rendezvous: From an initial starting position, use the device to scan for the location of the meeting point, then attempt to find this, using their own judgement at any path choices until meeting up with other participants.

Free walking: Repeat the task, but this time use their own choice of route to the same meeting point, with no feedback to guide them. Completing this second task allowed us to compare between participants' normal behaviour and that when using the system.

Measures

A large quantity of data from both logs and observations were collected from each trial, allowing us to measure the success of the system against each of our research questions. Each participant was also observed while using the system, and each group of participants was interviewed after the trial.

Viability: Measured as the percentage of rendezvous that were successful. In addition, we measured the time taken to meet up relative to the time predicted in the simulation.

Impact: We measured the impact of the system on normal behaviour by comparing the two routes participants took to the meeting point, both in time and distance, and the effect using the system had on each participant's percentage preferred walking speed (PPWS, see below). In addition, examining the time spent scanning for the centroid compared to the time spent walking without scanning allowed us to assess the impact of the act of scanning for feedback.

Walking speed as a measure of usability

Walking speed can be a useful proxy measure for the usability of a system. If someone is able to walk at their comfortable natural speed while interacting with the system, this suggests that the interaction is not seriously disturbing their normal moving behaviour. In the *Social Gravity* system this is particularly relevant, as the task requires pedestrian navigation throughout. The system would be of little use if the interaction was so distracting that users could not walk and navigate effectively. PPWS has previously been used as an evaluative measure to assess mobile interactions, and Petrie *et al.* [17] argue that it can be used as a measure of a device's effectiveness. Pirhonen *et al.* [18] found PPWS to be a sensitive measure of the usability of a mobile MP3 player, where an audio/touchscreen interface affected walking speed significantly less than the standard graphical version. Kane *et al.* [12] subsequently adapted the technique and used it to test 'walking user interfaces'. PPWS has traditionally been used as a summary statistic for a whole trial. We compensate for the increased variation in the outdoor environment by using a much higher resolution analysis than is commonly used, down to the level of individual steps. This allows comparisons of behaviour at different stages and conditions in the experiment.

Procedure

At the start of each session each participant was met by a researcher and introduced to the system and its purpose. Each participant's meeting location was separate to minimise effects from participants recognising each other before meeting up. Participants then used the system for a short training session, lasting no more than five minutes, in which they were able to feel example feedback and use the system while moving. Each participant was then taken to a location at the edge of the university campus. The five starting points were the same for each session, and were spaced evenly at the edges of an area of approximately 0.5km^2. When all participants were ready, each began the task while the researcher observed their behaviour from a short distance behind them. When all participants met up they were led back to their starting points and asked to make their way to the rendezvous point a second time, this time using their own choice of route rather than using the feedback to guide them. This provides a baseline measure of the best possible perfomance where users know exactly where to go and do not need to interact with the system. Finally, when all participants met up for the second time a short interview was conducted. At the end of the study each participant was rewarded with a bookstore gift voucher as a token of our appreciation.

RESULTS AND ANALYSIS

All participants successfully completed the *Social Gravity* guided rendezvous task. They were all also able to make their way back to the meeting point after being returned to their respective starting points. In the rendezvous case the participants took 13:05 minutes on average, with a standard deviation of 1:50. In the return case, where users were simply walking back to the (known) rendezvous point without guidance, they took 7:44 minutes (0:55 std. dev.). The mean time taken for the last participant to arrive back in the free walking case (a fairer comparison with the rendezvous task) was 9:45 minutes. This time difference is almost entirely due to walking further in the rendezvous case: 992m (193m std. dev.) for the rendezvous case, 573m (161m std. dev.) for the free walking case. This extra walking may be due to the limited ability of participants to plan efficient routes around obstacles in the rendezvous case.

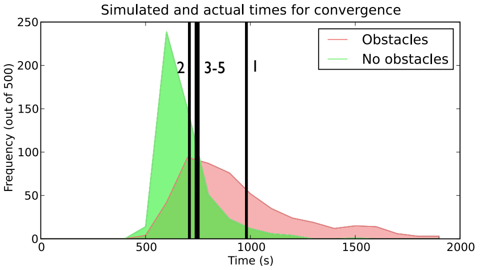

Figure 3. The simulator predicted times for convergence (from 500 runs) compared with the actual measured times (marked with vertical lines). The actual times fall well within the predicted times. Trials 3–5 are bunched very closely together. The long tails in the obstacle case are due to agents becoming stuck due to limited pathfinding ability.

Comparison to simulated results

Figure 3 compares the simulator results with the real field trial results. The times participants took to rendezvous agree well with those predicted by the simulator, falling between the estimated times for the model with obstacles and the model without. This suggests that the obstacle avoiding algorithm was, as expected, unrealistically simplistic in its choice of action. Humans are able to reason about ways to reach targets beyond complex obstacles so that an efficient path can be planned, which is something the simulator did not take into account. However, this does not detract from its usage in this situation: accurate estimations of behaviour have been produced that agree with the real world results.

Walking speeds throughout the trial

Participants' gait movement can be used to determine their walking speed throughout the trial, and was measured using the accelerometer in the SHAKE held during the tasks. To estimate the gait motion, the total acceleration is computed, which makes measurements insensitive to orientation variations. This is then filtered to remove drift and high-frequency noise, and a Hilbert transform is used to extract the gait phase angle [5]. Using this measure, the influence of factors on walking speed throughout the trial can be analysed in detail.

Figure 4 shows variation of walking speed in the rendezvous and free walking cases. There is relatively little difference, with slightly more slow walking in the free case, and more time spent at standstill in the rendezvous case. Figure 6 shows that most of this time is due to waiting around near the eventual meeting point while other participants arrived, which they did not have to do in the free walking condition. Figure 5 shows walking speed split into categories depending on whether the user was actively scanning for a target. Although users sometimes stopped to find a target heading, most of the time they walked and scanned at the same time.

Convergence behaviours

Figure 7 illustrates the paths chosen by participants for each rendezvous task, coloured according to their walking speed along the route. Walking speeds are similar for a large part of each route, with most participants slowing down only when near the centroid. Looking more closely at the individual routes taken by participants, we can see bottleneck areas

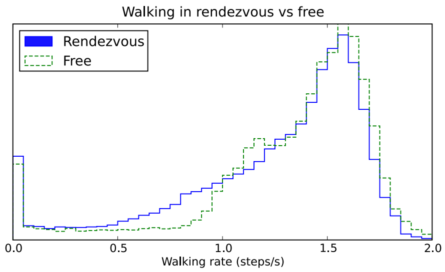

Figure 4. Histogram of walking speeds in the rendezvous and free walking cases, shown as a normalized density. The distribution of walking speeds is broadly similar.

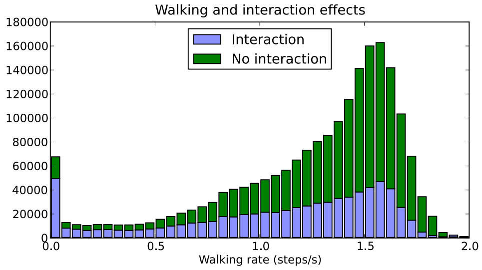

Figure 5. Histogram of walking speeds for when participants were using the *Social Gravity* system. Each bar is split into two regions; the lower shows steps where the user was engaged in probing for the target and the upper section shows steps where the user was simply walking. The bar sums to the total number of steps at that walking rate. Users were clearly able to walk quickly and interact.

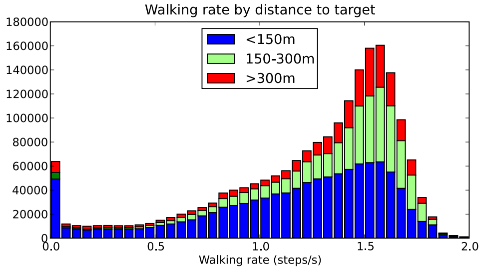

Figure 6. Histogram of walking speed for all participants in the rendezvous case, where each bar is subdivided into walking rates for three distance ranges from the centroid. Most of the time spent standing still is when the participant is close to the target (the peak at 0 steps/s is mainly where the user is $<$150m from the target).

where location constraints have limited path options, forcing participants along similar routes. When route options are more open, the paths taken are more varied. In some cases participants stopped or reversed their steps at various points during the route. This could be for a variety of reasons: encountering a dead end or obstacles, a change in the centroid position, or perhaps just a desire to take a different route. Regardless, we can see that participants have been willing to take alternative paths. This suggests that they trust feedback sufficiently that they are willing to take different routes, knowing that they will be brought back on target.

Participant observations and feedback

All except two participants held the device in their hand by their waist, while two chose to scan with their arm held almost horizontal. Several participants encountered locations where they were unable to directly follow the feedback due to an obstacle in their way. In this situation they kept scanning to one side for feedback for reassurance.

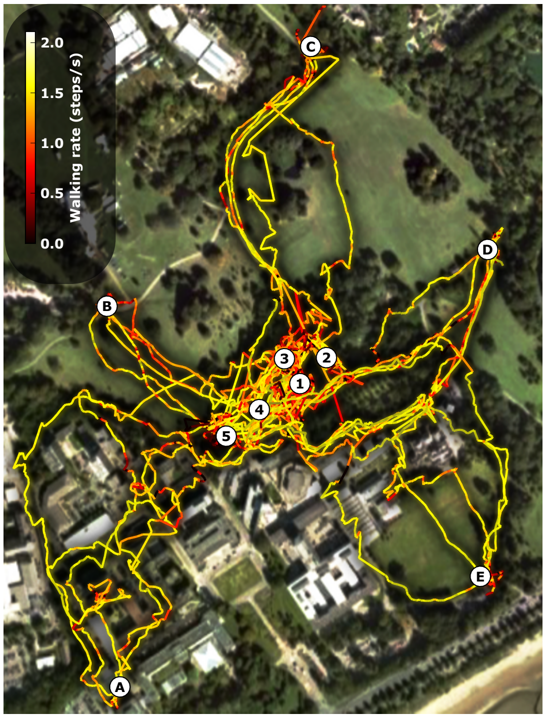

Figure 7. Paths taken by all 25 participants during the study, starting at A-E, ending at 1-5. Lines are coloured according to the participant's walking speed for each part of the route (steps per second, see key). Generally people walked at a constant pace, with occasional stops. Stopping was more frequent in the vicinity of the rendezvous point.

The interviews recorded after completion of the rendezvous tasks provided valuable insights, and many suggestions for potential use, including: festivals and large gatherings of people; when in unfamiliar locations; solo navigation to a fixed location; an adaption for blind users; as part of a large area game; in lieu of calling to avoid high costs. The majority of participants said they would definitely use the system if incorporated into their phone, but pointed out issues with the final part of the rendezvous task. When very near to each other, participants had often been given conflicting centroid directions due to GPS errors. A commonly-suggested solution was a form of alternative notification to alert them that they were near the final meeting point. Four participants were unfamiliar with the study area but had no problems using the feedback to navigate to the rendezvous point. Four participants noted issues in trusting the system – over time they could begin to trust it, but in a short study they were cautious in accepting the directions given. Two participants said they had had to adapt their behaviour to cope with the slight lag in vibration response, finding they needed to scan several times to get an accurate indication of the direction they needed to head toward. This highlights the importance of minimising feedback delays, which can be a problem on most mobile phones as their software architectures are often not compatible with tight real-time performance of vibration feedback.

Summary of results

The field trials demonstrated that users could meet up rapidly using just the vibration cues, even in a busy environment with many intervening obstacles. The comparison with the free walking condition indicated that the time for all five users to meet up was just under twice the time they would have taken had they been walking without guidance to a known destination. This is a very small overhead for non-visual navigation, and is especially compelling because the interaction requires no set up time. Most of the stopped or very slow movement occurred when users were almost converged – much of the time difference is due to users waiting around for others. Around 65% of the time when users were stopped occurred at the eventual rendezvous point. The rest of the time users were able to continue to interact while walking. The interaction did not significantly alter walking speed; users were able to interact and walk at the same time with little difficulty. Users interacted around half of the time they spent walking and actually interacted with the system more while moving faster. These results would be unlikely with a visual feedback system; people generally find it very hard to walk and look at a screen at the same time, even in lab environments [18, 21].

DISCUSSION

All participants were able to use the system to meet up with only directional vibrotactile feedback to guide them. The low resolution of the feedback did not greatly affect participants' behaviour, in line with our expectations from simulations. Participants used the cues to meet up, but they did so intelligently, without getting stuck behind obstacles. The use of minimal, appropriately uncertain feedback seems to have been successful in avoiding a "navigation blindness" situation, where users place excessive trust in guidance information. As might be predicted, the system was of less use for the final stage of navigation when participants were close to each other, mainly due to GPS inaccuracy. Interestingly, in trial 1 this provoked participants to follow the device's feedback despite having met up with other users. Previous research [7] suggested that users appreciate a constant update of positional information, and this was reflected in our results, with most participants preferring to constantly feel the feedback as they moved. This was presumably influenced by the specific request to use the device to meet up. Had we given a distracter task then scanning may have become a background task. In addition, users' lack of familiarity with the system, and with haptic feedback systems in general, no doubt affected their behaviour. If the participants had been more experienced it is likely constant monitoring would decline as their confidence increased. In our own pilot studies, experienced users successfully rendezvoused while scanning on an intermittent basis (90s approx.). Developing a simulator, using it to design test scenarios, and then evaluating the real system using these scenarios proved to be an effective technique in minimising the fiddly and time-consuming configuring of prototypes. The development cost of creating the simulator was far outweighed by the simplification of the subsequent prototype development. It also gave us the chance to work out concrete metrics and statistics and test the tools for analysing the data captured before the prototypes were even written. Data captured during the trials were recorded in as rich and as raw a form as possible, at the finest timescales available, so as to allow the widest range of possible analyses. Although this caused some difficulty in implementation due to the storage requirements of such detailed log files, it proved essential when verifying that subsequent analysis was accurate and valid.

Future extensions

One interesting extension for larger group sizes would be a system optimised to first bring small groups together, and then merge these, offering potential social and safety benefits. The structure of the meetup would then change from contraction to a point to that of a rubber sheet pulling together across of the obstacles of the environment. The waiting and wandering around the rendezvous location could be solved with conventional interaction techniques such as audio alerts or visual feedback. Developments could also use constraint models; for example, the meeting point could be directed away from the inside of buildings or towards well known landmarks. Richer communication via the 'tether', including sensing distance to goal, convergence rates of the group, or 'pinging' rendezvous partners could be included. Context-sensitive inference could alert potential participants to join or leave the rendezvous, and other participants could sense these events via the feedback from the tether.

CONCLUSIONS

This paper introduced a novel, privacy-preserving and dynamic mobile meetup system that helps groups navigate to a mutually-convenient meeting point. The system extends connections between socially related but geographically distant users using inertial sensing to ground users in a common space. Users successfully navigated through a complex environment while negotiating dead ends and moving meeting point positions. The system requires only occasional interaction and the feedback is non-visual and minimally intrusive. Users were able to interact while walking at a comfortable speed; navigation did not significantly disrupt their normal movement. Simply agreeing a mutually convenient meeting place via texting or voice channels would take many minutes for this number of people, and would be very difficult in unfamiliar or featureless environments. Groups of five people were able to meetup in only three minutes more than it took *to walk to a known point*. This overhead should be relatively constant as the number of users scale. We have also shown that only very crude angular feedback is needed for efficient target finding in pedestrian navigation, even in constrained environments. The use of simulators for studying likely behaviours in mobile trials has been clearly demonstrated; even a simple simulator made useful predictions about behaviour in the mobile environment at little cost, and allowed effective pre-design of evaluation metrics. *Social Gravity* sidesteps the effort of agreeing rendezvous location and preserves each user's privacy, without impairing ability to meet up quickly with a minimum of interaction. The social possibilities for such dynamic fused spaces are substantial and the the ability of users to engage via such spaces will grow as direct interaction techniques such as these are developed.

Acknowledgements

This work is funded by EPSRC projects EP/E042171/1 and EP/E042740/1. Nokia donated devices. Drawings by R. Donaldson. S. Strachan and P. Eslambolchilar helped in the early stages of the project, A. Ramsay wrote drivers for the SK6 devices. Data from the trials is available at `http://negotiatedinteraction.com/socialgravity`.

REFERENCES

1. J. Axup, S. Viller, and N. J. Bidwell. Usability of a mobile, group communication prototype while rendezvousing. In *Proc. CTS '05*, pages 24–31. IEEE, 2005.
2. M. Batty, J. Desyllas, and E. Duxbury. Safety in numbers? modeling crowds and designing control for the notting hill carnival. *Urban Studies*, 40(8):1573–1590, 2003.
3. M. Colbert. A diary study of rendezvousing: implications for position-aware computing and communications for the general public. In *Proc. GROUP '01*, pages 15–23. ACM, 2001.
4. M. Colbert. Rendezvousing at familiar and unfamiliar places. *J. Navigation*, 57(3):327–338, 2004.
5. A. Crossan, R. Murray-Smith, S. Brewster, J. Kelly, and B. Musizza. Gait phase effects in mobile interaction. In *Proc. CHI '05*, pages 1312–1315. ACM, 2005.
6. D. Dearman, K. Hawkey, and K. M. Inkpen. Rendezvousing with location-aware devices: Enhancing social coordination. *Interacting with Computers*, 17(5):542–566, 2005.
7. D. Dearman, K. M. Inkpen, and K. N. Truong. Mobile map interactions during a rendezvous: exploring the implications of automation. *Pers. and Ubiquitous Comput.*, Published online 4 March 2008, DOI: 10.1007/s00779-008-0195-2.
8. S. Holland, D. R. Morse, and H. Gedenryd. Audiogps: Spatial audio navigation with a minimal attention interface. *Pers. and Ubiquitous Comput.*, 6(4):253–259, 2002.
9. C. W. Johnson and L. Nilsen-Nygaard. Extending the use of evacuation simulators to support counter-terrorism: Using models of human behaviour to coordinate emergency responses to improvised explosive devices. In *Proc. ISSC '08*. System Safety Society, 2008.
10. L. Johnson and C. Higgins. A navigation aid for the blind using tactile-visual sensory substitution. In *Proc. EMBS '06*, pages 6289–6292. IEEE, 2006.
11. M. Jones, S. Jones, G. Bradley, N. Warren, D. Bainbridge, and G. Holmes. Ontrack: Dynamically adapting music playback to support navigation. *Pers. and Ubiquitous Comput.*, 12(7):513–525, 2008.
12. S. K. Kane, J. O. Wobbrock, and I. E. Smith. Getting off the treadmill: evaluating walking user interfaces for mobile devices in public spaces. In *Proc. MobileHCI '08*, pages 109–118. ACM, 2008.
13. M. Lin, Y. Cheng, and W. Yu. Using tactons to provide navigation cues in pedestrian situations. In *Proc. NordiCHI '08*, pages 507–510. ACM, 2008.
14. S. K. Nagel, C. Carl, T. Kringe, R. Märtin, and P. König. Beyond sensory substitution—learning the sixth sense. *J. Neural Engineering*, 2(4):13–26, 2005.
15. T. Nicolai, E. Yoneki, N. Behrens, and H. Kenn. Exploring social context with the wireless rope. In *Proc. OTM '06 Workshops*, LNCS, pages 874–883. Springer, 2006.
16. S. Olofsson, V. Carlsson, and J. Sjölander. The friend locator: supporting visitors at large-scale events. *Pers. and Ubiquitous Comput.*, 10(2-3):84–89, 2006.
17. H. Petrie, V. Johnson, S. Furner, and T. Strothotte. Design lifecycles and wearable computers for users with disabilities. In *Proc. MobileHCI'98*, 1998.
18. A. Pirhonen, S. Brewster, and C. Holguin. Gestural and audio metaphors as a means of control for mobile devices. In *In Proc. CHI '02*, pages 291–298, 2002.
19. T. Sokoler, L. Nelson, and E. R. Pedersen. Low-resolution supplementary tactile cues for navigational assistance. In *Proc. MobileHCI '02*, volume 2411 of *LNCS*, pages 525–539. Springer, 2002.
20. S. Strachan, P. Eslambolchilar, R. Murray-Smith, S. Hughes, and S. O'Modhrain. gpsTunes: controlling navigation via audio feedback. In *Proc. MobileHCI '05*, pages 275–278. ACM, 2005.
21. S. Strachan and R. Murray-Smith. Bearing-based selection in mobile spatial interaction. *Pers. and Ubiquitous Comput.*, 13(4):265–280, 2009.
22. D. H. Titterton and J. L. Weston. *Strapdown Inertial Navigation Technology (2nd Ed.)*. IET, 2004.
23. J. B. F. van Erp. Vibrotactile spatial acuity on the torso: Effects of location and timing parameters. In *Proc. WHC '05*, pages 80–85. IEEE, 2005.
24. J. B. F. van Erp, H. van Veen, C. Jansen, and T. Dobbins. Waypoint navigation with a vibrotactile waist belt. *ACM Trans. App. Percept.*, 2(2):106–117, 2005.
25. B. N. Walker and J. Lindsay. Navigation performance with a virtual auditory display: Effects of beacon sound, capture radius, and practice. *Human Factors*, 48(2):265–278, 2006.
26. J. Williamson, R. Murray-Smith, and S. Hughes. Shoogle: excitatory multimodal interaction on mobile devices. In *Proc. CHI '07*, pages 121–124. ACM, 2007.
27. J. Williamson, S. Strachan, and R. Murray-Smith. It's a long way to Monte Carlo: probabilistic display in GPS navigation. In *Proc. MobileHCI '06*, pages 89–96. ACM, 2006.