Shakra: Sharing and Motivating Awareness of Everyday Activity

Louise Barkhuus[1], Julie Maitland[1], Ian Anderson[2], Scott Sherwood[1], Malcolm Hall[1] and Matthew Chalmers[1]

[1]University of Glasgow
17 Lilybank Gardens
Glasgow, G12 8QQ
United Kingdom
[barkhuus, jules, sherwood, mh, matthew]@dcs.gla.ac.uk

[2]University of Bristol
Tyndall Avenue
Bristol BS8 1TH
United Kingdom
anderson@cs.bris.ac.uk

Abstract. The increasing health problems of the western world underline the need to encourage daily exercise and activity. This paper explores the use of an unaugmented everyday device—the mobile phone—to support increased socialization around exercise and activity. We describe an application that tracks the daily exercise activities of people carrying phones, using fluctuation in signal strength to estimate a user's movement. In evaluating the application that shared activity information amongst groups of friends, we found that awareness encouraged reflection on, and increased motivation for, daily activity. This application advances work on games and playful use of ubicomp technology, supporting enjoyable interactions that have secondary benefits. We describe some of the detail of the application's use, and discuss possibilities for future systems and studies.

1. Introduction

Where positioning technologies for mobile computing have long been a core part of ubiquitous computing research and applications, focusing on adapting these to relate to concerns such as privacy and context [5, 9, 19], similar technologies have to a lesser extent been explored for non-positioning awareness. To date the primary method of recognizing activities has been achieved via the introduction of additional sensor technology such as accelerometers [14, 19] or using location information such as GPS traces [17, 21, 22]. In this paper we explore a new application that uses off-the-shelve mobile phone technology to track and communicate users' exercise–related activity rather than geographical location. By detecting patterns in signal strength fluctuations and changes in the visibility of cells, the application can identify modes of activity such as being still, walking or traveling in a car. The application affords numerous uses, such as adapting phone functionality to the situation (for example switching to hands-free mode when driving), but the possibility that we explore in depth here is support of fitness and exercise.

Studies show that obesity is a growing problem in most of the western world, and while more exercise and higher levels of physical activity could counteract this trend and reduce obesity related illness such as diabetes and heart disease, there is little understanding of how to encourage this. Ubicomp technologies provide potential for not only supporting stationary exercise, e.g. on a machine in a gym, but also for motivating exercise such as walking and running outdoors through tracking and sharing details of movement [7, 25] and tailoring exercise programs [28]. Supporting moderate exercise as part of everyday life is particularly valuable in settings where people currently have low activity levels [24]. Not only does regular moderate activity improve health, it is also more manageable for people to incorporate into their everyday lives. We here focus on this type of increase in activity; although high-level exercise in collective stretches is in some cases more efficient for health improvement, the reality of many people's lifestyle means that moderate activity increase is start. Not only can such approach increase fitness level, motivating people to increase daily moderate activity can also make people more aware of their need for exercise.

For the purpose of motivating exercise and increase awareness thereof, we developed Shakra[1] an application that supports the tracking and sharing of everyday activity, on commonly available unmodified mobile phones. Shakra identifies patterns of movement from fluctuation in cell signal strength. A trained neural network is triggered by the rate of change in cell strength, and the change in the cells that are visible. In contrast to the use of pedometers, accelerometers and GPS, this application does *not* need special devices or need to be strapped onto the body. GPS-based applications also have limitations in that they rely on constant view of GPS satellites. Naturally, they rarely work indoors, but satellite 'shadows' appear frequently in city settings and also limit outdoor use. Shakra displays users' current activity, alongside a running total of their minutes of activity per-day and supports sharing this information amongst a group of friends. We here present a trial of the system amongst three groups of users; we discuss how the application encouraged reflection on users' levels of activity, as well as creating a playful application wrapped around a serious pursuit: activity awareness and health promotion.

In the next section we review current approaches to motivating exercise as well as reviewing other technical approaches to fitness tracking and motivation. Secondly, the technology and implementation of Shakra is presented. Finally, a user study evaluating the application is described, exploring the competition and game like nature of the application and attitudes amongst the participants towards motivation, awareness and exercise. We conclude by pointing to the potential utility of applications such as Shakra, and suggest improvements as part of future work on technologies supporting exercise and health.

[1] The application is named after the body's centers of spiritual energy, according to yoga philosophy.

2. 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 [23, 25]. 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], with levels of exercise actually falling across the western world.

Many approaches to increasing fitness propose an increase in *moderate* activity, such as brisk walking, in order to improve people's health [10]. 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. Health recommendations state that people should do moderate activity for at least 30 minutes on most days of the week [20]. This 30 minutes stipulation may be achieved incrementally, i.e. built up over a day instead of being done all in one exercise session, however much controversy exists in the literature [26]. For the purpose of our prototype we focus on incremented exercise, however, an adaptation of the application could easily measure 5-10-minute intervals of exercise. One important factor though, which is important to the design of the system, is that many people have difficulties making sure that their activity is in fact moderate and not just light, i.e. that they are in fact achieving the health benefits stated above [20]. It is therefore important for individuals to not only be aware of their overall amount of exercise but also its intensity.

2.1 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 [25]; 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 [7].

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 [16]. 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 tend to wear tighter clothes than men.

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 [3] 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” [15]. 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.

2.2 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 [29] and Social Cognitive Theory, based on the individual’s outcome expectancy and self-efficacy [23].

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 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 behavior, 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 than extrinsic motivation (external pressures or immediate rewards) increase the likelihood that the person will stick to a routine [18]. Examples of intervention using this approach include giving health advice over the phone, either by health professionals or via an automated service, and through an Internet service [12]. Studies showed that human interaction for example was successful in promoting increased physical activity among middle-aged and elderly when compared against a control group [11].

Other research has addressed social aspects of sharing information about activity and found that exercising together can also motivate individuals to do *more* activity; people increase their activity level as they engage in the light competition [27]. 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

[29]. It is evident that intrinsic motivation is influenced but not determined 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) [8]. The social cognitive theory addresses aspects of community, contrasting to the transtheoretical model although it focuses on internal values as main motivator to increase individuals' level of exercise. Interestingly enough, social factors such as poverty and neighborhood have been found to highly influence people's level of exercise [8]. Again, such categorisations abstract over individual cases, but it is reasonable to conclude that an individual's exercise is often 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 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.

3. The Shakra Application

Our key design goal for Shakra was that it could be carried around in a non-intrusive manner, and requiring little or no extra equipment for users. Where other activity measuring applications such as Locadio [13] and Sensay [22] relies on add-on technologies or technologies not implemented into the common consumer phone yet, such as accelerometers, GPS and WiFi, we aimed to develop an application for an actual off-the-shelve device: the mobile phone. In addition, 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 (e.g. 'walking') and travelling in a car, bus or train (collectively labeled 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 syncronise via the *Sync* menu option.

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 1a shows the *Compare 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 1b. 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 1c and 1d.

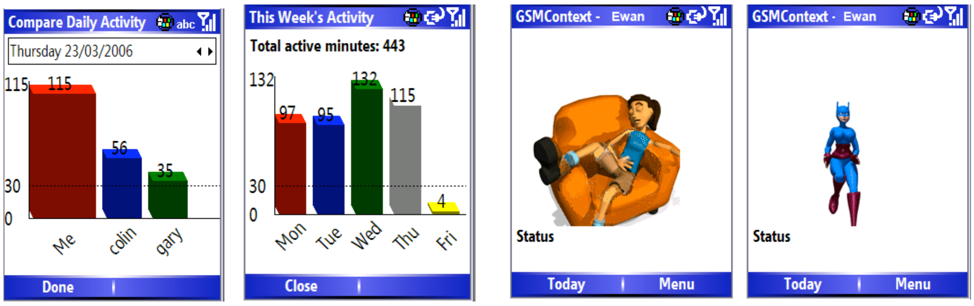

a) b) c) d)

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

3.1 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 [22] and the multi–modal sensor board of [14].

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.

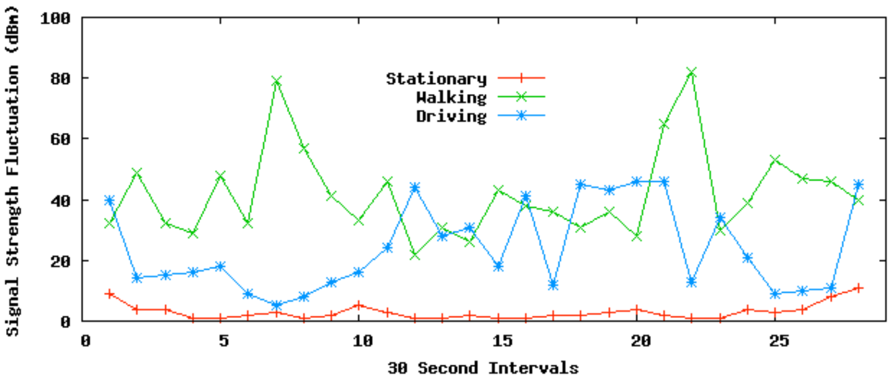

Figure 2: Distinguishable patterns of GSM signal strength fluctuation over successive 30second samples are used in identifying the activity levels *Stationary, Walking* and *Driving*.

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.

4. 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 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 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, cycling, and driving. In order to determine whether or not further personalization 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 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 days activity. This trained the system for the areas in which they usually go to throughout the course of a day. The training and the evaluation focused on weekdays, as weekend patterns of travel are often quite different from weekdays.

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.

4.1 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. This was one of the advantages of using phones in the study: it allowed communication about the study as well. The participants did not have to give us their own phone numbers, and calls on the phone were always related to the study. 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 the Shakra application.

The participants were recruited as groups of friends and/or coworkers who had daily interaction with each other and would enjoy sharing their exercise information. We also 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[2]. Table 1 provides an overview of the three groups.

	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

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 how the participants' experience of information sharing.

5. Reliability of Shakra in the Real World

Although previous tests (as reported in [2]) had shown highly accurate determination of activity, 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 labeling of activity. We chose three sample days

[2] 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.

for two different participants because their diary for that day was particularly comprehensive, i.e. six days in total. From the unfiltered data we analysed short stretches of 60 to 90 minutes with varied activity 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. The system would take three to five minutes stabilizing to the new transport mode such as walking, showing little half minute fragments of driving and walking until settling on walking. However since it would often have delay both before and after transportation, the measured minutes of exercise still ended up being real to the users. One more problematic finding was that running occasionally would register as driving. During one 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 minutes each day due to error. This was the maximum error we found from looking at participants' commutes.

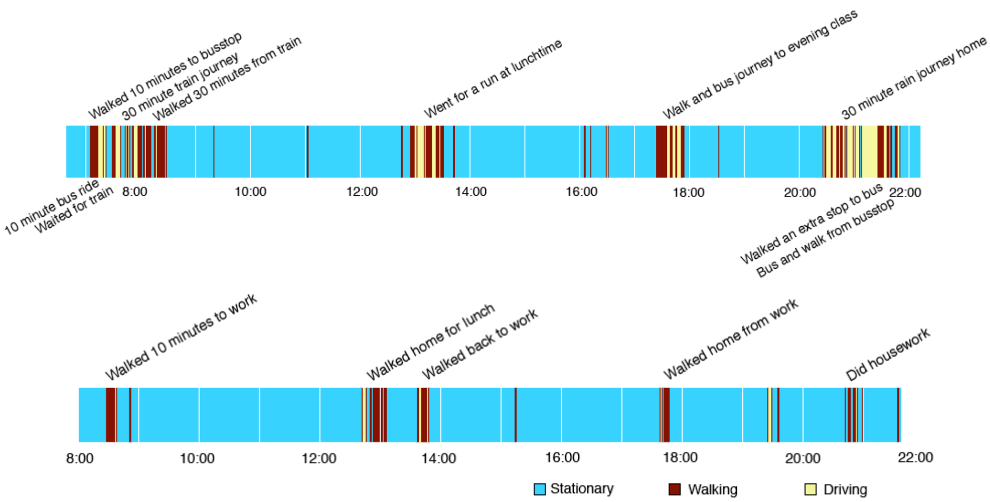

Figure 3: Example timelines of activity for two participant's days with colour showing the activity level and text showing the participants' diary annotations.

The participants confirmed the misinterpretations in the interviews. Three of them referred to two or three specific times where Shakra had shown the wrong icon, usually the driving icon, when walking fast or running. In the logs, it was easy to detect the two to nine false categorisations of driving activity per day. These were usually very small fragments, lasting a maximum of two minutes within longer stretches of walking activity or vehicle changes as mentioned before. Post-processing may be able to trim such errors, as such short drives are rare, but this is an area for future refinement—a topic we return to in a later section.

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 on 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 stoplights, etc. (When analyzing their timeline a few segments of still activity occurred within the walking stretch) Since the application is aimed towards increasing awareness rather then measuring physical exercise precisely, and offered usably accurate overall measures, we suggest that the small moment-by-moment lags and jitters in classification, e.g. two minute periods of activity, were not problematic.

6. 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 less 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. All participants met the daily 30-minute minimum moderate activity, except for four participants on one day each. The majority of the days, they accumulated much more than the recommendations, making the groups average at 53 (group 1), 64 (group 2) and 59 minutes (group 3); however, only three (one in each group) commuted by car, most of them therefore had long stretches of walk in connection with public transport commute. 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. Naturally, any kind of increased awareness of activity level is likely to initially result in people increasing their activity. As we discuss further in the Conclusion and Future Research section, there is a need for more long-term studies determining long-term use and effects.

6.1 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 minute one day, but when looking back though the diary, she realized 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 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 those three days a week for a while, without complete success. The weather had deterred him before sometimes, but with the 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 [28]. 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. Motivation and awareness are affected by other issues in return; therefore it should be related to social factors such as competition and collaboration – as the next section discusses.

7. 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.

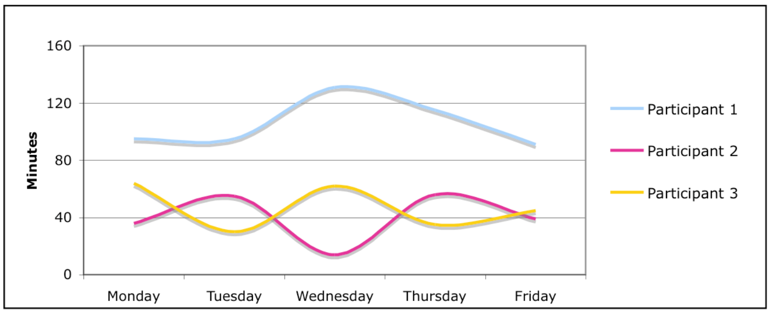

Figure 4: Group 2's total accumulated minutes per day

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. They mostly used the system to keep an eye on their own activity levels. Their lack of use of the application was probably due to initial hesitation about the system and that they were only two people sharing information rather than a larger group.

7.1 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 behavior among some participants, which not only shows their competitiveness but also how they wanted to 'play' the system. Playfulness was characteristic for the participants who had the flexibility to go out of their way to be more active. Group 1 for example had a very set daily life compared to the other groups and therefore mostly used it to monitor their own activity. Most of the participants in the other groups had time and flexibility to actually go out of their way to increase their activity level. In essence the application have game-like characteristics where input is moderate activity. One wins the game by doing more of such activity than others. Where traditional console-based games are sometimes criticised as promoting physical inactivity, we suggest that Shakra demonstrates that one can create a hybrid that is both a game and a health-promoting application.

7.2 Shakra as a Means of Collaboration

Finally, it was interesting to discover how the participants used the application to infer related information about their peer's whereabouts. On days where they did not see each other, two participants from group 3 said that it was nice to be able to see the other's numbers go up even at night. Then they knew that the other was still out and about.

A participant from group 2 used the activity sharing to infer whether one of the other participants had been to lunch or not. Depending on the number of active minutes he said he could see if the friend had walked his ten minutes to lunch and back. This assisted in their collaboration at work, since at least one of them had to man their work station all the time, i.e. they could not go to lunch together but would take turns.

Since the participants were only sharing with friends and co-workers that they chose to participate with, strong privacy within groups was not considered essential in design terms. However, we enquired whether participants felt uncomfortable with the level of information being shared. As one participant from group 2 noted: "The only thing was that I didn't get home 'til about half past twelve last night. So I clicked into the next day. […] They know what time I was getting in at [laughing]". He was not deeply uncomfortable with this, since he did not see this as highly sensitive information in his mind. However, it was clear that he had realised how the activity tracking could provide his friends with other, inferred information other than 'pure' or objective activity levels.

Such inference that does not give precise details is very useful to a tight group. None other than the friends of the user would know what a certain amount of activity mean and there is still a great margin of uncertainty, since the activity is not geo-referenced. Collaboration seemed to be an unforeseen benefit of the application for several of the participants and we therefore hope to explore this in future versions of the application as well.

8. Conclusions and Future Research

We set out to explore how the everyday technology of GSM phones could be used to make people more aware of their activity levels. We also set out to determine some of the detail of how our prototype application is used in a real environment. While it is clear that many of the participants felt more motivated, and were excited to see their own activity level in comparison to others, it is too early to claim that this approach is better in promoting awareness than other methods of recording and displaying 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. We also observed the use of time at work, sometimes as a resource and sometimes as a constraint on use—as was observed in studies of long–term mobile games. Although a preliminary study, we suggest that this study confirms the way that larger spatial and temporal scale of ubiquitous computing systems' use shows a wider view of context being taken into account. For example the competition within group 2 showed variation in use that was not only due to personal differences in attitude, but also technological differences, such as a wireless headset enabling walking while handling calls, and organisational differences—the two who sat watching their group colleague triumphantly walk by were also held back by the demands of their work.

The technology appears to have been less precise in distinguishing between different types of activity than in the previous controlled experiments of Shakra [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. Another change we intend to implement is to improve the accuracy of sensing activity when moving between disparate environments, in terms of cellular network infrastructures. One possible approach would be to use three different 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 by looking, over a longer period of time, at the pattern of changes to the list of neighbouring cells.

In this system the key distinction is between activity and non-activity such as driving and sitting still. Despite misidentifying walking as driving during a few 30-second time periods, the overall accumulation of active minutes was good when compared to the participant's diary and for the participants' use on a daily basis. Our system gives a simple, although still not entirely exact, indication of the carrier's general activity. Rather as has been noted by other researchers with regard to positioning systems of variable accuracy [4, 6], it is often the case that relatively inaccurate systems can still be accurate enough to support worthwhile individual and social interaction.

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 [14], 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 an application be well–integrated into the lives of the wider population sooner rather than later, as so many people already have and use mobile phones. In our own future work, 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.

Overall, in this paper we have shown that users became more aware of their activity level and enjoyed sharing activity information. They were also able to infer situational information about each other from activity information, without feeling intruded upon. As much as the technology in itself, participants' increased mutual awareness was key to their motivation. Naturally, this type of pilot study may be prone to novelty effects and says little about long-term effects. Nevertheless, we suggest that Shakra serves as an example of the way that understanding and designing with regard for the everyday social, organisational and technological context in which we live is a good way to aid motivation and continuation of healthy activity.

Acknowledgements

This research was funded by UK EPSRC (GR/N15986/01). We wish to thank all the participants for their time.

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. *Pervasive Computing,* Late Breaking Results. 2006.
3. Bell. M. et al., Interweaving Mobile Games with Everyday Life. To appear in *Proc. ACM CHI* 2006. (Also available at http://www.dcs.gla.ac.uk/~matthew/papers/CHIYoshiWithCopyright.pdf)
4. Benford, S. et al., The Error of Our Ways: The Experience of Self-Reported Position in a Location-Based Game, *Proc. Ubicomp 2004*, pp. 70–87.

5. Beresford, A. R. and Stajano, F. Location Privacy in Pervasive Computing. *Pervasive Computing*. January – March 2003. IEEE CS and IEEE Communications Society.

6. Brown, B. MacColl, I. Chalmers, M. Galani, A. Randell, C. Steed, A. Lessons from the Lighthouse: Collaboration in a Shared Mixed Reality System In *Proc. ACM CHI 2003*, pp. 577-584.

7. Consolvo, S., Everitt, K, Smith, I. and J. A. Landay: Design Requirements for Technologies that Encourage Physical Activity. In *Proc. of ACM CHI* 2006, pp. 457-466.

8. 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.

9. Iachello, G., Smith, I., Consolvo, S., Abowd, G.D., Hughes, J., Howard, J., Potter, F., Scott, J., Sohn, T., Hightower, J. and LaMarca A. Control, Deception, and Communication: Evaluating the Deployment of a Location-Enhanced Messaging Service. *Proc. UbiComp, 2005*. pp 213-231, 2005.

10. Johnson, M. F., and Sallis, J. F. and Hovell, M. F. Self-Report Assessment of Walking: Effects of Aided Recall Instructions and Item Order. *Measurement in Physical Education and Exercise Science.* Vol. 4. No. 3, pp. 141-155, 2000.

11. 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.

12. Korp, P. Health on the Internet: Implications for Health Promotion. *Health Education Research.* Vol. 21. No. 1. pp. 78-86, 2006.

13. Krumm, J. and Horvitz, E. LOCADIO: Inferring motion and location from wi-fi signal strengths. In *MobiQuitous*, pp.4-13, 2004.

14. Lester, J., Choudhury, T., Boriello, G., Consolvo, S., Landay, J., Everitt, K. and Smith, I. Sensing and Modeling Activities to Support Physical Fitness. Workshop paper in *UbiComp '05 Workshop: Monitoring, Measuring and Motivating Exercise: Ubiquitous Computing to Support Physical Fitness.* Tokyo, Japan, 2005.

15. 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.

16. 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.

17. Lukowicz, P., Ward, J. A., Junker, H., Tröster, G., Atrash, A. and Starner. T. Recognizing workshop activity using body worn microphones and accelerometers. In *Pervasive*, LNCS 3001, 2004.

18. 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.

19. Myles, G., Friday, A. and Davies, N. Preserving Privacy in Environments with Location-Based Applications. *IEEE Pervasive Computing*, Vol. 2. No. 1, pp. 56-64, 2003.

20. 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.

21. Patterson, D., Liao, L., Fox, D. and Kautz, H. Inferring highlevel behavior from low-level sensors, In *Proc. of UbiComp*, 2003.

22. 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.

23. 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.

24. 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.

25. 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.

26. Sporting Britain. Department for Media, Culture and Sport, UK, April 2004.

27. 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.

28. 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.

29. 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.