"Are You With Me?" – Using Accelerometers to Determine if Two Devices are Carried by the Same Person

Jonathan Lester[1], Blake Hannaford[1], and Gaetano Borriello[2,3]

[1] Department of Electrical Engineering, University of Washington
[2] Department of Computer Science and Engineering, University of Washington
[3] Intel Research Seattle
Seattle, WA, USA
{jlester@ee, blake@ee, gaetano@cs}.washington.edu

Abstract. As the proliferation of pervasive and ubiquitous computing devices continues, users will carry more devices. Without the ability for these devices to unobtrusively interact with one another, the user's attention will be spent on coordinating, rather than using, these devices. We present a method based on a coherence function, a measure of linear correlation in the frequency domain, to reliably analyze walking data recorded by low-cost MEMS accelerometers to determine if two devices are carried by the same person. We use inexpensive accelerometers and show that these sensors perform similarly to more expensive accelerometers for the frequency range of human motion, 0 to 10Hz. We also present results from a large test group illustrating the algorithm's robustness and its ability to withstand real world time delays, crucial for wireless technologies like Bluetooth and 802.11. We present results that show that our technique is 100% accurate using a sliding window of 8 seconds of data and the devices are carried in the same location on the body (we also present results for when devices are carried on different parts of the body), is tolerant to inter-device communication latencies, and requires little communication bandwidth.

For the past 30 years, the dominant model for using our computing devices has been interactive. This approach puts the human in a feedback loop together with the computer. A user generates input and the computer responds through an output device that is observed by the user who then reacts with new input. When the ratio of humans to computing devices was close to 1:1, this was a reasonable approach. Our attention was commanded by one device at a time, our desktop, laptop, or handheld. This was appropriate as our tasks often involved manipulating information on the computer's screen in word processing, drawing, etc.

Today, the conditions of human-computer interaction are rapidly changing. We have an ever-increasing number of devices. Moreover, they are becoming deeply embedded into objects, such as automobiles. Many of these devices have a powerful CPU inside of them, however, we do not think of them as computing devices.

There are two main implications of this explosion in the number of computing devices. First, the human user can no longer be in the loop of every interaction with and between these devices. There are just too many, the interactive model simply does not scale. To further complicate matters, each device will necessarily have a different specialized user interface due to their different functions and form-factor.

In the future, unless devices share information appropriately, these devices will end up demanding even more of our time as use becomes more casual. Second, as these new devices are embedded in other objects, we often are not even aware there are computing devices present, because our focus is on our task, not on the devices.

Invisibility is an increasingly important aspect of user interaction with the principal tenet being "do not distract the user from the task at hand". An example of this is package delivery, which now includes tablet-like computers to collect signatures, RFID tracking of packages, and centralized databases to provide web services to customers, such as the current location of their parcel. Delivery truck drivers, cargo handlers, and recipients do not want a user interface to slow down a package in reaching its destination. They prefer if the devices gather input, explicitly or implicitly, and communicate the data amongst themselves. There is no reason for users to take an interactive role with all the steps, nor do they want to. Another example draws on devices becoming so cheap that they are viewed as a community resource. In hospitals, nurses and physicians carry clipboards, charts, and folders that could provide more timely information if they were electronic devices connected to the hospital's infrastructure. Many individuals would use these devices as their paper versions are now. It would be much more efficient for the devices themselves who they were being carried by.

Motivated by these examples, we are investigating methods for devices to determine automatically when they should interact or communicate with each other. Our goal is to enable devices to answer questions such as:

- Is the same person carrying two devices? With what certainty?
- Are two devices in the same room? For how long?
- Are two devices near each other? How near?
- What devices did I have with me when I came in? When I went out?

Different applications will want answers to a different set of these and other questions. We are developing a toolkit of technologies and methods that can be used by interaction designers to create systems with a high degree of invisible interactions between many devices. This paper presents our work on developing methods for answering the first question.

We assume a world where a user will carry a changing collection of devices throughout a day. These might include a cell phone, a laptop, a tablet, and a handheld. In addition, the collection may include more specialized devices such as, RFID or barcode scanners, GPS receivers, wrist-watch user interfaces, eyeglass-mounted displays, headphones, etc. These devices may be tossed into a pocket, strapped to clothing, worn on a part of the body, or placed in a backpack or handbag. We posit that it will be an insignificant addition to the cost of these devices if they include a 3-axis accelerometer. We also expect these devices to have a means of communicating with each other through wireless links, such as Bluetooth or 802.11. Recent work in wireless sensor networks is demonstrating that the communicating nodes may become as small as "smart dust" and function with a high degree of power efficiency.

Here we consider the problem of how easily and reliably two such devices can autonomously determine whether or not they "belong" to the same user by comparing their acceleration over time. If the acceleration profiles of the two devices are similar enough, they should be able to conclude that the same person is currently carrying them. In practical situations, one of these devices may be a personal one (e.g., a wrist-watch or pager on a belt that is "always" with the same person) while the other is a device picked-up and used for a period of time. The question we posed is: "How reliably can accelerometer data be used to make the determination that two devices are on the same person?"

2 Related Work

There are several techniques that could be used to answer our question. We could use capacitive-coupling techniques to determine if two devices are touching the same person [1], however, this requires direct physical contact with both devices and is highly dependent on body geometry and device placement. A second approach is to use radio signal strength [2] to determine proximity of two devices. However, RF signal strength is not a reliable measure of distance and is also highly dependent on body orientation and placement of devices. Furthermore, these RF signals could be received by nearby unauthorized devices on another person or in the environment.

The approach we found most appealing is to use accelerometers to directly measure the forces acting on two devices and then compare them over a sliding time window. There has been much previous work in using accelerometers for gesture recognition and device association. We describe the contributions of three different pieces of related work: two from the ubiquitous computing research community and one from bioengineering instrumentation.

Gesture recognition using accelerometers has been used to develop a wand to remote control devices in a smart space [3] and a glove that uses sensing on all the fingers to create an "air keyboard" for text input [4]. This work is primarily concerned with using accelerometer data as part of the process in computing the position of an object, in these cases, a plastic tube or finger segments, respectively. By observing the variations in position over time, gestures can be recognized.

Device association is the process by which two devices decide whether they should communicate with each other in some way. Work at TeCO used accelerometers to create smart objects (Smart-Its) that could detect when they were being shaken together [5]. The idea was to associate two devices by placing them together and shaking the ensemble. Similar forces on the two devices would allow the connection to be established. The assumption is that it would be unlikely that two devices would experience the same forces unintentionally. Hinckley has developed a similar technique that uses bumping rather than shaking [6]. In both these cases, the analysis of the accelerometer data is in the time domain, which can be sensitive to latencies in communication between the devices. Both issues also have similarities in that the decision is strictly binary, instead of a probability that the two devices are being intentionally associated. However, the principal difference between this work and ours is more fundamental. While these two contributions exploit explicit user-initiated interactions (shaking or bumping), our focus is on making the determination implicitly and independent of the user's attention.

The work that is closest to ours, and provided much of our inspiration, used accelerometers to determine if the trembling experienced by a patient with Parkinson's Disease was caused by a single area of the brain or possibly multiple areas of the brain [7]. The key observation in this work is that Parkinson's related shaking is likely due to multiple sources in the brain that may be coupled to each other. Shaking in the same limb was found to be highly correlated, however, shaking across limbs was found to be uncorrelated. Physicians developed accelerometer sensors that were strapped to patients' limbs, data was collected, and an off-line analysis determined if the shaking was correlated in the frequency domain. We use a very similar approach, but with an on-line algorithm which can be running continuously within the devices being carried rather than strapped to the body.

Researchers attempting to identify activities in real-time have performed research to identify multiple activities, standing, sitting, walking, lying down, climbing stairs, etc. using more structured placement of multiple sensors and analysis methodologies like neutral networks and Markov models [10]. See [8] for a detailed overview of this work.

3 Our Approach

In order to provide a useful detection tool that doesn't require any user interaction; the input to our system must come from an existing, natural action. In this paper we focus on the activity of human walking.

Although there are a number of different actions that a person regularly performs, walking provides a useful input because of its periodic nature. Human motion is regulated by the mechanical characteristics of the human body [9] as much as conscious control over our limbs. This regular, repeated activity lends itself to an analysis in the frequency domain, which helps reduce the effect of problems like communication latencies, device dependent thresholds, or the need for complex and computationally expensive analysis models.

We have two aims in this paper. First, we want to assess the quality of acceleration measurements obtained from low cost accelerometers, to ensure that these devices are appropriate for this application and that their measurements have a physical basis, by comparing their outputs with higher quality sensors subjected to the same accelerations. Second, we want to determine whether there is sufficient information in the accelerations of two devices to determine whether they are being carried together.

4 Methods

Three different 3-axis MEMS accelerometers were used for our experiments, two were low cost commercial accelerometers from Analog Devices [4, 5] and STMicroelectronics, and the third was a calibrated accelerometer from Crossbow Technologies. Table 1 lists the accelerometers along with some specifications for each device.

Table 1. Chosen specifications from each of the three accelerometers used

Manufacturer	**Model**	**Axis**	**Sensitivity**	**Bandwidth**[1]
Analog Devices	ADXL202E	2-axis[2]	± 2G	~100Hz
STMicroelectronics	LIS3L02	3-axis	± 2G/6G[3]	~100Hz
Crossbow Technologies	CXL02LF3	3-axis	± 2G	50Hz

Since the accelerometers used in these experiments provide 3-axis outputs (X, Y, and Z) data can be processed on a given set of axes, or taken as a whole. We assume a random, and possibly continuously changing, orientation and thus we take the magnitude of the force vector by combining the measurements from all 3 axes using Eqn. 1 to derive a net acceleration independent of orientation. Each accelerometer is capable of measuring the 1G gravity field present on the Earth, and although this information can be used to estimate the direction of the gravity vector, we subtract this offset from our data to reduce the DC offset in the Fast Fourier Transform (FFT) frequency spectra.

$$A_{mag} = \sqrt{(A_x)^2 + (A_y)^2 + (A_z)^2} \quad \textbf{(1)}$$

Figure 1 shows the magnitude of the output of the accelerometers for a subject walking, and riding an elevator. For this experiment, as well as the majority of the experiments in this paper, the subject wore the accelerometer in a fanny pack worn around the waist. The use of the fanny pack ensured that the placement of the sensors was consistent between different test subjects. However, we also show results for carrying the devices elsewhere on the body.

[1] Bandwidth as configured in this experiment
[2] Two ADXL202E's are mounted at 90º to provide 3-axis (plus one redundant axis)
[3] The LIS3L02 has a pin selectable sensitivity between ±2G and ±6G

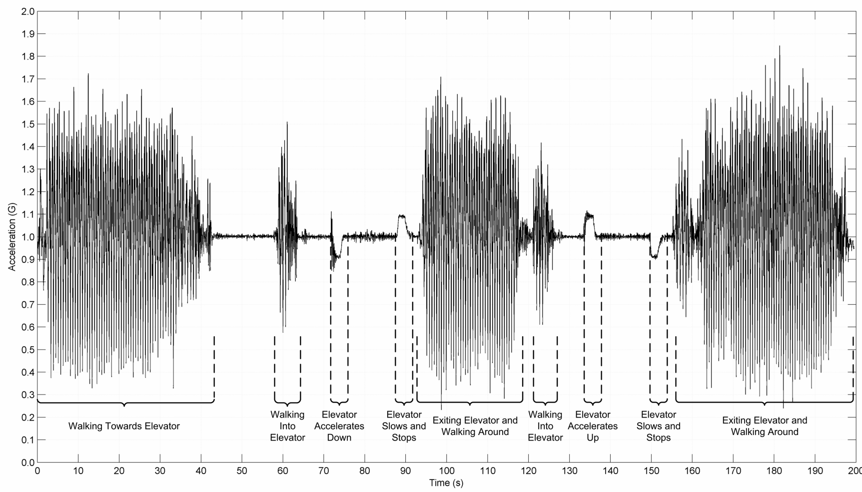

Fig. 1. Example of filtered magnitude data from accelerometer, similar to [10]. Data shows subject walking, riding down an elevator, walking, riding up an elevator, and walking

4.1 Accelerometer Characterization and Experimental Setup

The purpose of using a precise calibrated accelerometer is to characterize the performance of the low cost accelerometers; to ensure that the data provided is an accurate representation of the physical world. Without verification any experimental results could be based off of false data or rely upon artifacts present in the sensors and not in the physical environment.

The physical mechanics and low-level neural circuits of the human body controls the way humans walk, more than their conscious control of limbs. This automatic control makes walking a regular periodic activity, which has been studied extensively in the biomechanics community. It is widely accepted that the useful frequency spectra of human motion lies within the range of 0 – 10Hz [9]. Using this range as our region of interest we focus our study of accelerometers on the 0 – 10Hz range.

To characterize the accelerometers several experiments were conducted using different combinations of accelerometers. Data collection was carried out using a custom interface board, (approx. 9cm x 5cm) designed to record data through a serial connection to an iPaq, and a commercial data collector from Crossbow Technologies, the AD128. Our results showed that for low frequencies, 0 – 10Hz, the ADXL202E and LIS3L02 (the smaller and cheaper accelerometers) performed similarly to the CXL02LF3 (the larger and more expensive, but more robust, accelerometers), producing very consistent magnitude and FFT spectra. These results allow us to reliably conclude that low cost accelerometer are appropriate for collecting data in this frequency range and that they provide reliable data. Figure 3 shows an experiment performed with two accelerometers placed on a swinging pendulum, both traces are nearly identical in time and frequency domains.

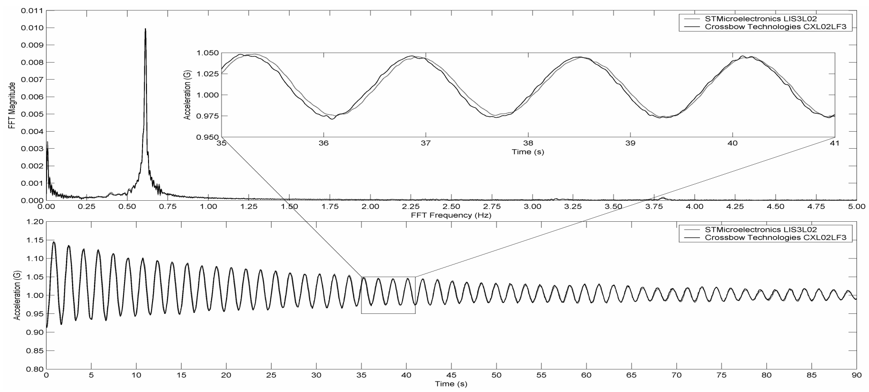

Fig. 2. Acceleration magnitudes and their corresponding FFT spectra collected during a swinging pendulum test showing that both signals are very similar in the time and frequency domains

4.2 Coherence

Given a time series of acceleration magnitude signals, like those found in Figure 1, it was necessary to find a reliable method for examining the data. Ben-Pazi *et al.*[7] used the coherence function to analyze the origin of rest tremors in patients with Parkinson's. By examining their use of coherence with biological accelerometer data, we reasoned that such methods could be expanded to other signals like the accelerations experienced by the body when walking.

Coherence, $\gamma_{xy}(f)$, is a normalized function of frequency derived from the cross-spectrum of two signals, S_{xy}, divided by the power spectrum of each signal, S_{xx} and S_{yy} [11, 12]. Coherence measures the extent to which two signals are linearly related at each frequency, with 1 indicating that two signals are highly correlated at a given frequency and 0 indicating that two signals are uncorrelated at that frequency [7, 11, 13]. Since the coherence is a complex quantity it is often approximated using the magnitude squared coherence (MSC), C_{xy}, shown in Eqn. 2. As all coherences used in this paper are real valued we'll refer to the MSC estimator simply as the coherence.

$$C_{xy}(f) = \left|\gamma_{xy}(f)\right|^2 = \frac{\left|S_{xy}(f)\right|^2}{S_{xx}(f)\,S_{yy}(f)} \qquad \textbf{(2)}$$

If the two signals used to compute the coherence are identical, then the coherence gives a unity result for all frequencies. Similarly if two signals are two completely uncorrelated random processes, the coherence is zero for all frequencies. For example, two separate audio signals recorded from an orchestra playing the exact same score would have a high coherence, near 1, for most of the audio spectra, whereas two orchestras playing completely different scores would have low coherences, near 0, for most of the audio spectra. And of course, two similar audio signals would have coherences spread throughout the range.

Eqn. 2 would always result in a unity magnitude for all frequencies (though the imaginary component may not be 1) if a single window were used to estimate the spectral density. The use of multiple windows in estimating the spectral densities allows the coherence magnitude to be non-unity. A commonly used windowing method is called weighted overlapped segment averaging (WOSA) and involves splitting two signals, X and Y, into equal length windowed segments. The FFT of these segments is taken and their results are averaged together to estimate the spectral density. As the name suggests segments can be overlapped to reduce the variance of the spectral estimate (an overlap of 50% is fairly common), however, overlapping is more computationally expensive. As our goal is to use these algorithms in environments with limited computing power we limit ourselves to two windowed segments, no overlapping segments, and a FFT size equal to the size of the input data [11, 14, 15].

Different window sizes/types as well as different sized FFTs can be used to obtain coherences with different attributes. Different window types have different characteristics and for simplicity a common Hanning window is used in our calculations. The number of windows used in a coherence calculation determines the significance level of the coherence output, that is, the more windows used the lower coherence has to be to signify a strong coherence. In general, more windows provide smoother coherences, with less variability, but require more computations.

Due to the fact that the coherence is a function of frequency, it was necessary to determine a method of computing a scaled measure of similarity based on coherence. A basic measure is to integrate the area under the coherence curve for a given frequency range. By leveraging the fact that physical human motion rests below the 10Hz range, we used a normalized integration over the range from 0 – 10Hz (i.e., by multiplying the resulting integral by 0.1) to get a 0 to 1 measure of the coherence (Eqn. 3). This result is expressed as P_{xy}, and is the measure of similarity between our two acceleration signals.

$$P_{xy} = 0.1\int_0^{10} C_{xy}\left(f\right)\ df \tag{3}$$

Although more complex methods are possible, the results presented in this paper only use this basic normalized integration of the coherence from 0 – 10Hz. This allows us to reliably quantify the coherence output using relatively fast and computationally inexpensive methods.

5 Experimental Data and Results

Two main experiments are discussed here. In the first experiment, the single-person experiments, six test subjects walked normally for ~30 meters wearing two accelerometers in a fanny pack around their waist. Each subject walked this distance eight times; twice with four different combinations of accelerometers and data collectors (listed in Table 2). Giving a total of 48 sets of recordings for the first experiment.

Table 2. Each subject completed the following four walking trials twice, a total of 8 walking trials from each subject

Walking Trial	First Sensor	Second Sensor	Data Collector
1, 2	ADXL202E	CXL02LF3	AD128
3, 4	LIS3L02	CXL02LF3	AD128
5, 6	ADXL202E	CXL02LF3	Custom Board
7, 8	LIS3L02	CXL02LF3	Custom Board

In the second experiment, the paired walking experiments, six subjects walked a distance of ~61 meters in pairs with another test subject. Again, each of the subjects wore two accelerometers in a fanny pack around their waist and data was collected using a custom interface board and an iPaq. During this experiment each subject walked with every other subject for a total of 15 pairs. Each pair of subjects walked the 61 meter distance four times, the first two the first two times they walked casually with the other person, and the last two times they walked purposely in stride with the other person. Giving a total of 60 sets of recordings for the second experiment.

The data from these two experiments was then manually trimmed so that only the walking data from each pair of recordings was selected. It is possible to use this method on a continuous stream of data by, for example, performing the calculations on 1-second wide windows of data and using the combined result over the past 8 seconds.

5.1 Single-Person Walking Experiment

Figure 3 shows a sample of the data collected from the first experiment. Fifteen seconds of acceleration magnitude data (bottom plot) were used to calculate the FFT spectra and the coherence (top plot). The two signals compared in Figure 3 were from data recorded on the same individual, during the same walking trial. The high coherence shown in Figure 3 indicates that the two signals are highly correlated at most of the frequencies. Notice also that the two FFT spectra appear similar as the two signals recorded the same walking pattern.

Figure 4 shows a similar figure, however, this time the two signals compared are not from the same individual. The data from the CXL02LF3 sensor in Figure 4 is the same 15-second segment shown in Figure 3 while the ADXL202E data is from a different person's walking segment, which was chosen because of its similar time plot. The wildly varying coherence shown in Figure 4 indicates that the two signals are not highly correlated, despite appearing similar in the time plots. The FFT spectra shown in Figure 4 are noticeably different as well.

To analyze the data in the first experiment we took the first 8 seconds of walking data from each trial and computed the coherence of each possible pairing of recordings. This created a 48 x 48 matrix (not shown) of coherence results. The diagonals of this matrix correspond to the (matching) coherences between pairs of sensors worn on the same person during the same trial. And the off diagonals correspond to the (non-matching) coherences between pairs of sensors worn on different subjects during different trials. Viewing the 0 to 1 range of the coherence integral as a percentage from 0 – 100% we found that the matching diagonal coherences had a mean of 90.2% with, a standard deviation of, 5.2%. By contrast the non-matching off diagonal coherences had a mean of 52.3% with a standard deviation of 6.1%.

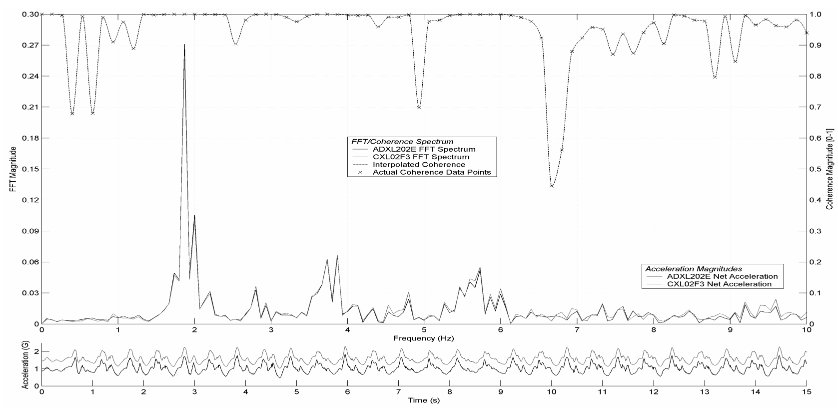

Fig. 3. Coherence and FFT spectra (top) of matching devices on the same person during a 15 second walking segment – (bottom) acceleration magnitudes as a function of time. Coherence is nearly 1 for the 0 – 10Hz range, indicating two highly correlated signals

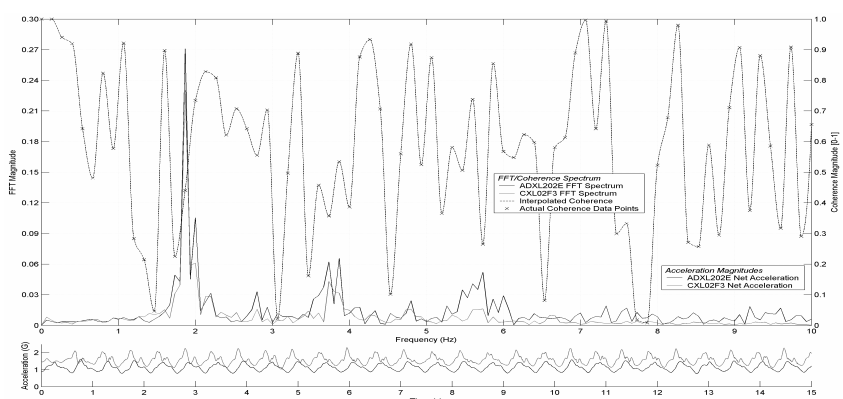

Fig. 4. Coherence and FFT spectra (top) of acceleration measured on different subjects during a 15 second walking segment – (bottom) acceleration magnitudes as a function of time. Coherence is not 1 for most of the 0 – 10Hz range, indicating two uncorrelated signals

5.2 Paired Walking Experiment

To analyze the data in the second paired walking experiment we again took the first 8 seconds of walking data and computed the coherence of each possible pairing of recordings. Unlike the first experiment, we now have four sets of accelerometer data, instead of two, because each of the two people walking together carries two accelerometers. Two accelerometers correspond to the first person in the pair, test subject A, and the other two correspond to the second person, test subject B[4]. We can therefore create two 60x60 matrices of coherence results, one corresponding to the two accelerometers on test subject A and the other to test subject B. Figure 5 shows the 60x60 matrix created using only the data from test subject A (test subject B's data creates a similar matrix and is not shown).

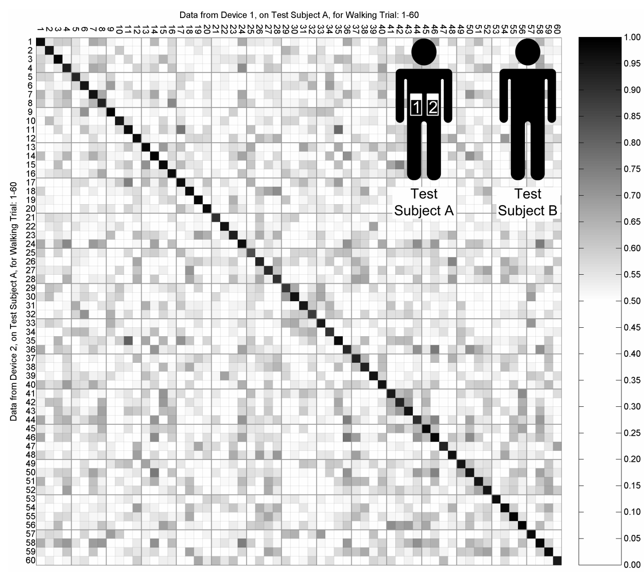

Fig. 5. Coherence values between all possible pairs of device 1 and device 2 worn on test subject A. In this matrix, coherences are calculated only with the first 8 seconds of walking data. The dark diagonal indicates high coherence between sensors worn on the same person at the same time

Again, the diagonals correspond to the coherence of matching pairs of accelerometers worn on the same person, test subject A, at the same time. As we would expect the diagonals show a high coherence with a mean of 95.4% ± 3.1% std. And the off diagonals show a low coherence with a mean 52.9% ± 6.6% std as they correspond to sensor data from different people recorded at different times.

---

[4] Test subject A and test subject B refer to the first and second person in each trial pair and not to any specific test subjects in the trials

Another useful matrix to examine would be the coherence between a device on test subject A and a device on test subject B. If coherence were affected by the two people walking together casually, or in step, we would expect to see high coherences on the diagonal of this matrix. Again, since each subject had two accelerometers on them we could construct more than one such matrix, we only show one of these matrices in, Figure 6. To create the matrix in Figure 6 we took the first accelerometer data from subject A and the first accelerometer data from subject B in each pair and computed the coherence of each possible pairing of recordings (other combinations of accelerometers on different test subjects produce similar results). In Figure 6 the diagonals correspond to the coherences of two devices worn on different individuals, who walked together at the same time. The off-diagonals represent the coherences of two devices worn on different individuals, recorded at different times. Overall the matrix in Figure 6 shows a low coherence, with the diagonals having a mean of 56.9% ± 8.4%, and the off-diagonals having a mean of 53.1% ± 6.4% (the entire matrix has a mean of 53.1% ± 6.8%).

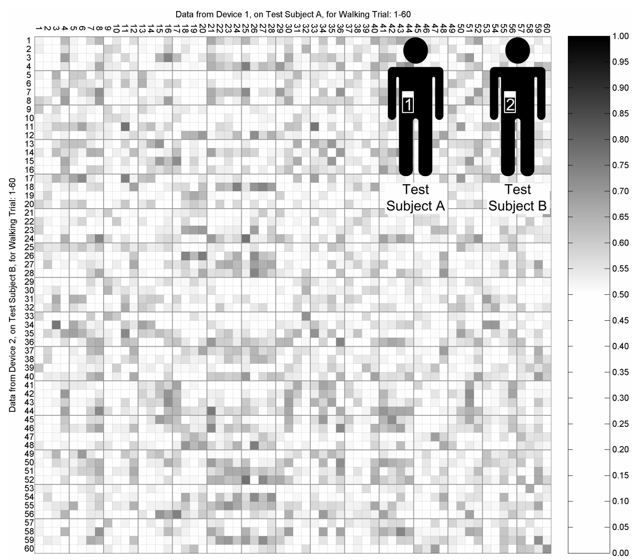

Fig. 6. Coherence values between all possible pairs of device 1, worn on test subject A, and device 2, worn on test subject B. In this matrix, coherences are calculated only with the first 8 seconds of data from the walking trials. The light shading throughout the matrix indicates low coherence between devices on different people. The second two rows/cols of each 4x4 box are the lock-step trials in which subjects intentionally tried to walk synchronously

In Figure 6 the 3rd and 4th row/column of each of the 4x4 boxes in the matrix represent data from when the subjects were walking together in step. As you can see there is little coherence difference between subjects walking casually together and those walking purposely in step.

6 Practical Considerations

For the results presented so far, only the first 8 seconds of walking data were used to generate coherence values. However, different sized segments of data can be used to calculate the coherence, with longer segments of walking data providing more information and shorter segments providing less information. The obvious question is how large do the segments need to be so as to get 'enough' information.

To address this question, we re-analyzed the data from our first, single-person, experiment by splitting the entire data set in segments of 1 – 15 seconds (in 1-second increments). We would expect there to be enough information in a segment if coherence is able to successfully pick the corresponding sensor pairs from the same segment. Since the walking data from the first experiment was typically only about 20 seconds in length this gives us 48 pairs, for 15 seconds, all the way up to 1,212 pairs, for 1 second segments. For each segment we computed all the possible coherence pairs and recorded the pair with the largest coherence. If the largest coherence belonged to the same pair as our chosen signal then the algorithm correctly determined the best fit pairing. We then repeated the process using the next segment until all segments had been processed. Figure 7, shows the results of this analysis in the trace labeled 'No Delay', which shows if 8-second, or larger, segments are used then the largest coherence is between the correct matching pairs of 100% of the time.

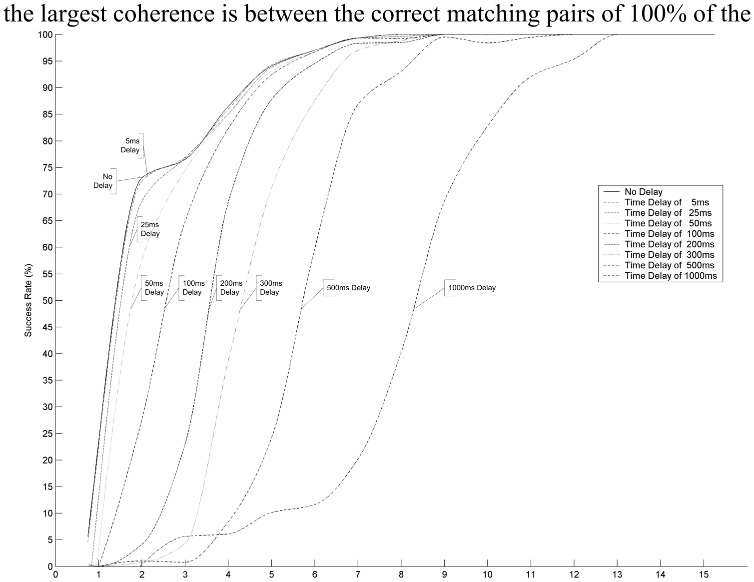

Fig. 7. Interpolated success rates showing the effects of time delay on the coherence results, using different length of segments. For segment lengths greater than eight seconds the success rate remains >95% even with a 500ms time delay

Another practical consideration in deploying this algorithm is understanding how sensitive it is to communication latencies between the two devices. Although data transfers would be done in large chunks, i.e. 1 or 8 second sliding windows, it is desirable if the system is relatively insensitive communication delays. Wireless synchronization methods do exist; however, they can require more power hungry radio communication. Given that our analysis uses magnitudes in the frequency domain, we would expect some insensitivity to phase shifts in the data. Figure 7 also shows how our results change as we vary the latency between the two devices. We simulated latencies ranging from 0 to 1000ms. Each curve shows the accuracy for different segment lengths at these different latencies. The results are only slightly degraded up to 500ms with more substantial differences above that latency. We claim that this is more than adequate for two devices likely to be in short-range radio contact. For example, Bluetooth packet latency is on the order of 50ms in the worst case.

Another practical consideration is the amount of data to be transferred to calculate the coherence. Because our frequency range lies in the 0 – 10Hz range, we can filter the signal and only sample at 20Hz. At this sampling rate it is only necessary to transmit 320 Bytes to the other device to send an 8-second acceleration magnitude (8 seconds x 20Hz x 2 bytes), which could fit in a single Bluetooth packet. This would enable devices with little computational power or battery life to have a more powerful device perform the calculation. However, if the sampling rate needed to remain high we could simply transmit the 0 – 10Hz FFT coefficients necessary for the coherence calculation. Figure 8 shows the situation where accelerometers are still sampled at 600Hz and the devices do a wireless handshake and then transmit the FFT data.

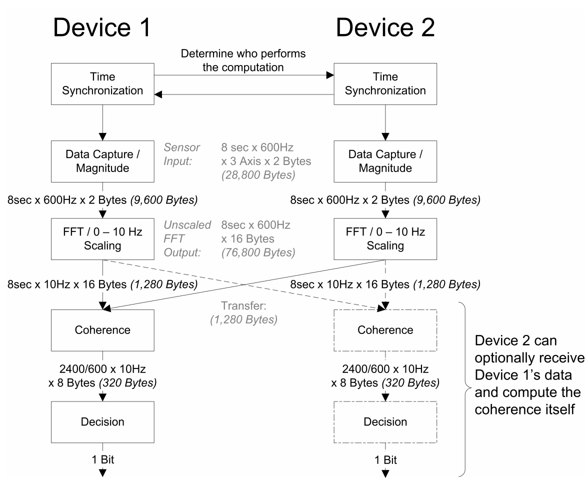

Fig. 8. Hypothetical data exchange as it would take place on embedded devices. Results in Figures 3 – 7 use higher resolution computations. Labels on the arrows indicate the data that must pass from one step to the next. Gray text in center indicates sensor input data, data generated internally, and transferred data between devices. Eight-second segments are used in this example, resulting in 4,800 magnitude data points and a 2,400-point FFT

The two devices can coordinate which of them will make the determination. The data need only be sent to one device, most likely the one with the processor capacity required, or the most battery power, etc. Alternatively, or possibly for security reasons, both devices can make the determination that they are likely on the same person. Of course, a completely secure implementation would require a more complex coordination process.

7 Discussion

As the activities of interest in our experiments are human motion we are able to focus on the frequency range of 0 – 10Hz. Our initial experiments using a pendulum provided a well-controlled environment for closely measuring the behavior of our accelerometers against significantly more costly calibrated accelerometers like the Crossbow Technologies CXL02LF3. The pendulum experiment clearly showed that the low-cost accelerometers produced similar acceleration magnitude data and very similar FFT spectra for the sampling rates we required.

Our first, single-person, experiments with different accelerometer pairs and data collectors confirmed that the accelerometers performed similarly in the frequency range of interest. Using these results we believe it safe to conclude that these inexpensive accelerometers are appropriate and well suited to the tasks of recording highly periodic human motion. We also believe that they provided measurements which were consistent with the physical world and were not merely the artifacts associated with a particular accelerometer.

Our first experiment also showed us that the coherence integral does provide a reasonable measure of the correlation of two acceleration magnitudes. We were able to show that correlated 8-second segments showed mean coherences of 90.2% and that uncorrelated 8-second segments showed mean coherences of 52.3%. This striking difference in coherence clearly indicates a separation between correlated and uncorrelated results. The second experiment, the paired walking experiments, further expand these results showing that devices on the same person have mean coherences of 95.4%, with uncorrelated results having a mean of 52.9% while devices recorded on two different individuals have mean coherences of 56.9% and 53.1%, even with people purposely attempting to confuse the results. The large separation between correlated and uncorrelated results is also clearly evident in Figures 5 and 6, clearly showing that we can indeed determine when sensors were worn on the same individual. There were no false positive results from the sensor readings from the same person (recorded during different walks), people walking casual together, and even people purposely trying to walk in lock step together.

8 Conclusions and Future Work

We set out with two goals for this paper: (1) showing that inexpensive accelerometers are adequate for our tasks and (2) showing that we can accurately determine when two devices are being carried by the same person using these accelerometers. We believe our experimental data strongly support positive conclusions in both cases. We have also answered our original question: “How reliably can accelerometer data be used to make the determination that two devices are on the same person?" by showing that we can reliably determine if two devices are being carried by the same person using just 8 seconds of walking data. We can even differentiate between devices carried by two people walking together.

As we are concerned only with the physics of human motion we are only interested in frequencies up to 10Hz. This is of great advantage for small, portable, and low-power devices as the accelerometers only need to be sampled at 20Hz to meet the Nyquist criterion. Low sampling rates will also help facilitate the use of wireless communication to offload computationally expensive portions of this algorithm to devices with more powerful processors and/or batteries.

Our results for determining whether two devices are carried by the same person are limited to people that are walking – either independently or side-by-side and even in lock-step. We use the values of frequency coherence between 0 and 10Hz to provide us with a rough likelihood that the two devices were being carried by the same person by integrating the area under the coherence curve. The results are not only very positive but they are also very robust. If we use an 8-second window for the comparison, our results are 100% accurate (no false positives, no false negatives). We feel that the most impressive result is that segments which should have high coherences have mean coherences in the 90 – 95% range with standard deviations on the order of 5 – 6% and segments that should have low coherences have means in the range of ~52% with standard deviations of ~6%. This substantial separation between correlated and uncorrelated signals illustrates our algorithm's robustness and even holds up when people attempt to purposely walk in lock-step to fool our algorithm.

If we narrow our data window, the results necessarily degrade, but only moderately. For example, with a 5 second window, we still have ~95% accuracy, even at 2 seconds, it is still in the ~75% range. This is encouraging as it is likely that other sensors can be used in conjunction with the accelerometers (e.g., microphones) to use additional data to make the determination.

The major limitations of our current approach is that we are requiring the user(s) to walk and that we only considered devices in a fanny pack. We have yet to extensively study the cases of other everyday motions. For example, sitting and typing at a desk, will likely require more analysis in the time-domain as there is less periodicity in the person's movement. Again, we see the technique presented here as only one of many analysis tools available to the interaction designer. For example, this method could be enhanced with additional sensor data, simple time domain analysis, etc.

Preliminary experiments have been conducted with devices at other locations on the body, with experiments with devices on the wrist, placed in one or both pockets, in a back-pack, and in a fanny pack. Figure 9 shows some preliminary experimental results with sensors worn in a pocket/fanny pack, pocket/backpack, and pocket/wrist. Table 3 lists the statistics gathered from this preliminary experiment. While the results from different locations for the devices are not quite as good as when both devices are in the same fanny pack, they still show promise. Coherence clearly is an important tool for solving the device association problem, in general. However, it is also clear that it will need to be enhanced with other techniques appropriate to the usage model.

Table 3. Stastitics from coherence analysis of three different preliminary experiments

Experiment	Diagonal	Off-Diagonal	Average Success Rate
Pocket / Fanny Pack	76.7% ± 4.6%	55.3% ± 8.8%	87.5%
Pocket / Backpack	77.1% ± 6.0%	54.9% ± 8.8%	81.3%
Pocket / Wrist	75.7% ± 6.5%	51.9% ± 6.9%	70%

Fig. 9. Data from three different experiments using sensors worn on different parts of the body, Pocket/Fanny Pack, Backpack/Pocket, and Wrist/Pocket

Our future plans are to incorporate these methods and others into a general purpose sensor platform that can be easily integrated with a variety of devices such as cell-phones, key-chains, digital cameras, and wrist-watches, as well as PDAs, tablets, and laptops. We are currently designing a new sensor board with accelerometers, microphones, light, and inertial sensors that will fit into a package roughly 2.8 cm^2. We will be using an embedded ARM7 processor that will be able to handle the computations easily along with Bluetooth RF communication.

When this new sensor platform is complete, we will also conduct user studies for typical use in two scenarios: borrowable cameras that can keep track of what other devices (and what person) they were being carried by when each picture was taken, and automatic synchronization of data between PDAs, laptops, and wrist-watches.

Acknowledgements

We would like to thank David Mizell for starting this project. Adam Rea and Ken Smith at Intel Research Seattle provided valuable support in getting our instrumentation in place. We would also like to thank Crossbow Technologies and STMicroelectronics for the use of their accelerometers. And most of all, we would like to thank our anonymous test subjects for their patience and willingness to participate in our experiments. Jonathan Lester was supported by a gift from Intel Research Seattle.

References

1. Partridge, K., Dahlquist, B., Veiseh, A., Cain, A., Foreman, A., Goldberg, J., and Borriello, G.: Empirical Measurements of Intrabody Communication Performance under Varied Physical Configurations. In: Symposium on User Interface Software and Technology, Orlando, Florida (UIST 2001) 183 – 190
2. Brunette, W., Hartung, C., Nordstrom, B., and Borriello, G.: Proximity Interactions between Wireless Sensors and their Application. In: Second ACM International Workshop on Wireless Sensor Networks and Applications, San Diego, California (WSNA 2003) 30 – 37
3. Wilson, A., and Shafer, S.: XWand: UI for Intelligent Spaces. In: Conference on Human Factors and Computing Systems, Ft. Lauderdale, Florida (CHI 2003) 545 – 552
4. Perng J., Fisher B., Hollar S., and Pister K. S. J.: Acceleration Sensing Glove. In: 3rd International Symposium on Wearable Computers, San Francisco, California (1999) 178–181
5. Holmquist, L., Mattern, F., Schiele, B., Alahuhta, P., Beigl, M., Gellersen, H.: Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. In: Proceedings of Ubicomp (2001)
6. Hinckley, K.: Synchronous Gestures for Multiple Persons and Computers. In: Symposium on User Interface Software and Technology, Vancouver, British Columbia (UIST 2003) 149 – 158
7. Ben-Pazi, H., Bergman H., Goldberg J. A., Giladi N., Hansel D., Reches A., and Simon E. S.: Synchrony of Rest Tremor in Multiple Limbs in Parkinson's Disease: Evidence for Multiple Oscillators. In: Journal of Neural Transmission, Vol. 108 (3) (2001) 287–296
8. Bao, L.: Physical Activity Recognition from Acceleration Data under Semi-Naturalistic Conditions. Master's thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science (2003)
9. Winter, D.: Biomechanics and Motor Control of Human Movement (2nd ed.). New York: Wiley (1990)
10. Lukowicz P., Junker H., Stäger M., von Büren T., and Tröster G.: WearNet: A Distributed Multi-Sensor System for Context Aware Wearables. In: Proc. of the 4th International Conference on Ubiquitous Computing, Göteborg, Sweden (2002) 361–370
11. Carter, C.: Tutorial Overview of Coherence and Time Delay Estimation. In: Coherence and Time Delay Estimation – An Applied Tutorial for Research, Development, Test, and Evaluation Engineers. IEEE Press (1993) 1–23
12. Carter, C.: Coherence and Time Delay Estimation. In: Proc. of the IEEE, Vol. 75. (1987) 236–255
13. Shulz, M., and Stattegger, K.: Spectrum: Spectral Analysis of Unevenly Spaced Paleoclimatic Time Series. In: Computers & Geosciences, Vol. 23, No. 9 (1997) 929 – 945
14. Carter, C., Knap, C., and Nuttall, A,: Estimation of the Magnitude-Squared Coherence Function Via Overlapped Fast Fourier Transform Processing. In: Coherence and Time Delay Estimation – An Applied Tutorial for Research, Development, Test, and Evaluation Engineers. IEEE Press (1993) 49 – 56
15. Welch, P.: The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms. In: IEEE Transactions on Audio and Electroacoustics, Vol. 15, No. 2 (1967) 70 – 73