Increasing the Awareness of Daily Activity Levels with Pervasive Computing

J. Maitland, S. Sherwood, L. Barkhuus, I. Anderson, M. Hall, B. Brown, M. Chalmers, and H. Muller

***Abstract*—Public health promotion technology should be accessible to the general public at which it is aimed. This paper explores the potential for use of an unaugmented commodity technology—the mobile phone—as a health promotion tool. We describe a prototype application that tracks the daily exercise activities of people carrying phones, using fluctuation in signal strength to estimate a user's movement. In a short-term study of the prototype that shared activity information amongst groups of friends, we found that awareness encouraged reflection on, and increased motivation for, daily activity. We describe some of the details of the pilot study, and conclude with our intended plans to develop the system further in order to carry out a longer-term clinical trial.**

***Index Terms*—Health Promotion, Commodity Technology, Physical Activity, Collaborative System**

I. Introduction

The decreasing levels of daily activity undertaken by the general public form an ongoing challenge for those involved in public health, and is of concern to both primary and secondary healthcare. The benefits of physical activity are well documented and widely acknowledged and yet the World Health Organisation state that 60% of the worldwide population are not active enough to profit from these benefits [26]. Pervasive and ubiquitous computing technologies are well-suited for use within the healthcare industry and have the potential to be far-reaching and effective. This paper presents a prototype application that runs on arguably the most pervasive computing technology of all, the mobile phone. By detecting patterns in signal strength fluctuation and changes in the visibility of GSM[1] cells, the application can infer whether the carrier of the phone is sitting still, walking, or travelling in a car. This information is then used to calculate the carrier's daily activity level, which can then be shared with and compared to the activity levels of others. Rather than being driven by experimental hypothesis or outcomes, we took a more exploratory approach. A short-term pilot study gauged information regarding usability, user response and attitudes toward the prototype, and will inform the future refinement of the system so that it is suitable for longer-term clinical trials.

Augmentation of traditional exercise technologies and practices with more pervasive and ubiquitous computing is becoming an established area in both research and commercial arenas. Gym-based equipment such as treadmills and rowing machines has been complemented with virtual reality environments to motivate and stimulate users during their workout (http://www.fpgamerunner.com). Some of this equipment also facilitates the downloading of workout data onto an individual's PDA, so that personal workout programs may be monitored and adapted (www.acrocat.com). Positioning technology is now also being used to extend such services beyond the gym environment to benefit walkers, cyclists and road-runners. These technologies assist individuals who have already taken steps to get fit or remain healthy. However, there is relatively little in the way of assistive or motivational technologies that are aimed at the more sedentary adult/child. Pedometers are a widely used fitness-related technology that does not demand a vested interest in health in order to be used. Although their accuracy may be volatile, they have been found to motivate individuals taking early steps towards a more active lifestyle [21, 22]. As pervasive technologies advance, the ability to detect and monitor the physiology and physical activity levels of an individual or community are of increasingly fine granularity. A multi-modal sensor board can now distinguish between 8 physical activities [12], and commercially available technology can be worn on the body to monitor blood pressure, heart rate, and stress levels. These technologies are unarguably useful, but their specialist nature may prove to be a barrier to widespread adoption and utilisation.

The recommended level of activity for an adult is at least 30 minutes of moderate activity, five times a week. Although prolonged periods of activity are most advantageous, the daily amount of 30 minutes can be accumulated throughout the day in shorter periods of 10 minutes or more [6]. Most adults who do not currently reach this level of activity may be able to achieve this target by making small changes to their everyday routine. By capturing and acknowledging everyday activity in an accessible and non-invasive manner, and facilitating the sharing and comparison of that information between peers, we hope that awareness would be raised in such a way that it motivates users to become more active on a day-to-day basis.

Manuscript received October 2, 2006. This work funded by UK EPSRC (GR/N15986/01).

J. Maitland, S. Sherwood, L. Barkhuus, M. Hall, B. Brown, and M. Chalmers are with the Department of Computing Science, University of Glasgow, Glasgow, G12 8QQ (phone: +441413302333; email: {jules, sherwood, barkhuus, mh, barry, matthew}@dcs.gla.ac.uk).

I. Anderson and H. Muller are at the Department of Computer Science, University of Bristol, BS8 1UB (phone: +441179545264; email: {anderson,henkm}@cs.bris.ac.uk).

[1] GSM is used here since it is the most common mobile phone technology in many parts of the Western world including Europe and increasingly in the United States. In theory other technologies such as CDMA can be used for the same inference of activity.

Shakra[2] is the first prototype of such a system that runs on an unmodified mobile phone. Although not everybody owns a mobile phone, it is the most uniformly adopted technology throughout all classes [9], and so this platform hopefully overcomes the aforementioned barrier to adoption and, therefore, effect.

The following section reviews conceptual approaches to behavioural change in physical activity, alongside current technical approaches to fitness tracking and motivation. The resultant design and implementation detail of Shakra follows, before the pilot study is presented and discussed. Following the discussion of our study, implications for further development are presented, which are aimed at improving the Shakra system for a larger clinical trial. These are also presented as guidelines for future development of similar health related systems.

II. The Problems of Motivating Exercise

Numerous studies show how just a minimal amount of daily activity can increase general health, such as lowering blood pressure, and can lead to weight loss among overweight people—not to mention the related social and mental health benefits [16, 20]. Having a lifestyle that promotes regular exercise seems to be a challenge in the western world, since our daily lives are busy and many of us draw upon transportation systems, such as cars, trains and buses. Studies suggest that around 70% of the UK population fails to meet minimum recommendations for physical activity [1].

In view of the aforementioned public health recommendations for minimum physical activity levels, many approaches to increasing fitness propose an increase in *moderate* activity, such as brisk walking, in order to improve people's health [7]. Moderate activity is generally defined as when a person's heartbeat is increased to 55–69% of maximum heart rate, which for many people would occur when walking at about 4 miles per hour. One important factor for consideration is that many people may have difficulties making sure that their activity is in fact moderate and not just light, i.e. that they are achieving the health benefits stated above. It is therefore important for individuals to not only be aware of their overall amount of exercise but also its intensity.

A. Tracking and motivating fitness and moderate activity

Many technical methods have been developed to measure physical activity. One common device is the pedometer, a small device that measures each stride the wearer takes. One recent report indicates that just the presence of the pedometer can motivate people to be more active [21, 22]; another study showed that sharing daily activity information within a small group of friends was more satisfying and motivating compared to a control group who measured but did not share their information [5].

One of the most advanced commercial technologies in this area is the BodyBugg (www.bodybugg.com), also known as SenseWear. The BodyBugg measures an array of values such as relative body temperature, step count and acceleration, in order to estimate how many calories the wearer is burning. It has been shown to work reliably in controlled tests for measuring calories burned, with an accuracy of 89–98%; however it is limited in its determination of the actual context of the wearer [14]. Also, it has to be worn on the upper arm for 24 hours a day; it can therefore easily disrupt sleeping and collide with everyday clothing—a particular disadvantage among women who often wear tighter or lighter clothes.

A less direct means to motivate activity is taking part in mobile games. Most mobile games involve infrequent play over a relatively short period, with limited health benefits, but some games such as Mogi Mogi (www.mogimogi.com) and Feeding Yoshi [4] take place over a longer period of time and are more 'interwoven into everyday life'. A study of Mogi Mogi showed that players would frequently take detours from their normal routes, and that "many alight at an unusual metro station on the way home if they notice an object on their mobile screen, even if this means walking much further to get home. Many players also said they went out at night because the mobile screen had indicated objects in the vicinity" [13]. Similarly, Bell et al. report that players adjusted their everyday routines of work and travel so as to spend more time playing the game, often walking a good deal more than they would do normally. A disadvantage of this approach is its relative lack of clarity or precision about the exercise undertaken. While players increase their activity as part of playing the game, this is not directly connected to or encouraged with the game—instead it is a useful but indirect benefit of the game.

B. Theories and Studies of Change in Activity

Numerous studies have explored how to motivate people in increasing their activity level, and there are two well-cited theoretical approaches: the Transtheoretical Model, where behaviour change is described as a multi-stage process [17] and Social Cognitive Theory, based on the individual's outcome expectancy and self-efficacy [3].

The *Transtheoretical Model* is one of the more common theories referred to in the health literature. It focuses on the individual stages people go through with regard to physical exercise regimes, such as pre–contemplation, contemplation, preparation, action and maintenance. Although it is possible to determine people's individual stage at a given time with a standard questionnaire, the theory does not account for individuals' different *levels* of exercise and it does not address the possibility for individuals to skip between the stages. One critique has also been that it is focused on attitude rather than behaviour, although in observational terms both seem to be significant. For example, it has been pointed out that the difference between the stages of pre–contemplation and contemplation only refers to a change in attitude rather than actual change in physical activity. Moreover, recent research points to the theory's weakness in showing long-term changes [1].

The *Social Cognitive Theory* focuses on increasing the individual's self-efficacy by different means, in relation to keeping fit, leaning on studies that show how intrinsic motivation (enjoyment, feeling good about the exercise) rather

[2] Shakra is a composite of the words Sharing and Chakra, the latter being the body's centres of spiritual energy according to yoga philosophy.

than extrinsic motivation (external pressures or immediate rewards) increase the likelihood that the person will stick to a routine [15]. Examples of interventions using this approach include giving health advice over the phone, either by health professionals or via an automated service, and through an Internet service [10, 11].

Other research has addressed social aspects of sport participation and physical activity, finding that sharing information about activity and exercising together can increase interest, enjoyment, and motivate some individuals to do *more* activity [5, 8, 20, 25]. Similarly, when people receive tailored information that is personally relevant, it is more likely to stimulate change, adding to people's self-efficacy and outcome expectancy [23, 24]. It is evident that, although not determined by social factors, intrinsic motivation is affected by wider social interaction.

Behavioural change is difficult to promote, and many researchers point to the combinations of internal and external influences that are complicated to trace, target and categorise in individual cases. One critique that has been made of the physical activity literature, for example, is that it does not separate between individual environmental values (such as age, social class, health status) and social environmental values (such as family, school/work and community) [7]. The social cognitive theory addresses aspects of community, in contrast to the Transtheoretical model, although it focuses on internal values as main motivator to increase individuals' level of exercise. Although such categorisations abstract over individual cases, it is again reasonable to conclude that an individual's level of physical activity is affected by interactions with his or her surrounding group. In our work, we therefore focus on social and communal aspects of exercise; the light pressure from the surrounding community is a great motivational factor not to be underestimated in relation to intrinsic motivational factors. Also, rather than taking a broad survey and relying on social categories such as class, in our evaluation we focus on the details of particular individuals' experience. Based on our understanding of related theories and studies, our system design is directed towards a long–term goal of achieving greater public health. We assume that a member of the general public is likely to make only minor behavioural changes, and that this will be based on individual awareness as well as social interaction.

III. The Shakra Prototype

Our overall aim is to design and implement a system that will help to motivate adults who do not currently achieve the minimum recommended daily activity level, and who can benefit from a raised awareness of their current levels of activity: a system that can track and categorise an individual's daily activity into accumulative time spent in inactivity, light, moderate, and vigorous activity. In acknowledgement of the influence that social networks can have on the actions of an individual, the system should facilitate the sharing and comparison of data between peers. In order to evaluate user response to such a system and general usability, a basic prototype was created that determined whether a user was active or inactive, accumulated daily totals, and allowed the sharing and comparison of the daily totals.

Our key design goal for Shakra was that it could be carried around in a non-intrusive manner, requiring little or no extra equipment for users. Minimal user intervention is required in order for it to function effectively; the system tracks the activity of the user without direct manual input. The application tracks users' general level of activity, showing the current mobility state: no movement ('stationary'), moderate activity ('walking') and travelling in a car, bus or train (collectively labelled here as 'driving'). The moderate activity is then used to display a 'minutes of activity per day', with a historical view supporting comparison of activity across the previous week. This supports a user monitoring his or her own activity and exercise levels, with the exception that stationary exercise (such as working out at a gym) is not tracked.

When running the application for the first time, the user is prompted to provide a name, used to identify him or her within the system and to other users. The application records up to seven visible GSM cells and their signal strengths, once per second. The current activity of the user is then classified every 30 seconds by the application's neural network, as described in more detail below. Using a web service, each phone uploads the recorded activity of the user via GPRS and stored on a MySQL database, while simultaneously downloading information about other participants for later review. The system updates this shared information automatically every hour. If a user does not want to wait for an update, he or she can manually synchronise via the *Sync* menu option.

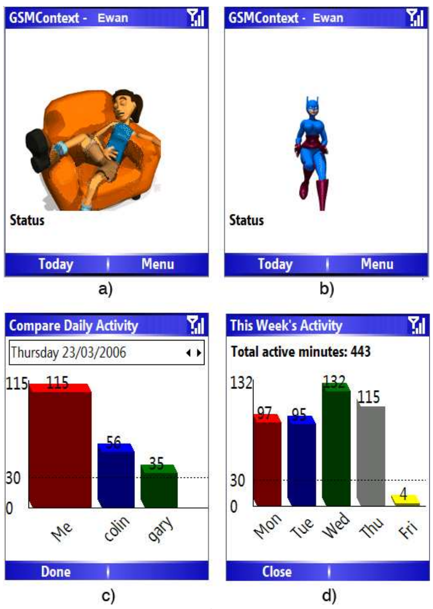

Figure 1: The phone interface. Images a) and b) show two of the screens showing the estimated current activity level: *Stationary* and *Walking*. Images c) and d) show screens for examining relative and individual activity levels: *Compare Daily Activity* and *This Week's Activity.*

Users specify in advance the peers they wish to share results with, but at any time they can change the list of peers whom they wish to exchange information with. Figure 1c shows the *Compare Daily Activity* screen that users can view to assess their performance in relation to their peers. For a week's overview of their own activity, users may use the *Week's Activity* screen shown in Figure 1d. In order to provide real time feedback to the user an animated representation of the user's current mode of activity runs continuously on the main screen of the application, this is shown in Figure 1a and 1b.

A. Sensing Activity

The current activity of the user is inferred using patterns of fluctuation in GSM signal strength and changes to the ids of detected cells. This method has been demonstrated as a reliable and unobtrusive way of sensing current activity [2], and has the advantage over the more traditional approach of using an accelerometer in that it does not require additional sensor hardware as in Sensay [18] and the multi–modal sensor board of [12].

Rather like a traditional accelerometer, when a mobile phone is moved the levels of signal strength fluctuation change. For example, Figure 2 shows the total signal strength fluctuation across all monitored cells during successive 30-second time periods whilst walking, remaining still and travelling in a motor car. The figure illustrates that it is relatively easy to distinguish between moving and remaining stationary, but at times, the pattern of fluctuation whilst walking will match that of driving and vice-versa. This is due to the stop–start nature of both walking and travelling in a motor car in urban areas. When driving, a greater geographical distance will typically be covered over a given time period when compared to that of running or walking. As such it is possible to use the rate of change of neighbouring cells to infer travel by car.

To classify these patterns we use an artificial neural network. The network inputs are: the sum of signal strength fluctuation across all monitored cells and the number of distinct cells monitored over a given time interval. The network consists of a single layer of eight hidden neurons; weights are learnt using back propagation. The network outputs the currently sensed activity for the given input values. The network is trained by repeatedly presenting data collected during each method of movement.

The current activity of the user is conditionally dependent upon their previous activity. In order to provide instant feedback to the user interface, the neural network deliberately does not model this behaviour. Instead, when determining if any additional minutes have been earned, we apply task knowledge based upon the output from the neural network over the previous two and a half minutes. This enables noise to be filtered out and a more accurate representation of the users activities achieved. For example, periods of low signal strength fluctuation such as stopping at traffic lights whilst driving can be ignored when placed between periods of high fluctuation where many distinct neighbouring cells were monitored. It could be argued that activity would be more accurately inferred if a longer rolling filter had been applied to the GSM data. Introducing longer filters would have increased the likelihood of active minutes 'disappearing' from the users' activity totals. A decision was made that for the purpose of this study priority would be given to user experience; with the intention that this trade-off would be addressed in future work.

IV. The User Study

The Shakra application was evaluated with three groups, to detail its use, to determine whether it increased users' awareness of their activity level and if this could potentially motivate them to be more active, and to derive implications for future work. Naturally, a longitudinal clinical study lasting months or years would be needed to rigorously assess long–term changes in users' behaviour and health, but our one week trial served as a pilot evaluation of a potentially powerful activity promoting application. The focus was on the users' experiences with both the activity tracking and the sharing feature; it was important to find if sharing information was good for increasing awareness and motivate a more active lifestyle.

Before the trial, a base neural network had been constructed by using GSM data collected by the development team while sitting still, walking, and driving. In order to determine whether or not further personalisation of the network was required for each of the trial participants, the system was given to each participant for a two day training period. During this period, the participants were asked to

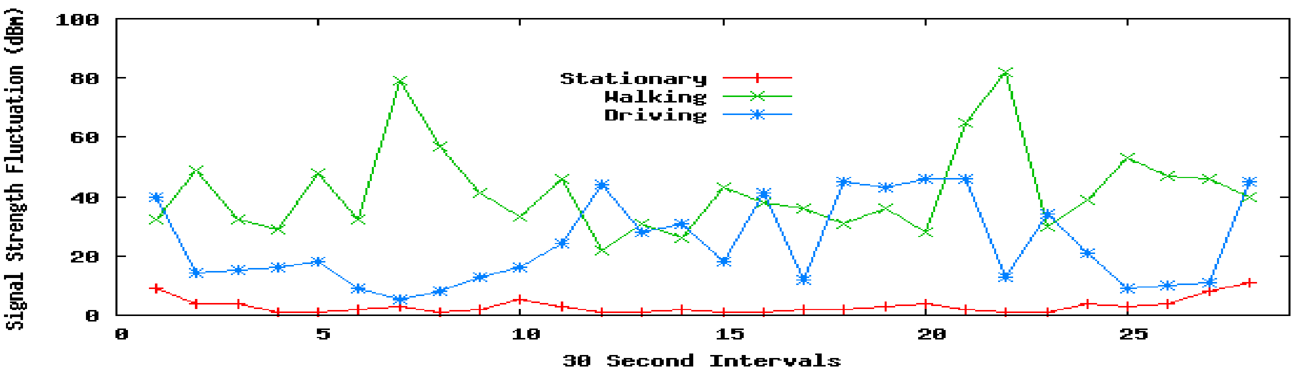

Figure 2: Distinguishable patterns of GSM signal strength fluctuation over successive 30 second samples are used in identifying the activity levels *Stationary, Walking* and *Driving.*

record whenever their activity mode changed. Functionally, this was a simple task supported in the application's main interface that users learned to do quickly. For the training days, we asked the participants to take the phone with them as they went about a normal day's activity. This trained the system for the areas in which they usually go to throughout the course of a day.

Following the initial system-training period, the data collected by the trial participants were analysed. We found that only minor changes to the previously trained neural network were required by three of the nine volunteers. This was due to them living and working in urban areas that exhibited different levels of signal fluctuation to those where the initial training data had been collected by the research team.

A. *Method*

Overall, the trial took place over ten days. The participants initially filled in a simple activity diary for three days, to determine their present level of activity and to compare activity to the week of using the application. Immediately after, they trained the system for two days and then finally used the system for a five-day working week, filling in a diary describing their use of the system and whereabouts for each day. We kept in touch with the participants by phoning them once during the week, and sending text messages in the few cases where it looked like the phone was not uploading properly. At the end of the study, each participant was interviewed individually to expand on the use and reflect on the experiences with and opinion of Shakra.

	Group 1	**Group 2**	**Group 3**
N	2	3	4
Age range	52-54	28-30	19-37
Sex	Female/male	Male	Female
Activity level	Fairly inactive	Two moderately and one highly active	One inactive, two moderately active and one highly active
Occupation	Teacher and administrator	Technical administrators	Manager, administrative staff and student

Table 1: Participants in each group

The participants were recruited as groups of friends and/or co-workers who had daily interaction with each other and would enjoy sharing their exercise information. Although the target users for the final system are inactive people, it is unrealistic to expect that only inactive people will use the system. This is especially true when the system is aimed towards peer groups who will naturally include individuals of differing levels of activity. We therefore aimed to study the use of the system among a diverse set of people, and the nine participants varied in the degree of their normal activity. Two were highly active, with purposeful exercise at least three days a week, four were moderately active people, working out one to two times a week, and three were fairly inactive, walking but not doing any purposeful exercise[3]. Table 1 provides an overview of the three groups.

After the study, the system logs were analysed. First of all, the activity times were compared to the self-reported diaries and the interviews, to make sure there was a fair level of accuracy in measuring activity. Secondly, the logs were scrutinised to see how participants used the application, how often they compared their activity to others', and how often they looked at their weekly chart. The interviews were transcribed immediately and the parts were categorised according to major topics and themes. They were used to elaborate on the diary, such as precise times of commute, actual transport methods and more detailed experiences and impressions of the application during the week. In the next section, we report the results in relation to three topics, one relating to precision or reliability of the application's measurements, a second looking at users' individual experience, and the third exploring the participants' experiences of information sharing.

V. RELIABILITY OF SHAKRA IN THE REAL WORLD

Although previous tests had shown highly accurate determination of activity [2], the real test of the application would be using it in an uncontrolled environment among many different people. We did not expect to get as high accuracy, because of the unstructured and diverse behaviour of people leading their everyday life. Overall, the application showed very good determination of activity and the participants found it very useful as a tool for measuring their activities. After analysing the diaries and annotating them with information gained through interviews, we compared each day of each participant with a log–generated activity timeline. It was easy to see participants commute to work, break for lunch, and commute back from work; two examples, with diary annotations, are shown in Figure 3. A rough analysis was done to determine the rate of correct labelling of activity. We chose three sample days for two different participants because their diary entries for those days were particularly comprehensive, i.e. six days in total. From the unfiltered data we analysed short stretches of 60 to 90 minutes with varied activity; this was done to refrain from considering the long hours of inactivity, which occurred during their workday where participants were mostly sitting at their desk. Including this would have given unrealistically optimistic numbers. Results showed a minimum of 70% accuracy during users' commute when fluctuations are highest. The misinterpretations often occurred during change between different methods of transportation such as getting off a bus or a train. However since there would often be a delay both before and after transportation, the misinterpretations would cancel each other out, correcting the accumulated minutes of exercise. One more problematic finding was that running occasionally would register as driving. During one

[3] Naturally this is a very broad characterization from the participant's own statements and diary reports. It is not necessarily a true reflection of their level of health or level of fitness.

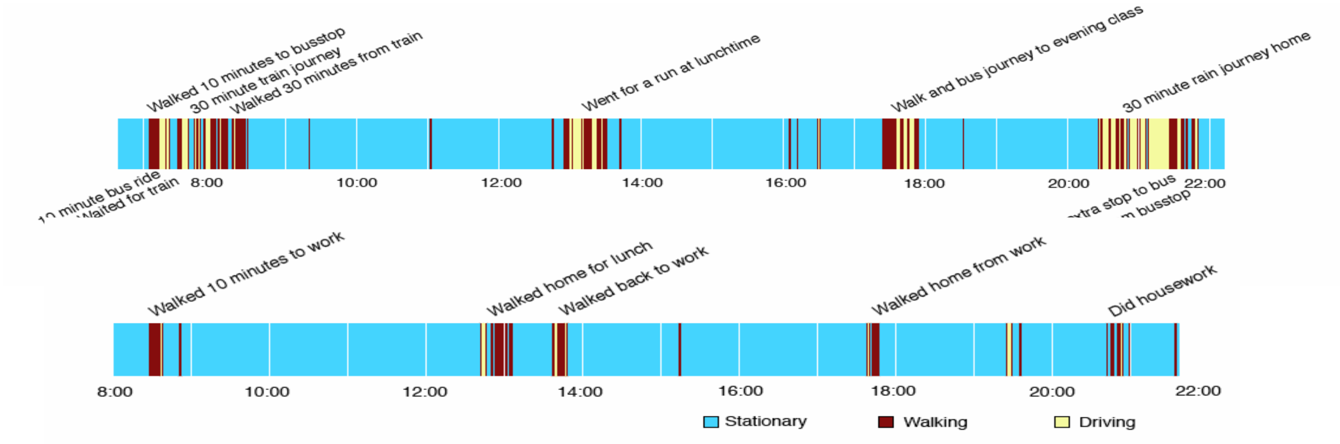

Figure 3: Example timelines of activity for two participant's days with colour showing the activity level and text showing the participants' diary annotations.

participant's 45-minute lunch run, 15 of the minutes were registered as driving. For another participant with a long commute for example, it meant that he gained a maximum of seven active minutes each day due to error. This was the maximum error we found from looking at participants' commutes.

Some of the diary entries assisted in showing when still or walking activity was misidentified. For example, one woman from group 3 explained that she went on a walk for 30 minutes, but had only increased her overall activity count by 22 minutes when she returned. It should be noted that this particular participant lived in the countryside where we knew that the neural network would be less accurate in the present version. Similarly, a male participant reported that his 10 minute walk to work sometimes only gave him 7 to 8 minutes of activity. This may in part be attributable to a lag in activity determination, as well as the participants stopping at road crossings, etc. Since the application is aimed towards increasing awareness rather then measuring physical exercise precisely, and offered useably accurate overall measures, we suggest that the small moment-by-moment lags and jitters in classification were not problematic. Post-processing may be able to trim such errors, but this is an area for future refinement—a topic we return to in a later section.

VI. User Experience

The participants all took the phones with them every day, carrying the phones around with them wherever they went for the vast majority of the day. The application was found to be both reliable and stable overall, and everyone found it easy to use. Where group 2 had the chance to use it during most of their working day and therefore checked it and compared extensively (between 11 and 34 times a day), the other groups had busy days where they would mostly check their numbers and compare in the evening, therefore checking fewer times (between 1 and 20 times).

Participants reported that the application was fun to use and gave them good— and sometimes surprising—awareness of their activity level. Two participants (from group 2 and 3) reported it to be highly 'addictive', in particular the sharing aspect. Another participant repeatedly explained how it made him see how 'lazy' he was. Although only four of the nine participants reported doing more activity than usual in the interviews (and attributed it to the application's sharing functionality as well as more general competitiveness), the diaries show that the other participants were also more active compared to the initial three day 'base' diary. The short-term nature of this pilot study does not allow for observation or inferences to be made about the initial novelty value of the system. As we discuss further in the Conclusion and Future Research section, a longitudinal clinical trial will determine long-term use and effects.

A. *Individual Use and Motivation*

The participants described how they would enjoy checking how much walking and running activity they did during the day. Most of them checked their own minutes regularly and were astonished how they gained minutes during busy days. One woman from group 3 was surprised that she had accumulated 177 minutes one day, but when looking back though the diary, she realised that she had been busy commuting between two different work places (which involved walking to and from a bus and a ferry), as well as walking her dogs in the morning and evening. We were able to detect most of her activities in the data log, except for some of her transport that had a few small gaps of 30 seconds walking when she was in fact driving. This error, however, did not add more than seven minutes of walking to the whole day. This participant was busy and already highly active, and did not feel the application had made her change her activity level during the study.

One participant from Group 2 on the other hand, was very active that week in particular, and attributed this to the application. He explains how he increased his activity that week:

> [I]t probably encouraged me to go running Monday, Wednesday and Friday, because I always have the

> intention of going running at the beginning of the week. […] and I sort of set out Monday, okay right, I will take my stuff and I will go, you know, just Monday, Wednesday and Friday. [It also encouraged me to] just walk a couple of extra bus stops […]

He was very keen on increasing his activity level, and had tried to get into running three days a week for a while, without complete success. The weather had sometimes deterred him before, but with Shakra, he went out every planned day despite it being very rainy two of those days.

Although the participants seemed to be motivated from just the awareness of their activity, the effect was not unanticipated; often merely the knowledge that others can detect one's activity (either from a fill in diary or a tracking system) makes one more active. However, it was important to explore whether the use affected users' awareness and attitude towards moderate exercise. Behavioural change is a slow and often long-term process, but the necessary first steps have been taken here, in that awareness and motivation increased. Other issues affect motivation and awareness in return; therefore it should be related to social factors such as competition and collaboration—as the next section discusses.

B. *Shared Experience*

The groups did not only enjoy the increased awareness of their individual activity levels, they also enjoyed competing among themselves. Group 2 were quite determined in their competition, in particular one participant who would spend much of his working day walking around taking calls on his wireless headset, much more than he usually did. One of his group members explains:

> "… [W]e would be sitting in calls and he would be walking by [showing the phone to us]. Maybe there was a meeting round that side of the building (pointing), he would walk all around the building to get there (the building is donut shaped) … Me and Colin would sort of check more often to see. Ewan just rubbed it in front of our noses, how far he went".

This group enjoyed the competition despite a very different number of accumulated active minutes as figure 4 shows. Since the 'overachiever' described above had a wireless headset and was not confined to his desk, he could work while walking around—or walk while working. The other two group members were more confined to their desks during the day and only reached about half of his minutes every day. Where the first of these two said that he realised how 'lazy' he was. The participant second explained that he did not care that much, since he worked out at the gym about three times a week. He was quite content with his activity level, and did not see his ten-minute walk to and from work as 'exercise'. In this case there was more concern from the less active of the two, who was in the category that the application is most focused on, although he was constrained in changing this awareness into greater activity—at least during the trial.

Group 3 also started competing, with two women particularly competitive with each other. One wanted to beat her very active friend. For example, one evening when she came back from a run with 112 minutes, she saw her friend had 177 minutes of activity. In an attempt to catch her friend up, she asked her neighbour if she could take the latter's dogs for a walk. She therefore managed to get 137 minutes—not quite enough to beat her friend, but a respectable amount of exercise to say the least.

Group 1 did not compete much, but they did enjoy the fact that they could see each other's activity when they were apart. The oldest of the study participants and also a married couple, they mostly used the system to keep an eye on their own activity levels.

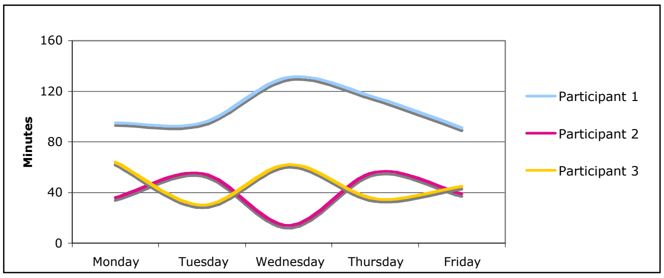

Figure 4: Group 2's total accumulated minutes per day

C. *Sharing the Fun*

One distinct difference between our application and mobile games is the designed purpose of promoting exercise versus promoting play. However, we found that the difference in use is not necessarily so distinct. Where other games have been shown to promote exercise, we found that playfulness can be a side effect of health-focused applications.

Participants had fun competing as described above and they did not only use it for teasing each other and as a conversation topic: some of them saw it as a game. One of the participants commented that his buddy "wanted to win so much. Before we could even get it to a certain level, he was flying", he said. 'Walking around' with the sole purpose of gaining active minutes was common behaviour among some participants, which not only shows their competitiveness but also how they wanted to 'play' the system. In essence, the application has game-like characteristics for those who like to play: the winner is the person who accumulates the most activity in a day.

VII. Conclusion

The development of Shakra is a first step towards creating a low cost physical activity monitor and health promotion application that is easily accessible to the general public. Shakra's real-time collaborative aspects and its lack of sensors beyond the mobile phone differentiate it from other research and products in this area. The initial reaction to Shakra during its pilot study was extremely encouraging; however some issues with accuracy, feedback, privacy and awareness were raised and must be addressed in future implementations.

All of the study participants responded positively towards the system and were tolerant of the momentary lags and jitters in activity classification (as discussed in Section V). Many of the participants were excited to see their own activity level, expressing higher motivation and displaying some increase in physical activity. We observed some of the same features that have been seen in more traditional collaboration in exercise to lead to more exercise being done, such as encouragement among ‘buddies’ and, in some cases, strong competition. The way in which the application was used varied between individuals and groups: it was used variously as a mutual awareness tool, a self-monitoring device and as a game. This highlights the need for a degree of ambiguity within the design of a health-promoting system that has a broad user demographic; enabling individuals and groups to use the system in such a way that suits and benefits them.

The technology appears to have been less precise in distinguishing between different types of activity than in the previous controlled experiments [2]. It is to be expected that accuracy might be reduced when used ‘in the wild’, yet it is our belief that the system can be improved to be more accurate. In the current system, for example, the training period might run for longer so as to account for more of the areas that users go to, or the neural network could be trained dynamically over the course of the application’s life.

One issue discovered was the decrease in accuracy when moving between disparate environments. We are currently experimenting with ways to improve this. One possible approach would be to use three different gated artificial neural networks, each of which has been calibrated for specific types of environment, i.e. rural, suburban and metropolitan. This would reduce activity-sensing errors associated with driving and walking between different classes of environment. Initial data analysis suggests that it should be possible to infer the type of environment that the user is currently located in (and hence the network to use) by looking, over a longer period of time, at the pattern of changes to the list of neighbouring cells.

As with any pilot study, there are limitations to the validity of any resultant claims made. It is not possible, for example, to claim that over a longer period of time the participants would remain enthused and continue to feel motivated by the system. What we do infer from the pilot study is that Shakra is usable and can initiate such positive responses, and suggest that with further development the system may prove to be an effective health-promotion tool. The final section discusses the future work planned in order to create such a tool.

A. Future Work

In addition to improving the accuracy of activity inference, the granularity of activity inference must be increased if we are to achieve our overall goal of implementing a system that can categorise the various levels of activity intensity. We believe that this will be possible as the accuracy improvements are made, and more rigorous neural network training is implemented. Instead of training a neural network to recognize walking at any speed, the network should be trained at the various intensity cut-off points, e.g. low intensity below 4 mph, moderate intensity above 4 mph. Similarly work will be done including other activities such as cycling, with equivalent distinction between low, moderate, and high intensity cycling. If any remaining types of activities or contributing factors to intensity are to be acknowledged by the system then we envision the need to utilise additional technology. We intend to explore new sensing and analysis techniques that can run on commodity phones, especially as they evolve to contain such previously exotic hardware as WiFi, GPS and, in phones such as the Nokia 3220, accelerometers.

As the focus of the system is to primarily encourage small changes in behavior, no attention has been made so far to the minimum recommended session length of 10 minutes. This could be easily introduced by a post-processing 10 minute rolling filter, earning users additional accreditation when a 10 minute session is completed. Another potential avenue of exploration is that of an adaptive system that evolves alongside a user’s activity pattern; the 10 minute session accreditation being introduced when the system detects substantial levels of intermittent activity throughout the day.

Once completed the system will be the subject of a clinical trial to determine the extent of any resulting changes in attitude, behaviour and health. We expect to use both qualitative evaluation techniques to assess these changes in objective terms, and qualitative evaluation techniques, to explore the detail of individual and social interaction around the system. In particular, we are interested in how people weave such technology into everyday life [4], and expect users to develop to tactics and strategies for use beyond our expectations, appropriating or even ‘hacking’ the technology to suit their own goals, desires and contexts. There is clearly great potential in technical explorations using highly accurate specialised assemblies of hardware and software, such as the multi–modal sensor board and iMote of [12], but our study illustrated the pragmatic advantages of a lightweight application running on a mobile phone with no such specialised sensors, and no cumbersome attachments, e.g. being strapped to the body. We suggest that a commodity platform will help such a health promoting application be more readily integrated into the lives of the wider population sooner rather than later.

Acknowledgments

We wish to thank all the participants for their time, and Viki Penpraze of the IBLS department of the University of Glasgow for her feedback on our initial thoughts and ideas. This work was funded by UK EPSRC (GR/N15986/01).

References

1. Adams, J. and White, M. Why Don’t Stage-Based Activity Promotion Interventions Work? *Health Education Research.* Vol. 20. No. 2, pp. 237-243, 2005.
2. Anderson, I and Muller, H. Context Awareness via GSM Signal Strength Fluctuation. In: the 4th International Conference on Pervasive Computing, Late breaking results, pages 27--31. Oesterreichische Computer Gesellschaft, May 2006.
3. Bandura, A. *Social Foundations of Thought and Action: A Social Cognitive Theory.* Prentice Hall., Englewood Cliffs, NJ, 1986.
4. Bell. M. et al., Interweaving Mobile Games with Everyday Life. *Proc. ACM CHI* 2006, pp. 417-426.

5. Consolvo, S., Everitt, K, Smith, I. and J. A. Landay: Design Requirements for Technologies that Encourage Physical Activity. To appear in *Proc. ACM CHI* 2006.
6. Department of Health. At least five a week: Evidence on the impact of physical activity and its relationship to health. *London: DoH*, 2004.
7. Duncan, S. C., Duncan, T. E., Strycker, L. A., Chaumeton, N. R. A Multilevel Approach to Youth Physical Activity Research. *Exercise &Sport Sciences Reviews.* Vol. 32. No.3. pp. 95-99, 2004.
8. Frederick-Recascino, C., M. Competition and intrinsic motivation in physical activity: a comparison of two groups. *Journal of Sport Behavior.* Vol. 26 Issue 3, pp 240-255, 2003.
9. Ipsos Mori, e-MORI Technology Tracker, December 2002 http://www.mori.com/technology/techtracker.shtml [last accessed 29th June 2006]
10. King, A.C., Friedman, R. Marcus, B., Castro, C., Forsyth, L., Napolitano, M. and Pinto, B. Harnessing Motivational Forces in the Promotion of Physical Activity: the Community Health Advice by Telephone (CHAT) Project. *Health Education Research.* Vol. 17, No. 5. pp 627-636, 2002.
11. Korp, P. Health on the Internet: Implications for Health Promotion. *Health Education Research.* Vol. 21. No. 1. pp. 78-86, 2006.
12. Lester, J., Choudhury, T. and Borriello, G. A Practical Approach to recognizing Physical Activites. *Proc. Pervasive 2006*. pp 1-16, 2006.
13. Licoppe, C. and Inada, Y. Emergent uses of a Multiplayer Location-Aware Mobile Game: the Interactional Consequences of Mediated Encounters. *Mobilities*. Vol. 1, No. 1. 2006.
14. Liden, C. B., Wolowicz, M., Stivoric, J., Teller, A., Vishnubhatla, S., Pelletier, R. and Farringdon, J. Accuracy and Reliability of the SenseWear Armband as an Energy Expenditure Assessment Device. *BodyMedia* 2002. http://www.bodybugg.com.
15. Mullen, E. and Markland, D. Variations in Self-Determination Across the Stages of Change for Exercise in Adults. Motivation and Emotion, Vol. 21. No. 4, pp 349-362, 1997.
16. Pate, R.R., Pratt, M., Blair, S.N., Haskell, W.L, Macera, C.A. and Bouchard, C. (1995). Physical activity and public health: a recommendation from the Centers for Disease Control and Prevention and the American College of Sports Medicine. *Journal of the American Medical Association*, 273, 402-407.
17. Prochaska, J. and DiClemente, C. Transtheoretical therapy: toward a more integrative model of change. *Journal of Consulting and Clinical Behaviour*, *51*. 390-395, 1982.
18. Siewiorek, D. P., Smailagic, A., Furukawa, J., Krause, A., Moraveji, N., Reiger, K., Shaffer, J. and Wong, F. L. SenSay: A Context-Aware Mobile Phone. In *Proceedings of 7th International Symposium on Wearable Computers, ISWC.* pp 248-249, 2003. IEEE Computer Society.
19. Sparling, P.B., Owen, N., Lambert, E.V. and Haskell, W. L. Promoting Physical Activity: The New Imperative for Public Health. *Health Education Research.* Vol. 15. No. 3. pp. 367-376, 2000.
20. Speck, B. J. and Harrell, J. S. Maintaining Regular Physical Activity in Women – Evidence to Date. *Journal of Cardiovascular Nursing.* Vol. 18, No. 4. pp. 282-291. 2003.
21. Speck, B. J. and Looney, S. W. Effects of a Minimal Intervention to Increase Physical Activity in Women – Daily Activity Records. *Nursing Research.* Vol. 50, No. 6, pp. 374-378, 2001.
22. Sporting Britain. Department for Media, Culture and Sport, UK, April 2004.
23. Tudor-Locke, C. Taking Steps Toward Increased Physical Activity: Using Pedometers to Measure and Motivate. *President's Council on Physical Fitness and Sports.* June 2002.
24. Van Sluijs, E., van Poppel, N.M., Twisk, J.W.R. Chin A Paw, Calfas, K. J. and van Mechelen, W. Effect of a Tailored Physical Activity Intervention Delivered in General Practice Settings: Results of a Randomized Controlled Trial. *American Journal of Public Health.* Vol. 95. No. 10. pp. 1825-1831, 2005.
25. Vandelanotte, C. and De Bourdeaudhuij, I. Acceptability and Feasibility of a Computer-Tailored Physical Activity Intervention Using Stages of Change: Project FAITH. *Health Education Research.* Vol. 18. No. 3. pp. 304-317, 2003.
26. Weiss, M., R. Motivating Kids in Physical Activity. *The President's Council on Physical Fitness and Sports Research Digest.* Series 3. No. 11, 2000.
27. World Health Organisation: Move for Health http://www.who.int/moveforhealth/en/ [last accessed 29th June 2006]