Lessons Learned from a Personal Sensing Architecture

Mik Lamming, Denis Bohm[1], Robert N. Mayo
Consumer Applications and Systems Laboratory
HP Laboratories Palo Alto
HPL-2005-150
August 18, 2005

minders, SPEC, ubiquitous computing, pervasive computing, sensing human behavior

For the past two years we've been using a personal sensing system we call SPEC. Because it has different goals it has a different design philosophy, and thus its architecture is quite different from the more common environmental sensing system. It was designed to monitor a user's activities for their own purposes, in order to support everyday information processing activities. In other words, it is not a surveillance system. It uses collections of personal, private sensors, called SPECs, which can be worn, or deployed where each individual user needs them. This P2P system is designed on the principle that for widespread adoption the users must feel confident that whatever information they sense about themselves is kept under their own control, and not some central agency – however apparently benign. This places it in sharp distinction to RFID technology for example. After two years experience building applications with our system, we are convinced that this approach has a lot to offer, and we are embarking on a program to build an improved version, and take this opportunity to review the consequence of our design philosophy, and architecture, and identify areas for further research.

* Internal Accession Date Only
[1] Firefly Design, Los Altos, CA, USA
Approved for External Publication

Lessons Learned from a Personal Sensing Architecture

Mik Lamming[1], Denis Bohm[2], Robert N. Mayo[1]

[1]Hewlett-Packard Labs
Palo Alto, CA, USA

[2]Firefly Design
Los Altos, CA, USA

Abstract

For the past two years we've been using a personal sensing system we call SPEC. Because it has different goals it has a different design philosophy, and thus its architecture is quite different from the more common environmental sensing system. It was designed to monitor a user's activities for their own purposes, in order to support everyday information processing activities. In other words, it is not a surveillance system. It uses collections of personal, private sensors, called SPECs, which can be worn, or deployed where each individual user needs them. This P2P system is designed on the principle that for widespread adoption the users must feel confident that whatever information they sense about themselves is kept under their own control, and not some central agency – however apparently benign. This places it in sharp distinction to RFID technology for example. After two years experience building applications with our system, we are convinced that this approach has a lot to offer, and we are embarking on a program to build an improved version, and take this opportunity to review the consequence of our design philosophy, and architecture, and identify areas for further research.

Introduction

In 2003 at UbiComp[1] we described an architecture for Personal Sensing using tiny personal computers called SPECs, emphasizing a proof-of-concept reminding application we had developed to explore our ideas. Over the past two years we have built two more applications to get a feel for other aspects of this novel architecture, and now are in a position to report our experiences.

Our SPEC system was designed with the following philosophies in mind:

1. It is designed to **sense the user's activities, for the user alone** – what we called the "Sense me for me" philosophy. In some respects it is similar in concept to a broad range of personal sensing devices, e.g. the wearable heart rate monitor used by athletes. For example: SPEC is wearable; designed to monitor various human

activities in the background, continuously and for very long periods, and has a very simple user interface.
2. What makes it different from a simple wearable device is that the sensing system can be **disaggregated**, each individual sensor being replicated and located where it can do the best job of providing 24x7 sensing.
3. **Sensors may be shared** amongst many simultaneous users, and/or applications.
4. Nevertheless the system can ensure that the **user's data will remain private to that user**.
5. Finally, a complete ensemble of sensors can have a **value proposition that makes sense for a single user application**.

The essential strength of personal sensing systems is that they can remain vigilant while our own attention lapses; they can immediately make sense of the data they perceive when it might be beyond our own capabilities or patience; and they can store and recall sensor data many minutes or years after it has become inaccessible to our own fallible, or failing memory. We saw a huge market for personal sensing applications ranging from things as simple as reminding you of a person's name just before you met them again; understanding our own patterns of behavior and how they affect our fitness, or wellness; through to home care of the aging baby boomers. However, users are wary of such monitoring systems, justifiably worrying that they will have little control over the data collected. For this reason we included user privacy as a requirement of our system.

In many ways, the "Sense me for me" structure is the opposite of an RFID system. In an RFID system, sensors in the environment monitor objects flowing through a supply chain. Location and ownership of products changes from time-to-time, and each owner is interested in, or at least indifferent to, participation in an efficient tracking system. RFIDs are efficient and effective for this purpose. Once a product is purchased, however, ownership transfers to a person that mostly likely wants tracking limited to his personal applications, or perhaps disabled entirely.

If ubiquitous infrastructure exists to track RFID tags, this makes it difficult to employ RFID technology for private applications. Carrying an RFID tag for one application, such as allowing the person to track his own movements within his residence, compromises his privacy if he encounters the readers of another RFID application. For instance, our home user would be detected when he enters a retail store that uses RFID tracking technology, even though the user only intended to track his own movements within his own house.

In contrast, our system allows individual users to monitor themselves, and the environment surrounding them, recording the data they collect in their personal devices for their personal use. This dichotomy between the environment surveilling you versus you surveilling the environment has been written about extensively in the context of video surveillance; see for example the writings of Steve Mann [2]. In our experience, the privacy of personal information is important to users; if privacy is not guaranteed, users are much less likely to find value in the system and may actively oppose its adoption [3].

Two years on we are sufficiently convinced of the utility of this architectural approach to personal sensing that we are contemplating a new design that builds upon what we learnt from the first simple prototype. Here we present what we have learnt.

SPEC Operation

Although our architecture is described in [1] a recap is appropriate:

SPECs are tiny autonomous sensors that can be attached to people, places or things. Each SPEC broadcasts a unique 32-bit identifier (ID32) every 2 seconds (to conserve power this interval is increased automatically when the set of SPECs in proximity isn't changing). They also listen continuously for the ID32 broadcast from nearby SPECs. When a new SPEC is sighted, a sighting record is created, time stamped with the start time, and stored in a history. Each record describes an interval during which a particular ID32 was repeatedly sighted. If sightings cease for more than a specified interval (2 minutes in the current system), then the sighting record is closed. As we shall explain shortly, the sighting history can be analyzed locally to see if the user should be alerted to any noteworthy events, or it can be uploaded, via a SPEC portal, to a trusted archive for analysis or storage. Portals also provide access to a time service enabling mobile SPECs to set their real-time clocks, and a means to download patterns to support real-time applications.

A very simple sighting history pattern recognizer is used within SPECs to detect noteworthy situations. The pattern recognizer uses a byte code form that is downloaded to the SPEC from a portal. Once downloaded the pattern recognizer runs independently within the SPEC. The pattern language is declarative and consists mainly of time interval and ID set operations. Functions are composed to create reminder expressions. If the expression result is true then the reminder is considered active, and an alert message may be issued.

We have given high priority to small size and battery life and in consequence, sacrificed communications and computational power, storage and user interface capability. We aspired to achieve power budgets in the prototypes that would allow small field trials of about a week to be completed without a battery change, and in practice exceeded that by a significant margin.

Primary Hypotheses

We've built a system around a new architecture, and have two years of experience with it. We now turn our attention to evaluating it. To be specific, we wish to evaluate these primary hypotheses:

1. The "sense me for me" approach allows for easy and incremental deployment and configuration, along with sharing of sensors among users.
2. Our system can "sense key aspects of human behavior".
3. The "users control their own data", so that they are more likely to adopt the system.

Sense Me for Me

We had considerable experience building activity sensing systems using a centralized architecture: environmental sensors connected to a central analyzer [4,5,6]. We wanted to see if our "sense me for me" approach provided the benefits we sought. Using our new architecture we built three example applications: the "Kyle" reminding system [4], an automatic diary generator similar to the Pepys system, and an Alzheimer home care simulator, from which the following results are obtained.

All of the applications produced highly informative, sometimes surprising, and often very encouraging results. The Kyle reminder system worked particularly well and demonstrated to our satisfaction that with just a small number of SPECs very useful systems could be deployed beyond our lab., and in the world at large. We placed SPECs in our houses, and cars, in our cubicles at work, and with permission, in public areas like our local Starbucks, and on things we wanted to keep track of, like laptops, and book bags. We also attached them to things which when moved, were strong indicators of certain activities – like a gym bag, and the family dog (to resolve disputes about who's turn it was to walk Fido). Most importantly we discovered that even with incomplete data useful work could be done if the application was designed to take that into account.

Deploying SPECs where they are needed

The "sense me for me" philosophy is pragmatic, and was born out of the realization that deploying a pervasive system to support everyday information processing tasks can be prohibitively expensive for small research projects like ours. Our earlier attempts to build personal sensing systems had required, and assumed that some agency or enterprise would install the necessary infrastructure throughout an environment. The infrastructure was expensive to install and necessarily had to be shared, and thus was only put in places where the benefit to the community was clear. For individuals, this meant that only part of their life could be covered (typically work – which is only about 30% of the day), and the rest of their life (home, sports, driving, etc – which is the other 70% of the day) was left unsupported. For everyday activities this is a devastating state of affairs. Memory support for example is of little use if either the ability to automatically capture episodes, or to retrieve previously captured and stored episodes is not available in locations we frequent. Indeed if the system is only active 30% of the time, then it will only be useful about 30% x 30% of the time = 9%. Our goal with SPEC was to create an architecture that could be deployed by the user wherever he, or she needed it as individuals – typically their office, home, car, gym, club or other public, or private venue, and which could grow incrementally, (and hopefully virally), and be shared as needs arose.

SPECs are extremely easy to deploy. In our proof of concept prototype a ten year-old student was able to attach SPECs to his belongings (book bag, scooter etc.), and places (home, school) he visited, and make a wearable for himself. As additional needs were spotted, he added extra sensors to places and things. Since the sensors were battery powered, and wireless they could simply be stood on shelves, or picture

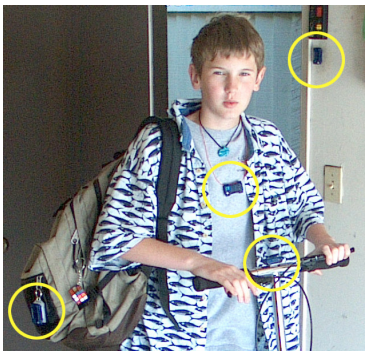

rails, or attached to things with glue, or Velcro. In another instance, several people in our lab set up the system in their own homes, and encountered few problems.

However, there was one unsolved problem related to deployment, and that was configuration: assigning functions and names, to the SPECs deployed. For instance, if SPEC 452 is placed on a shelf in the bedroom, we'd like to know that fact so that at the very least we can turn the ID into a user-readable description. In a centralized system this configuration would typically be performed by a knowledgeable system administrator, but in our user centric system the user had to perform this process manually themselves. This is a burden on the user that we would like to minimize or eliminate.

Although we provided a configuration wizard, we aspired to solve this problem algorithmically. Our idea was to analyze patterns of movement for a few days, and from that determine the role each SPEC was playing. For instance, a SPEC in the bedroom would typically see somebody stay in that room for a long period each night, with occasional visits during the daytime. We applied sophisticated pattern recognition technology to this problem, with small initial success, but not the quality of result that we were hoping for. A human carefully inspecting the activity logs could generate ad-hoc rules that did a much better job of determining the roles of the SPECs. Based upon this information, we consider this to still be an open problem that may have better solutions to the ones we tried.

We found that our desired synergy-effect, or viral growth effect, was in fact real. During deployment of the SPECs in the homes and offices of several lab members, each of us encountered other people's SPECs when we visited them, their offices, cars, etc. In addition, since many of us visited the local espresso stand, we had enough traffic to justify placing a SPEC there. No single user found such a placement compelling, but since there were a number of us we found it was worth the cost.

Keeping my data with me.

One of the most powerful wins of this architecture is that data accumulated about an individual user is not held in some centralized "Big-Brother" system, but within the system that that user bought and owns. In the complete implementation of our privacy architecture, these components are the SPECs themselves and a user's trusted archive, which may reside on his PC. However, portals and other networking components are not trusted. Consequently there is no requirement for the user to place trust in infrastructure installed by possibly untrustworthy third parties, or apparently benevolent commercial or government agencies that may in fact have unstated reasons for accumulating data about you. Of course the user may, if he wishes, loan the data to a trusted analysis service if the he deems it trustworthy.

We will discuss issues concerning interpreting the data, and privacy in the next two sections. The main issue we need to report here was that users sometimes ran out of memory in their wearable to store all their observations. Although the SPEC had an excruciatingly small amount of memory, this happened rarely. But when it did it was catastrophic because it created significant gaps in what was supposed to be a comprehensive record of events. It is difficult to have confidence in a technology designed, among other things, to provide support for human memory, if it gets bad attacks of amnesia itself. Even regular uploads via IR to a trusted archive did not entirely solve this problem. Fortunately the problem can be solved almost entirely by advances in technology: the processor chips we now use have very much more memory, and we now use RF to perform uploads which is much more reliable, faster and does not require line-of-site to a portal, and so can happen much more frequently. We anticipate that many weeks of observations can now be stored before an upload is required, and the upload will now take just a few seconds whenever the wearable is within a hundred feet or so of any portal.

One obvious consequence of adopting an architectural philosophy where a device on the user senses the environment, instead of the environment sensing the user is of course the need to carry around an active device. Coming up with an acceptably small device that could be worn 'unconsciously' all the time was a challenge upon which we decided to compromise for our first iteration. The big question we faced was where to wear them. Our solution was to leave that ambiguous, but make the devices small so users could fashion their own solution. The most common solutions employed by users were to attach them to their bodies like a pendent, or pin them to their shirts. None of these were completely acceptable, as they were socially unacceptable - they drew undesired attention to the person. Our ten-year old user was successful in packaging his SPEC to look cool, but the rest of us did not have much success.

We expect the form-factor problem to be a real concern in many applications. Social acceptability is needed for any consumer application such as reminding. If we apply this technology to medical applications such as the caretaking of Alzheimer's patients another set of considerations comes up. In addition to being socially acceptable, the device must function in what amounts to a hostile environment, with water being present, vigorous banging of the device, and perhaps occasional attempts to damage it or cover it up.

The very smallness of the device made it difficult to incorporate much of a user-interface directly. The SPEC user interface was minimalist consisting of one button and one green LED. At one extreme we considered having no UI at all, but we anticipated it would be hard to check on simple system errors like flat batteries, or even to switch it off between experiments. The single LED was barely adequate for our reminder application since it was not visible in bright sunlight and could not draw the user's attention. If our reminder application had needed to distinguish more than one condition then we would have been forced to use some non-intuitive method, such as blinking the LED in some characteristic pattern – clearly not very user friendly.

From our simple experiment we confirmed two rather obvious things: that we needed a more reliable way of attracting the user's attention; and that on occasions we needed to communicate more information than a tiny platform could reasonably

support. We also realized that as SPECs became even smaller, it would be even harder to incorporate user interface elements into the design, and so we anticipated the need for a recruitable user-interface (RUI) capability. Future SPECs will be augmented with a means to detect when there are devices with suitable user interfaces nearby, and a standard mechanism for recruiting them to attract the attention of the user. A suitable UI might be on a watch, or a cell phone. As the dialog complexity escalates we expect that the user's focus of attention may need to be moved to ever more powerful devices like a TV, or a wall-sized display and in this case there will be complex issues involving privacy.

Sense Key Aspects of Human Behavior

Using a centralized architecture, we had had some success constructing human recognizable descriptions of human activity from raw sensor data. We wanted to see if our new architecture could give comparable results.

Ability to generate human recognizable descriptions

The Pepys automatic biographer, the Forget-me-not system for recalling everyday events, and Sellen's context-based reminding system were benchmarks for the kind of applications we wanted to support. Each of these systems, tried to generate a model of the user's activities that mimics the user's episodic memory [7,8,9] for the same events. Important features of this kind of structure are, the sequence of episodes, and the context of the episode: where it occurred, and who or what else was present, e.g. "I had my purse when I met with Sue at Starbucks, but not when I got back into the car". We wanted to be sure that our architecture could capture and describe these events in a similar manner. Notice that the important feature here is to try and report events in the same way that the user might themselves report them. Two people in close proximity, but separated by a wall, would not report that event as a meeting. We coined the term social proximity to describe the effect we were seeking.

As an aside, note that activity diaries can also provide a valuable context for interpreting other sensor data in clinical trials and healthcare research – both of which we are exploring. DataEdge reports that "25% of all clinical trials for pharmaceutical products use a patient diary." However, the reliability of patient-generated data is often suspect. A 2004 Drug Information Journal article[12] states that patients are "apt to forget to complete the diary altogether, and can record invalid and inappropriate responses even when the diary is designed with check boxes or instructions for recording responses within specific parameters." Moreover, a 2003 Lancet article[13] asserts that "patients faked entries in medication logs" and that one study indicates that "actual compliance was only 11% although participants returned diary cards corresponding to 90%." Clearly automatically constructed activity diaries may offer considerable value, and we are exploring the extent to which SPECs might be valuable.

In designing SPEC we had decided to concentrate on gathering social proximity data, which we would approximate by sensing the presence of nearby SPECs.

Because we wanted to sense when users were close to each other, or near important possessions, or in particular rooms we made SPECs wearable, which forced them to be small and battery powered. The need to have small batteries that lasted for weeks or months was perhaps the most important decision we made, and it impacted all corners of our design. For simplicity we decided to make all SPECs, regardless of their intended role, behave in an identical manner emitting a unique identifier every two seconds, and listening out for others all the time. To simplify development we decided to allow SPECs to upload their data to a PC for archiving, and analysis, though in principle we intended for all analysis to be done on board the SPEC, only the results being communicated to the user (perhaps via a portal, and suitable recruitable user interface) when needed.

Our decision to use IR for determining social context was a tradeoff, and performed about as well as expected. However, the IR emitter must be worn in a conspicuous location on the user – which is socially undesirable in many cases. As we expected, IR was somewhat unreliable. The primary reasons for this were: that the sensors occasionally became obscured by clothing or the wearer's body position, and that the sensors were washed out by bright sunlight, such as when riding in an convertible automobile with the top down.

In a centralized system, all of the sensor data can be analyzed to produce a picture of activity for each person. In the new architecture, often for our self-imposed reasons of privacy, each sensor would have to perform its own analysis based only upon the information it could sense directly, or obtain from other sensors nearby. We wanted to understand what impact this limited view of the world might have upon our ability to construct accurate descriptions and learned that good results could indeed be obtained (Figure 1.) We are confident that this kind of analysis can be done on the SPEC without requiring upload to any external device or archive.

23:04:01 Bedroom for 7 hours, and 48 minutes
06:52:01 Bathroom 12 minutes, and 40 seconds -- taking a *shower*
07:04:54 Car for 56 minutes, and 59 seconds, -- *out for breakfast*
including excursions elsewhere totaling -- *in company with Laura, left car*
38 minutes, and 24 seconds with Laura -- *and went to restaurant*
08:17:54 Dining Room for less than a second -- *return home, to collect bag*
08:17:58 Home-Desk for 2 minutes, and 15 seconds
08:20:35 Gym-Bag for 6 seconds -- *carry bag (which has a tag on it) to car*
08:21:10 Car for 15 minutes, and 25 seconds -- *drive to work*
08:40:53 Gym-Bag for 36 seconds -- *carry bag into work to desk*
08:42:05 Work-Desk for 9 minutes, and 57 seconds
08:52:02 Elsewhere for 8 minutes, and 6 seconds -- *no tags in sight*
7 minutes, and 6 seconds with Denis -- *encounter with Denis*
…
15:37:58 Work-Desk for 2 hours, and 34 minutes
18:15:32 Gym-Bag for 2 minutes, and 8 seconds -- *session at work gym*
19:13:38 Gym-Bag for 3 minutes, and 1 second
…
19:18:48 Car for 9 minutes, and 40 seconds -- *stopped on way home to buy milk*
19:31:13 Car for 6 minutes, and 52 seconds
…
20:02:55 Sitting Room for 10 minutes, and 51 seconds -- *Watching TV*
…
21:39:52 Bathroom for 31 seconds

Figure 1. Selections from a personal diary generated automatically by SPEC system deployed at home, work and in car, and worn by family members. Comments in italics added manually for purposes of this paper.

We made a decision to reduce each SPEC's ID transmit interval according to it's expectation of activity. In other words, if it didn't expect to see any changes, because the set of beacons it was observing hadn't changed for a while, it would transmit its own ID less often. This decision was made to reduce the average power consumption. For example, a SPEC in an empty room would reduce it's transmission interval when nobody was around, thereby extending battery lifetime. One problem with reducing the ID transmit interval is that it increases the time it takes to recognize a change. Occasionally there were situations when the user made rapid trips between rooms, that observations were missed because the devices involved had backed-off their ID transmit interval too far, and thus rapid movement was interpreted as no moment at all. This caused some missed sightings in our system, but gave us longer battery life. We anticipate that this problem can be alleviated by the addition of accelerometers to adjust the ID transmit interval whenever the user moves significantly.

One difficulty we encountered was that when SPECs died – usually of a flat battery, this often went unnoticed for a while. We devised a scheme where the SPEC would flash it's LED when its battery had only a couple of days life left, but this was not good enough. When SPECs were dead, and the analyzer didn't know, it obviously made bogus assumptions, and produced an inaccurate log. We intend to add a service that will model the battery life of each SPEC and provide a reliable way to notify the user when it is time to change a battery.

To generate a proper interpretation of any SPEC's observation log, a SPEC has to know what role each other SPEC it observes is playing. For the Kyle reminder application the roles were known to Kyle's wearable, and that application worked very well. Later on we deployed several personal systems in the team member's houses, and around the lab. and tried to generate activity diaries. It was not unusual to encounter a SPEC sighting from a completely different system for which you had no role recorded. Obviously this made interpretation difficult, if not impossible. We are adding the ability for friendly SPECs to share this type of information under the consent of the users.

Efficiency

Central to this style of system is the need to be able to efficiently translate the sensor data into human recognizable/meaningful descriptions of events. In our case we were concerned about the impact on our power budget.

One advantage of using IR proximity was that its physical properties are a good proxy for the social proximity characteristics we were looking for. We considered using RF signal strength as a proxy, but discovered that at the frequencies we could accommodate, the human body could cause wild fluctuations. The fact that RF went through walls indicated that a building survey would be required to construct a room database, and SPECs would need to perform some triangulation before consulting the database to figure out if they were in the same room. Communication is relatively power hungry. We rejected GPS for a number of reasons: it exceeded our size and power budget, didn't work inside buildings, and would have the same communication overheads as RF signal strength. IR had this convenient property that it didn't go

through walls and ceilings, but would normally bounce off them, and naturally solved the problem of social proximity without heavy communication, or processing.

To map the IR proximity observations into human recognizable events we devised a pattern matching language which was standard on all SPECs. SPECs could be tailored to a particular role by downloading particular patterns they were required to look for.

There were many facets to our use of patterns, and most of these worked well. The simple pattern language we developed was sufficiently powerful to implement reminding applications. One aspect we could improve on, however, was the sensor data that was used as input to the patterns. We decided to merge all identical sensor observations that occurred over a 2-minute interval into a single observation spanning that interval. This stream of merged observations was then fed into he pattern recognizer. We have since discovered applications where 2- minute is not the best choice, and we believe that this parameter needs to be adjustable.

The downloadable nature of our patterns turned out to be a good decision. By downloading the patterns, the devices could be reconfigured for new applications or to try out new variants of the patterns while developing the applications. Patterns were written and then compiled into byte codes, and this had two main benefits: the compilation step allows for optimization, similar in spirit to other pattern compilation systems [14], and the interpreted nature of the byte codes ensured that only pattern-recognition operations would be executed by the downloaded code. In future work we would like to extend the patterns to generate events that can be used by other services, as well as exploring the use of patterns provided by friendly services that can share high level information without compromising the privacy of low level observations. For example, a sport equipment service might provide patterns to analyze activity data from SPECs embedded in running shoes, to recommend what sort of running shoes the user might want to buy next time. We believe this can be done without the user sharing any details of their own sport activities.

Users Control Their Own Data

As mentioned previously, we believe the privacy aspects of personal sensing systems are critical. People use systems because they gain overall value from them. No matter how valuable a personal sensing application is to the user, a lack of control over their information will be a detriment that subtracts from the overall value. Our experience, and that of others [15], shows these to be critical concerns, and if not actively addressed can often prevent adoption. Not providing privacy effectively deprives potential users of the benefit they could have received had privacy been addressed.

During deployment in our lab we verified, again, privacy concerns. Our peer-to-peer design was partially chosen to ameliorate privacy issues. While effective to some degree, we still got resistance from some of our users to the collection of data. This is true even though we are a small and friendly group, and the data was in their possession at all times. However, it did point out that any ambiguity as to who controls the data is detrimental. In the design of any system such as ours, it is key

that the user have a clear, obvious, and accurate model of how much control he as over his data and how it will be used.

A more troublesome problem is that users never really have 100% control of the data. In our case, users estimated that there was a small, but non-zero, probability that social pressures within the group would force them, in order to be team players, to turn over their data for research purposes. This is similar in nature to "rubber hose" [17] attacks on privacy, where an authority demands that you turn over your data or else suffer physical coercion. In our case there was no coercion, but the potential social pressures amounted to a mild form of a rubber hose attack with the same resultant problems.

Another sort of privacy violation occurs when a person is tracked in an unwanted manner. By analogy, consider a college dormitory where a student wishes to put up a public webcam showing video of his activity. We would expect no objection if this camera was focused upon just that student. But if other students were in the field of view, much opposition could be possible. There's a corresponding problem with personal sensing systems.

SPECs transmit IDs to the environment in order to enable applications. This, however, poses a privacy threat in that these IDs can be tracked by adversaries. There are also problems that occur even without an adversary. Consider a SPEC application, such as reminding, or sensor logging, used by a specific person for personal purposes. Also assume a similar system employed by an unrelated person. Neither person would expect their system to interfere with the other person's system, nor would they expect to exchange data between the systems. Some mechanism is needed to prevent this interaction as well as to foil adversaries.

The solution we are currently implementing is to scramble the IDs in a manner that gives the user control over who can make sense of them [15]. The ID a SPEC transmits is changed at a fixed interval, currently set at 15 minutes. The new ID is derived from the previous ID using a secret key, which is only available to the SPECs participating in a particular application as well as any associated archive or sense-making services. As a result, SPECs participating in other applications have no way of interpreting the IDs broadcast from another application, since they lack the key.

This solution has properties that are important to us, mostly related to energy use. Transmitting a changed ID uses no more bits than if we kept the ID the same. Every extra bit transmitted uses a significant amount of energy, so this is highly beneficial. In addition, our solution does not require any round-trip communication with another device. Round-trip communication is difficult to do well in an error-prone medium such as IR, unless even more bits are used to do error correction. The energy cost of cryptography is limited by our selection of a 15-minute ID change interval. Our experiments and calculations lead us to believe the energy use for this interval is acceptable.

In evaluating this solution, we find that regular changing of IDs does not eliminate all ID-related privacy threats. During a specific 15 minute period, an adversary can treat an ID as a pseudonym that can be tracked, revealing partial information that may be detrimental to the user. We expect this problem to disappear quickly over the next few years as the energy consumption of cryptography decreases due to better hardware. As this improves, we will be able to change IDs more quickly, eventually reaching a point where an ID is changed after a single use. At that point the

pseudonym tracking problem is eliminated, although the mere fact that an ID is transmitted does reveal some information, namely that a SPEC system is in use. If in the future we wanted to solve that problem, we could employ physical-layer mechanisms, such as spread spectrum, that allow physical sensing of IDs only by those that have the key for those IDs.

While we've found the privacy benefits of a non-centralized system to be great, there are some downsides. One is that the user is responsible for the safekeeping of his own data. This entails undesirable data management activity, such as organizing and backing up files. We have started to address these issues using a friendly peer based backup service that is transparent to the user.

Another problem is that it is difficult for a user to prove to a 3rd party facts contained within user's data. In a centralized system there is often a party trusted by all that collects and holds the data, thus allowing one to make assertions that can be verified by that party. For example, if a user wants to prove he was in a certain location, he can have the trusted party verify this. A SPEC user, however, can't prove he was in a given location just by showing his data. We don't envision this need in any of our applications, but we concede that it may appear as we experiment further.

Open Research Problems

Our research group doesn't have the capacity to even scratch the surface in many areas, but we want to inspire specialist in other areas to explore with us. Here we identify just a few of the challenges we see ahead.

We intend to ultimately move away from IR in favor of other technologies. We would like to have a technology that does not have to be worn in a visible location on the user, and can shrink to very small dimensions. Fortunately there are some new RF based ranging systems that appear promising.

We are investigating ways of detecting nearby SPECs in a low-power manner without lengthening the ID transmit interval. If successful, this will allow us to keep discovery responsive while consuming less power. Methods we are considering include: 1) a better clock, so SPECs can be synchronized and transmit only at predictable times, 2) a timing reference signal to help coordinate clocks [18], and 3) wakeup radios specifically designed for this purpose.

We are looking forward to the development of wakeup radios by other groups. An ideal wakeup radio would alert us quickly to the presence of other nearby devices in the social context, while drawing power in the range of tens or hundreds of nanowatts. Our application is a peer-to-peer system of identical devices, thus the power we wish to minimize is the sum two sorts of energy: that used to detect the presence of others, and that used to alert others to our own presence.

The form factor for the SPEC received little thought beyond what was necessary to put a box around a circuit board to prevent it getting broken when in use. Using IR meant that the device had to be worn where it was visible to other nearby SPECs. Various ingenious means were found for attaching it, decorating it, and/or disguising it, but it was nerdy[19], and marginally acceptable at best. Our next version is intended to be used for an application to support home care-givers of the elderly. We




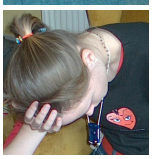

are working on a new more capable device with more sensors, and a better user interface that can be worn on the wrist as we have learned from health care workers that these patients are more likely to accept a wrist-worn device. We also have plans for a device that uses RF rather than IR which can be worn under clothing, in a pocket for example. But again this is an interim solution. We believe the proper imperatives are to reduce the size of the SPEC (to the size of a speck :-), while extending the battery life – ultimately moving to nanotechnological solutions to enable SPECs to be embedded in almost anything. This of course raises issues about where the user interface can be accommodated.

A major weakness of our current system is the user interface, which was designed for technologists to monitor the system and not for end user applications. Our envisioned applications require a method of alerting the user, and our single LED is not sufficient, even for system maintenance purposes. In the short term we plan on using other user alerting mechanisms, such as a beeper or pager vibrator, but we believe the proper solution is for all SPECs to incorporate the capability to discover and recruit user interfaces in the environment. Whether the remote user interface resides on a watch, a cell phone, or a shared device like a TV, or a public display in a supermarket is a matter for the application designer. We are exploring this concept in our Recruitable User Interface initiative which is defining an architecture for this purpose.

We believe a system must offer strong privacy guarantees in order to give real overall value to users. Thus, we are working on secure methods that protect the privacy of users against adversaries. Because of the low power nature of our devices, along with intermittent connectivity and high cost of communication, there are many opportunities for new protocols and organizations that perform much better in our application than the solutions found in current networks and systems.

We have hardly scratched the surface on sense-making and visualization tools. While specific ad-hoc data analysis tools can often do a good job of analyzing the human activity data, we really need a more disciplined approach, and look forward to collaborating with experts in this field. Equally, presenting complex time series data is a challenge, even for debugging, and there is some interesting research to be done on how best to present human episodic data.

From the point of view of researchers who participated in the arrival, and/or development of the personal computer, and who are looking for a new, if not more challenging frontier to explore, this technology is tremendously exciting. Personal sensing applications offer the opportunity to address the highly valued information processing needs of literally every man, woman and child on the planet, and at the same time push the envelope of design of computer systems in ways that we would never have previously imagined.

References

1. Lamming, M., Bohm, D. *SPECs: Another Approach to Human Context and Activity Sensing Research, Using Tiny Peer-to-Peer Wireless Computers*. Ubicomp 2003: 192-199
2. Mann S. *Sousveillance, not just surveillance, in response to terrorism*. http://www.chairetmetal.com/cm06/mann-complet.htm
3. http://www.spychips.com/
4. Newman, W., Eldridge, M., Lamming, M. *Pepys: Generating Autobiographies by Automatic Tracking*. In Proceedings of the second European conference on computer supported cooperative work. Amsterdam: Kluwer Academic Publishers. 1991.
5. Lamming M. and Flynn M. *"Forget-me-not" Intimate Computing in Support of Human Memory*, Proceedings of FRIEND21, '94 International Symposium on Next Generation Human Interface, Japan, February 1994
6. Sellen, A. J., Louie, G., Harris, J. E., Wilkins, A. J. *'What Brings Intentions to Mind? An In Situ Study of Prospective Memory'*. Memory, Vol. 5, No. 4, 483-507 1996.
7. Tulving, E. (1983). *Elements of Episodic Memory*. Oxford University Press.
8. Reiser, B.J., Black, J.B., & Kalamarides, P. (1986): *"Strategic memory search processes",* in D.C. Rubin (ed.): Autobiographical Memory, Cambridge University Press, Cambridge, 1986, pp. 100-121.
9. Smith, S., Glenberg, A.M., & Bjork, R.A. (1978): *"The influence of environmental context on recall and recognition"*. Memory & Cognition, vol. 6, no. 4, 1978, pp. 342-353.
10. DataEdge 1999. http://www.pdacortex.com/invivodata_study.htm.
11. McKenzie, S. et al. *Proving the eDiary Dividend.* Applied Clinical Trials, June 2004 www.invivodata.com/pdf/ACT_eDiary_ROI_2004.pdf
12. Shea, Heidi E. *Electronic Patient Diaries in a Clinical Trial-The Holistic Approach.* Drug Information Journal(2004), v.38, no.3, pp. 225-238.
13. Powsner, S., Spitzer R. *Sex, lies, and medical compliance*. The Lancet; Jun 14, 2003; 361, 9374;
14. Stubblebine, A. *Regular Expression Pocket Reference*. 1st Edition. Sebastopol, CA. O'Reilly, 2003. Series: Pocket References. SBN: 0-596-00415-X.
15. Harper, R. *Why people do and don't wear active badges: a case study*. Comput. Supported Coop. Work, 4(4), pp 297-318, Kluwer Academic Publishers, 1996.
16. Bohm, D. Mayo, R. et al. *Controlling who tracks me.* Twelfth International Workshop on Security Protocols. Cambridge, England. 26-28 April 2004 http://www.lamming.com/mik/Papers/who_tracks_me.pdf
17. Soghoian, C. Defending against Rubberhose Attacks. JHU Systems Seminar, March 9 2004. http://spar.isi.jhu.edu/~chris/presentations/rubberhose.pdf
18. Mayo, R. HP Labs. Tech report. In preparation.
19. Seuss, Dr. *If I ran the zoo*. New York: Random House, 1950