

Monitoring Body Positions and Movements During Sleep using WISPs

Enamul Hoque
Department of Computer
Science
University of Virginia
eh6p@virginia.edu

Robert F. Dickerson
Department of Computer
Science
University of Virginia
rfdickerson@gmail.com

John A. Stankovic
Department of Computer
Science
University of Virginia
stankovic@cs.virginia.edu

ABSTRACT

Sleep monitoring is very important for elderly people as inadequate and irregular sleep are often related to serious diseases such as depression and diabetes. In many cases, it is necessary to monitor the body positions and movements made while sleeping because of their relationships to particular diseases (i.e., sleep apnea and restless legs syndrome). Analyzing movements during sleep also helps in determining sleep quality and irregular sleeping patterns. This paper presents a sleep monitoring system based on the WISP platform - active RFID-based sensors equipped with accelerometers. We show how our system accurately infers fine-grained body positions from accelerometer data collected from the WISPs attached to the bed mattress. Movements and their duration are also detected by the system. We present the results of our empirical study from 10 subjects on three different mattresses in controlled experiments to show the accuracy of our inference algorithms. Finally, we evaluate the accuracy of the movement detection and body position inference for six nights on one subject, and compare these results with two baseline systems: one that uses bed pressure sensors and the other is an iPhone application.

Keywords

Assisted Living, Activity Monitoring, WISP, RFID

Categories and Subject Descriptors

H.3.4 [Systems and Software]: User Profiles and Alert Services

General Terms

Algorithms, Performance, Experimentation

1. INTRODUCTION

Sleep plays an important role in quality of life, and is an important factor in staying healthy, active, and energetic. Having inadequate and irregular sleeping patterns

has a serious impact on our health, and can lead to many serious diseases like cardiovascular disease, diabetes, depression, and obesity ([2]). Besides the amount of sleep, it is also necessary to have sound sleep. Despite sleeping for a sufficient amount of time, people can still feel fatigued and cannot concentrate during the day. This may be caused by interrupted sleep, such as having frequent periods of restlessness during sleep. Moreover, in many cases, particular body positions should be maintained or avoided. For example, patients with obstructive sleep apnea should avoid sleeping on their back ([18]).

Sleep monitoring systems are important to recognize sleeping disorders as early as possible for diagnosis and prompt treatment of disease. They can provide healthcare providers with quantitative data about irregularity in sleeping periods and durations. They can also provide detailed sleeping profiles that depict periods of restlessness and interruptions such as bed exits and entries due to visiting the bathroom. This information helps find trends that correlate to certain diseases. Moreover, it enables monitoring effectiveness of treatments to sleep-related diseases. Many studies ([18]) are focused on finding correlations between body positions during sleep to various breathing problems (e.g., sleep apnea). So, if a sleep monitoring system can provide fine grained information about body positions during sleep, it would help such studies.

To date, there are very few low-cost, unobtrusive sleep monitoring systems. Among the existing ones, the most accurate and reliable are polysomnography (e.g., electroencephalogram, electrooculogram, electromyogram) devices [8]. There are drawbacks to using them, since they need to be worn, and require professional monitoring and thus are expensive to use. A less obtrusive way of monitoring sleep is an actigraph ([21]), a device that can be attached to any of the limbs (e.g., wrist) to provide data on movement; however, they still need to be worn. There are systems based on audio and video signals for sleep monitoring, but they raise privacy concerns for people. Some sleep monitoring systems have been developed that work by sensing body pressure/vibration; these systems are unobtrusive, but can be uncomfortable to sleep on.

We propose a sleep monitoring system based on Intel WISPs ([22]) (Wireless Identification and Sensing Platform). Our system does not require any additional action from the users outside their daily routines. We attach WISP tags to the bed mattress and collect accelerometer data reported by them. We analyze these data and infer body position of the users and movements they make while on the bed. We can then

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Wireless Health' 10, October 5-7, 2010, San Diego, USA
Copyright 2010 ACM 978-1-60558-989-3 ...\$10.00.

record entries and exits from the bed, and movements and body positions during sleep. Recently WISPs ([7]) have been used to recognize various daily activities. Our system complements such activity recognition systems. We compare the performance of our system with two baseline systems. Firstly, to evaluate the performance of the WISP devices as close to ground truth as possible, we compare it with a baseline system that uses pressure pads to measure the movement levels. We validate the pressure sensor based baseline system by comparing it with video data recorded by a camera during sleep. Secondly, as our system uses accelerometer data, we compare its performance with another sleep monitoring system that uses an accelerometer, choosing the popular iPhone based sleep monitoring application “Sleep Cycle” ([1]).

The main contributions of this paper are: 1) a novel sleep monitoring system based on WISPs; 2) an empirical study with 10 subjects on three different mattresses to evaluate the accuracy of the body position inference algorithm of our system; 3) a realistic evaluation of the movement detection and body position inference algorithm of the system with one subject for six nights and comparison of our system with a baseline system using bed pressure sensors 4) performance comparison with a popular iPhone based application “Sleep Cycle” to show that our system can help in detecting transitions between sleep cycles and thus can help in understanding sleeping patterns.

The rest of the paper is organized as follows. Section 2 summarizes existing sleep monitoring systems and their advantages/disadvantages with respect to our system. Section 3 presents a brief description of WISPs and how we use them in our sleep monitoring system. Section 4 describes short-term controlled experiments and their results, which show that our system provides fine grained monitoring of body positions. Section 5 presents the algorithm and experiments to show how we detect body movements made while lying on the bed. Section 6 details the realistic overnight experiments we conducted to verify the algorithms and their results. We conclude in Section 7.

2. RELATED WORK

Physiological signals are regarded as the most accurate means to differentiate between sleep and wake phases. Electroencephalogram (EEG), the measurement and frequency analysis of brain waves, shows the difference between sleep and wake stages ([8]). Electrooculogram (EOG) and electromyogram (EMG) are also standard technologies for sleep monitoring. Electrocardiogram (ECG) is used to derive heart rate which is well known to decrease at sleep onset. [23] studies heart rate variation over different sleep stages. [20] and [14] use respiratory-derived features together with ECG-derived features for classifying different sleep stages automatically.

But all these techniques have major limitations. For example, they must be performed by trained professionals and in clinical environments. Many commercial sleep clinics are available for this purpose. But their expense limits the monitoring to one or two nights. Besides, these techniques require equipment to be attached to the patients. This limits the movements of the patients which may be uncomfortable. Also, these physiological signals do not support monitoring body positions during sleep.

Temperature regulation in a body can also be used to

monitor sleep quality. [15] shows that skin temperature increases during sleep onset and decreases during wakeup. But these temperature variations can only be measured under controlled laboratory conditions. [25] uses infrared triangulation distance sensor to detect movements of different body parts without attaching any device to the body. But it does not provide any information about body position.

To overcome the limitations of the above techniques, there are many systems that enable sleep monitoring in home environments. Actigraphy ([21]) is a commonly used technique for sleep monitoring that uses a watch-like accelerometer based device attached typically to the wrist. The device monitors activities and later labels periods of low activity as sleep. There are many commercial products like [4] that have been designed based on actigraphy. [6] is another commercial product for sleep monitoring in home environments. It is a headband that users need to wear each night so that it can detect sleep patterns through the electrical signals naturally produced by the brain. But all these products are very expensive and users need to wear the device.

Another method used for sleep monitoring is to instrument a mattress pad with sensors and passively infer body movements and sleep quality. [3] is such a commercial bed pressure sensing pad that monitors change in body pressure on the pad to detect movements. In [11], the authors use pressure and temperature sensors laid out in a grid pattern in the mattress to determine quality of sleep. NAPS ([17, 16]) is a low-cost physiological sensor-suite that can passively acquire important physiological and environmental characteristics. The NAPS suite allows subjects to simply lie on a mattress pad, embedded with vibration sensors, to obtain multidimensional data (e.g., body temperature, heart rate, respiration rate, positional mapping and movement). One might also use tiny sensor motes with accelerometers in place of vibration sensors. The main advantage of all these solutions is that users do not need to wear any device. But, in some cases batteries are needed and it may also be uncomfortable to sleep on a pad and thus, they can affect sleep quality.

Audio and video signals can also be used to determine sleep quality accurately. In [19], a combination of heart rate, audio and video sensors is used to infer a sleep-awake condition. But such systems raise privacy concerns among the users.

[1] is a popular iPhone based application that uses the accelerometer in the iPhone to monitor body movements and determine which sleep phase the user is in. The user just needs to put the iPhone in a suitable place on the bed. But it does not provide fine-grained body position monitoring. Also, the iPhone can accidentally fall off the bed and it needs to be connected to the charger for the whole night.

Recently WISP tags have been used for recognizing daily activities. In [7], the authors instrument everyday objects (e.g., glass, plate, books) with WISP tags that have accelerometers. Later, daily activities are inferred from the traces of objects that are moved. In [9], WISPs are used for sensing and monitoring of exercises involving free weights. The authors instrument free weights (e.g., dumbbells) and body parts (e.g., wrists) with WISPs. Then the accelerometer sensor readings from the tags are used to infer the exercise being done and the association between the user and the particular weight(s) being used. [26] presents a wireless neural interface that uses WISPs. It provides the neurosci-




Figure 1: WISP Tag



Figure 2: Speedway Reader

entists a wireless, battery-free method of monitoring neural signals.

In summary, the advantages of our WISP-based sleep monitoring system are that users do not need to wear any device, they do not need to sleep on any mattress pads instrumented with sensors, no batteries are needed, the system is wireless, and it avoids privacy violations of video solutions.

3. WIRELESS IDENTIFICATION AND SENSING PLATFORM

The sensing element in our system are WISP ([5, 22]) active RFID tags that combine passive UHF RFID technology with sensors to balance energy independence with sensing potential. A WISP tag is shown in Figure 1. The device’s antenna and power harvesting circuitry enable off the shelf EPC “Gen 2” RFID readers (shown in Figure 2) to power and read from it.

To a RFID reader, a WISP appears as a normal RFID tag, but inside the WISP, the harvested energy is operating a 16-bit fully programmable ultra low-power microcontroller. The microcontroller can sample a variety of sensing devices including 3-dimensional accelerometers, lights, and temperature sensors. In our system, we only use 3D accelerometer readings. The WISP tags report these readings by encoding them as part of their identifiers that are read by a RFID reader. While WISPs are currently assembled from discrete components that have a cost of roughly \$25, they are intended to be mass manufactured like RFID tags at price points closer to \$1 [7].

3.1 Hypothesis: WISPs Help in Sleep Monitoring

When a WISP tag remains stationary, its 3D accelerometer measures the acceleration force due to gravity with respect to each axis. The x, y and z axes values are dependent on the device’s alignment with respect to gravity. As long as the alignment remains the same, the reported values should remain the same (plus/minus some noise). But, if the alignment along any of the three axes changes, then the accelerometer reading along the corresponding axis also changes. We exploit this property to infer whether the bed is empty or not. If the bed is not empty then we can infer




Figure 3: Accelerometer Reading Variation for Empty/Lying/Sitting

body position of the person on the bed using this property.

For example, when we attach one WISP tag along the edge of the mattress (similar to the position of Tag 1 in Figure 6) then the variation along the y-axis (shown in Figure 3) can be used to differentiate among the cases when the bed is empty, someone is lying on it, or someone is just sitting on the bed watching television or reading. When the bed is empty, the y-axis accelerometer of the tag is aligned perpendicular with respect to gravity, but when someone lays on the bed, because of the impact of the body on the mattress, the orientation changes. Again these two different orientations are different from the one when someone watches television sitting on the bed. The accelerometer readings returned by WISP are a scaled version of the actual values normalized by the full scale 10-bit ADC reading (1024). When static, an axis perpendicular to gravity reads about 50%.

Using the accelerometer readings, we can also distinguish the following four positions: lying on the back, stomach, left, or right sides (shown in Figure 4). In the Figures 3 and 4, we show the accelerometer readings along the y-axis. The readings along the z-axis(which is parallel to gravity) also show similar variation. The readings along the x-axis do not show too much variation, but if we combine them with the readings along the y and z-axes, together they can differentiate among the above four positions. So, by combining readings along all the axes and by using more tags, we can increase the accuracy of our prediction, which we show in the next section.

While on the bed, each time someone moves, the accelerometer readings change rapidly. So, by looking at the change in readings we can detect each movement. Our system monitors how many times a person tosses and turns during the course of the night and how many times the person leaves the bed. If someone is lying on the bed and does not move for a significant amount of time, then we can assume the person is asleep. In this way, the frequency of movements can help in determining whether a person is asleep or not. Frequency of movements is also different for different sleep stages and thus can be related to which sleep stage a person is currently in ([13]). Transitions between different sleep stages also correspond to change in frequency of body movements. Thus, based on a summary of movements




Figure 4: Accelerometer Reading Variation for Different Lying Positions




Figure 5: Heights and Weights of the Subjects

made during each night, doctors can infer quality of sleep and irregular sleeping patterns.

4. CONTROLLED EXPERIMENTS

In this section, we describe the controlled experiments we conducted to test whether WISP tags can be used to infer body position of a subject on the bed. We then present the results and analyze them.

4.1 Design

For our controlled short-term experiments, 10 graduate students volunteered as subjects. The subject population was diverse in both height and weight (shown in Figure 5). All participants were volunteers, and were informed of the experimental procedures and the study’s goals prior to participation. We conducted our experiments in three different beds to consider how different mattresses affect the measurements. Five subjects were evaluated on a twin-size bed in our University’s medical testbed ([24]). Five other subjects participated in each of the other two beds that were in a graduate student’s apartment. All three of them were twin mattresses. So, for each bed, five subjects participated in the experiments. For two of the three beds the participating subjects were the same.

For each experiment, we attached 3 WISP tags to the




Figure 6: Experimental Setup



Figure 7: Body Positions while Lying on the Bed

mattress of a bed. Figure 6 shows such a bed along with the positions of the tags. We placed the tags in such a way that when someone lies on the bed, there is one tag on each side of his body and one tag near the legs. We used two antennas for reading from the tags. The reader sends 10 read requests per second. The read rate from each of the tag was 4 – 7 reads per second during all our experiments. If we use one antenna, then read rate of one or more tag falls much lower. One disadvantage of using the WISP tags is that they need to be placed within 1 – 2 meters of the antenna of the reader. To meet the read-range requirement and to keep the equipment away from obstructing a resident’s movement, we placed the antennas below the bed. The antennas were wired to the reader which was connected with the laptop.

Each subject laid on the bed in the following four positions: on the back, on the stomach, on the left side and on the right side. These four positions are shown in Figure 7. Each subject also sat on the bed with his or her back on the wall and face towards the camera (Figure 6). This position resembles the way someone lies when watching television or reading a book while sitting in bed. For each position, we recorded data for two minutes. For each WISP tag, we obtained the acceleration along the x, y and z axes. From the readings of all three tags, we get a 9-tuple. Note that, all three tags do not report their acceleration values synchronously. We combine the readings from the three tags within each second and construct each possible 9-tuple. We associate all the 9-tuples collected during these two min-

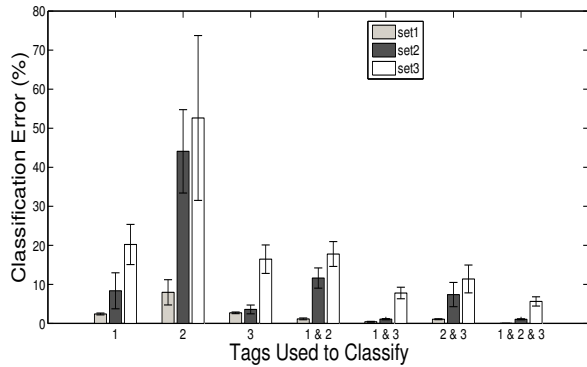


Figure 8: Average Classification Error for five Subjects for One of the Beds

utes to that particular body position. We also recorded the readings from the tags when the bed was empty. We use the collected data to train our system.

After the training phase, the subject repeats the tasks again and we record data for 30 seconds for each position. Our system then classifies the new data based on previous training. For training and classification, we use the open source software “Orange Canvas” [10], which supports a number of classifiers. We decided to use the Naive Bayesian Classifier. Note that for each subject, first we train our system based on the subject’s training data and then classify his or her remaining data.

4.2 Results

For each subject, we classify the collected data under three different settings. In the first setting, we test whether it is possible to differentiate between the bed being empty and someone lying on it (in any position). So we label all data collected during a subject lying on the bed in four different positions as lying. We do not include the data when the subject was sitting on the bed. In the second setting, we include the data for sitting and test whether it is possible to differentiate among the bed being empty, or someone lying on it (in any position), or someone sitting on it. In the last setting, we test whether it is possible to differentiate among all six cases: empty, lying on back, lying on stomach, lying on back, lying on left side, lying on right side and sitting. We name the above three cases as “set1”, “set2” and “set3”.

For each setting, first we train and classify based on the data collected from one tag only (tag no. 1 or 2 or 3 of Figure 6). Then we use data from a combination of two of the three tags. Finally, we use data from all three tags. Our goal is to test how increasing the number of tags helps in reducing classification error. The results of our experiments on one of the three beds is summarized in Figure 8. Five of the 10 subjects participated in the experiments on this bed. For each case, the y-axis shows the average of percentage classification errors for all five subjects. The error bars represent the standard deviations of the errors for each experiment.

As we see from Figure 8, if we increase the number of tags, the classification error decreases. When we use data from only one tag, the performance of tag 2 is the worst. This is expected, because it is placed near the leg, and so




Figure 9: Average Classification Error for All Mattresses

it fails to capture enough of the variation of body impact on the middle portion of the mattress for different positions. When we use data from any two of the three tags, we see that the combination of tags 1 and 3 performs best. This is because both of them are placed in the middle parts of the two opposite edges of the mattress. When we use data from all three tags, the error for “set1” becomes almost zero. For “set2” and “set3”, average percentage errors are 1.06% and 5.64%, respectively. For the other two mattresses, we also observe similar trends, i.e., increasing the number of tags increases classification accuracy.

We also check how classification error varies over different mattresses. Figure 9 shows average classification error for all mattresses. Here we calculate the average over the classification errors for all subjects that participated in the experiments on a particular bed. As we see from the figure, classification error for “set1” is almost zero for all mattresses. But for the other two sets, classification error is greater for mattress 3 than the other two mattresses. This mattress is the one that is in our university testbed. The testbed quality is different than the other two. It is hard and inflexible. So, the impact of the body weight does not change the orientation of the WISP tags immediately. As mentioned earlier, we classify the body positions for 30 seconds of data for each subject. Later we used the data from the last 20 seconds and the classification error went down significantly and was approximately same as the other two mattresses. So for such mattresses, we need to classify the body position after the body settles in to a new position.

Now, we analyze what body positions are misclassified most. Here, we consider misclassifications for “set3” only. For every mattress, the case when the bed is empty is classified correctly. For the other positions, the average of misclassifications for each mattress is shown in Figure 10. Here, we see that classification error is most prominent for the two body positions where a subject lies on back and on stomach. The reason is that sometimes one of these is classified as the other. For both these positions, the impact of body weight on the mattress remains almost same. For the other three positions, the classification error remains less than 10% for each mattress.

Note that, for these controlled experiments, the training

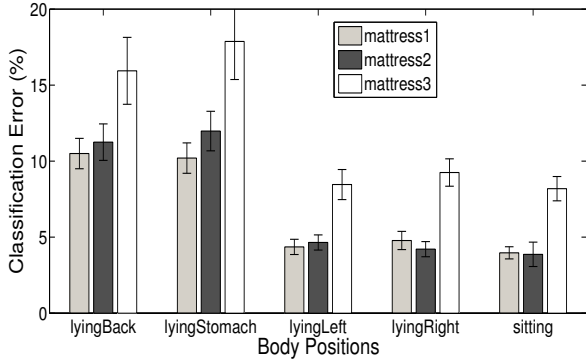


Figure 10: Average Classification Error for Different Body Positions for All Mattresses

period is only two minutes for each body position. For practical use, we need to train the system for longer periods. During our realistic overnight experiments, we train our system for several nights (about seven hours per night) and then run the system. The results are much better and shown in Section 6.

5. MOVEMENT DETECTION ALGORITHM

As we see from Figures 3 and 4, when a subject lies on the bed in a particular position or when the bed remains empty, the accelerometer values returned by the WISP tags remain within a noise level of a particular value. This is true for acceleration values along each of the three axes. To find the maximum deviation in the readings, we calculate the derivative of all the readings when a subject remained in a particular position. The derivatives show that if the subjects remain in a particular position or if the bed is empty, the deviation remains in the interval $[+a, -b]$. The values of a and b vary for different tags, axes and mattresses, but remain same for different subjects. We calculate these values from the data collected during the controlled experiments of the previous section.

If the subject moves to a new position or makes significant movements while remaining in the same body position, the derivative of the accelerations of all three tags along both y-axis and z-axis become higher than the corresponding $+a$ or lower than the corresponding $-b$. So during the movements, the derivatives of y and z acceleration values cross the threshold values ($+a$ and $-b$) several times. Figure 11 shows y-axis accelerometer readings during such a move. Here the values of both a and b are 1.

Our algorithm to extract movement events from derivatives of y and z-axes accelerations of the three tags is as follows: For each axis of each tag, we record timestamps when the reported reading is outside the interval $[+a, -b]$. We consider each of these moments a possible movement. Note that the three tags do not report values synchronously. We calculate the total number of movements reported by the three tags within each two second time window. If the total number of movements within a time window is less than a predefined threshold, we consider those as discrete movements that do not affect sleep quality. We then cluster the other time windows, when a significant number of move-




Figure 11: Accelerometer Reading along y-axis during a Movement

ments take place, using the DB-SCAN clustering algorithm ([12]) to compute discrete movement events. The clustering also ensures that discrete movement events that happen within a short amount of time are combined as a period of restlessness. For each cluster, we set the movement level as the maximum of movement levels of all the time-windows belonging to that cluster.

Figure 12 shows the number of movements for each 2 second time window during 70 minutes of a controlled experiment. During the experiment, the subject got on the bed, laid there for 70 minutes during which he made several movements and finally got off the bed. Some movements were from one body position to another and in some cases, the subject made significant movements while remaining in the same body position. We normalize the y-axis by dividing the number of movements for each time window by the maximum number of movements in any time window to get the movement level. We use 0.3 as the threshold to filter out the time windows where movement level is insignificant.

Figure 13 shows the discrete movement events as clustered by DB-SCAN. All the discrete movement events during the controlled experiment were successfully detected by our system. As we can see from Figure 13, some movement events span several minutes. During these movement events, the subject made a number of movements in quick succession. We comprehensively validate the performance of our movement detection algorithm by realistic overnight experiments that we present in the next section.

6. REALISTIC OVERNIGHT EXPERIMENTS AND RESULTS

In this section, we present the realistic overnight experiments we conducted to evaluate the accuracy of our movement detection and body position inference algorithm.

To evaluate the performance of our movement detection algorithm, we compare it with a baseline system that uses pressure pads to measure the movement levels. The pressure sensor we used was a USB-interface Multi-Platform Dance Dance Revolution (DDR) pad typically used in the popular DDR video game series. The configuration of the pad is shown in Figure 14. Two pads were tiled to cover the area of a twin size bed. Data collected from the DDR pad is a bit-vector of size 16 representing which of the 16 buttons




Figure 12: Number of Movements per each Two-second Time-window during 70 Minutes of Controlled Experiment




Figure 13: Discrete Movement Events during 70 Minutes of Controlled Experiment

are activated. Our algorithm examines a time window, and takes the sum of the number of changes occurring in this bit-vector in that window. We chose a window size of two seconds, same as we did for WISP tags. After calculating the number of movements during each two-second time window during the night, we clustered them in the same way as discussed in the previous section.

We also compare the performance of our system with an iPhone-based sleep monitoring application “Sleep Cycle” ([1]) that uses accelerometer data to infer sleep quality. The application requires the iPhone to be placed on a suitable position of the bed (e.g., beside the pillow) all night and it collects data from the accelerometer of the iPhone for the whole night. Based on the data, it produces sleep quality related data that includes transitions between different sleep cycles. Durations of different sleep cycles over the course of the night are part of a person’s sleeping pattern. So monitoring the transitions between sleep cycles helps in identifying irregular sleeping patterns. Our hypothesis is that transitions between the sleep cycles will correspond to higher number of

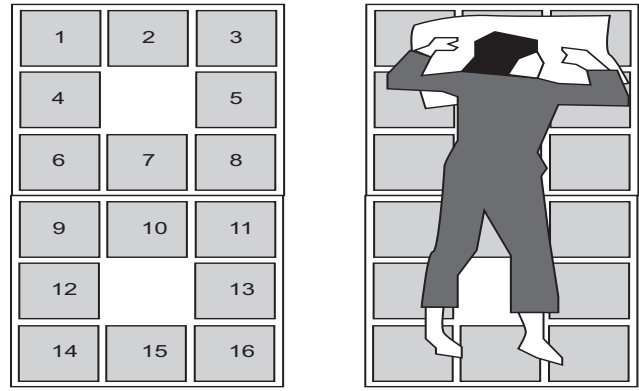


Figure 14: The DDR Pad has 8 binary contact buttons around the side, but the middle portion of the pad does not have a button. We tiled two pads can cover a twin size bed.

movements per time window. So from our overnight report of number of movements during each time window, we can infer the transitions between sleep cycles and the duration of each of them. We test our hypothesis in this section.

The study participant slept on the same bed for six nights. We collected and logged data from the DDR pads and the WISP devices simultaneously, and also placed an iPhone on the bed (beside the pillow) during each of these six nights. The “Sleep Cycle” application recorded sleep quality data and produced a report for each night. We also videotaped the sleeping period of the subject for each night after being given the subject’s consent. We first validated the performance of the DDR pads in detecting movements during sleeping by comparing it with the video data for the first three hours of the recorded data for the first night. The validation result confirmed that the DDR pads can be used as ground truth to detect movements during sleeping. For evaluation, we use a cross validation approach. For each evaluation set, we choose five nights’ data to train our system and evaluate the performance for the remaining night’s data. So, there are six possible sets of training data. Thus, we have six sets of evaluation.

For each evaluation set, training of the movement detection algorithm includes calculating the thresholds of rate of change of acceleration values (i.e., values of a and b) along each axis for each tag and also the threshold to filter out the time windows where movement level is insignificant. During training, we consider movements detected by the DDR pads as ground truth. Training of the body inference algorithm includes training the Bayesian classifier with the accelerometer readings collected during the five nights with the corresponding body position. Collecting the actual body position for each time instant of each these five nights is challenging. One option was to monitor the recorded video for each night and assign body positions accordingly. But this requires significant effort. To reduce effort, for each night, we watch the initial body position from the video and from then on we assume that unless there is a movement detected by the DDR pads, the position remains unchanged. When the DDR pads detect a movement, we fast forward to that time instant and see the new body position from the video and we continue in this way. Thus, we collect the ground




Figure 15: Movement determined by Our System during One Night’s Sleep of Evaluation Set 1




Figure 16: Movement Determined from the DDR Pad during One Night’s Sleep of Evaluation Set 1

truth for body position.

6.1 Evaluation of Movement Detection Performance

Figures 15 and 16 show the movement events during one night’s evaluation (from the first evaluation set) of the subject as detected by the WISP tags and DDR pads, respectively. If we compare these two figures, the first and last movements on both the figures represent the events when the subject got on and off the bed, respectively. Our system reported all movement events detected by the DDR pads. The timings of the movements are same in both figures. There was one movement that our system reported, but the DDR pads did not. It happened just after 7:00 AM in the morning. To investigate this incident, we fast forwarded to that specific time of the recorded video and observed that there was no significant movement during that time. So it was indeed a false positive.

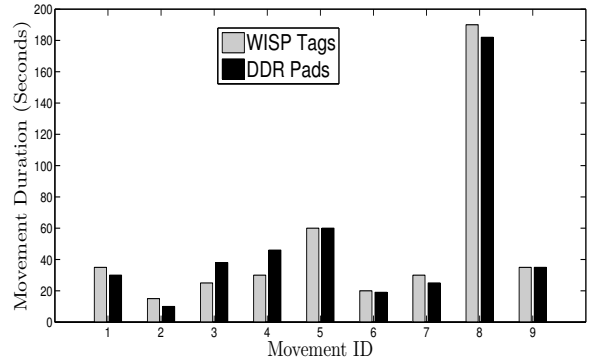


Figure 17: Durations of the Movements Detected by Our System and DDR Pads

Another notable difference occurred just before 9:00 AM in the morning, when our system reported two movement events and the DDR pads reported one movement event. However, the two events reported by our system are very close to each other and can be considered a part of the same movement. The duration of the nine movements during this night that both system detected are shown in Figure 17. From this figure, we see that there are no notable differences between the duration of movements calculated by both the systems. We present a summary of results and their implications for all six evaluation sets at the end of this section.

Figure 18 shows the report produced by the iPhone application “Sleep Cycle” to show the sleep quality for the same night as shown in Figures 15 and 16. The application shows various sleep stages like ‘awake’, ‘deep sleep’ and ‘dreaming’. These sleep stages are irrelevant for our comparison. The application recorded data up to 8:00 AM in the morning. The vertical bars show when movement events are reported by our system. As we know, transitions between different sleep cycles correspond to movements made by a person. From the figure we can see that the timings of the movement events match to those of transitions between sleep cycles. There are no vertical bars for two transitions: one that happened between 7:00 and 8:00 AM and the other in between 4:00 AM and 5:00 AM. During the latter one, the subject was in deep sleep stage before and after the transition. So this is why there were no major movements. We explain the reason of lack of movements during this transition at the end of this section. But this result proves our hypothesis that from the frequency of movements reported by our system, it is possible to infer transitions between sleep cycles. In addition, our system provides fine grained body position monitoring which the “Sleep Cycle” application does not.

6.2 Evaluation of Body Position Inference

Figure 19 shows the body positions as inferred by our system for the same night that was considered in Figures 15, 16 and 18. If we compare these four figures, we see that during each transition from one body position to another, there was a discrete movement event detected by our movement detection algorithm. Also, for the last three movements, the body position did not change. To ensure robustness against discrete erroneous classifications, we consider that the subject changed his body position if 20 successive instances are

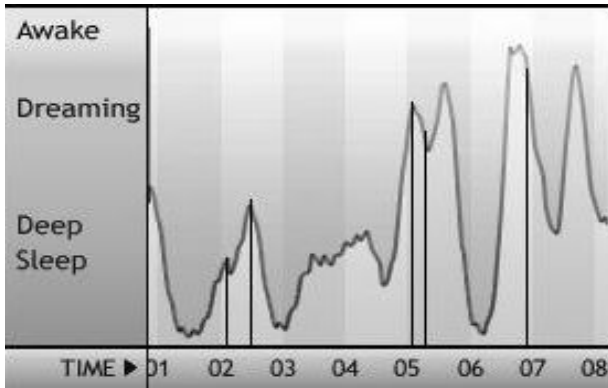


Figure 18: Sleeping Quality Report Produced by the iPhone Application




Figure 19: Body Positions During One Night's Sleep of Evaluation Set 1

classified as the new body position. Also, if the movement detection algorithm detects that a movement is taking place, the body position is considered to be the same as it was before the movement until the movement is complete.

To evaluate the performance of the body position inference algorithm of our system, we generate 10 random instances of time for each night and check the subject's body position during each of those instances. For each night, the time instances are uniformly distributed over the course of the night. We define the accuracy of our inference algorithm to be the percentage of time instances when the body position inferred by our system match to the actual body position as seen from the recorded video data. We present the accuracy for each night as part of the summary of all results next.

6.3 Summary of Results

Table 1 presents a summary of results for our six sets of evaluation. False negatives refers to the number of movement events that are detected by the DDR pads, but not by our system. Similarly, false positives refers to the number of movement events that are detected by our system, but not

Evaluation Set	1	2	3	4	5	6
False Negatives	0	0	0	0	0	0
False Positives	1	0	0	1	0	0
Avg. Error in Movement Duration	6.9s	6.2s	2.2s	5.2s	4.1s	5.2s
Sleep Cycle Detection Accuracy	71.4%	75%	80%	75%	90%	80%
Body Position Inference Accuracy	100%	100%	100%	90%	100%	90%

Table 1: Summary of Results for Six Evaluation Sets

by the DDR pads. For each night, we define 'average error in movement duration' as the average of absolute differences between the movement durations calculated by our system and the DDR pads. Sleep cycle detection accuracy refers to the percentage of sleep cycle transitions (as shown by the iPhone application) that correspond to increased number of movements detected by our system.

From Table 1, we see that for each set, our system detected all the movement events detected by the DDR pads. Average error in calculating movement duration is less than six seconds for each night. But, for two nights, we observe one false positive each in our system. This may be due to the threshold in change of acceleration that we selected to filter insignificant movements. We believe by training the system for more nights, we can get rid of such false positives. Overall, our system shows 100% accuracy in detecting discrete movement events and calculates the durations of each movement with reasonable accuracy. The accuracy of the body position inference algorithm is at least 90% for all sets. Our evaluation was based on 10 randomly selected time instances that are uniformly spread over one night. More detailed evaluation is necessary to guarantee its performance. Therefore, we can say that, with proper training, our system performs as well as a system that uses pressure sensors and also is more comfortable for the users and completely unobtrusive. Moreover, our system provides fine grained body position monitoring which no existing pressure sensor based sleep monitoring system provides.

We also compare our system with the popular iPhone based application "Sleep Cycle". Comparison results show that by only looking at the movement reports of our system, it is possible to identify most of the transitions between sleep cycles. Among the transitions that were not possible to identify, most of them were during deep sleep stages. The pressure sensors also did not identify them. So, these types of transitions do not correspond to significant body movements. We need to lower the value of the threshold for filtering out insignificant body movements which was set assuming the DDR pads' detected movements as ground truth. Therefore, we can say that, by training our system with the transitions detected by the iPhone application, it is possible to detect all the transitions between sleep cycles by our system.

7. CONCLUSION

In this paper, we have shown how sleeping movement and position can accurately be monitored using active WISP tags. We show a summary of various sleep monitoring sensors and their tradeoffs in Table 2. The WISP tags have a

Device	Comfort	Privacy	Accuracy	Price
WISP tag	High	Good	High	Moderate
Pressure Sensor	Medium	Good	High	Moderate
Camera	High	Very Low	Excellent	High
EEG/ActiWatch	Very Low	Good	Excellent	High

Table 2: Qualitative Comparison Across Devices

high comfort level as they do not affect the sleeping surface and no device needs to be worn on the body. The WISP tags themselves are cheap and even when including the reader, are only a fraction of the cost of professional monitoring devices such as EEGs. Moreover, if WISP tags are used for recognizing other daily activities, the reader cost will be amortized. The recognition accuracy is similar to pressure sensors, but at a lower price and intrusiveness. Finally, our system provides fine grained body position monitoring which none of the existing systems offer. In this paper, our experiments are limited to twin mattresses. Our future goal is to run these experiments on other types of mattresses. We also aim to infer sleep stages from these body movement related data.

8. ACKNOWLEDGEMENTS

This work was supported, in part, by NSF grant EECS-0901686. The WISPs used in this research were donated by Intel as part of the WISP Challenge.

9. REFERENCES

- [1] Sleep cycle. <http://www.lexwarelabs.com/sleepcycle/>.
- [2] Sleep related problems. <http://www.sleepfoundation.org/articles/sleep-related-problems>.
- [3] Bed alarm sensor pad. http://www.1800wheelchair.com/asp/view-product.asp?product_id=3060.
- [4] Philips actiwatch. <http://www.actiwatch.respironics.com/>.
- [5] Wisp: Wireless identification and sensing platform. <http://www.seattle.intel-research.net/wisp/>.
- [6] Zeo personal sleep coach. <http://www.myzeo.com/>.
- [7] M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing daily activities with rfid-based sensors. In *Ubicomp*, 2009.
- [8] M. Carskadon and W. Dement. *Principles and Practice of Sleep Medicine*, chapter Normal Human Sleep: An Overview. W.B. Saunders, 1989.
- [9] R. Chaudhri, J. Lester, and G. Borriello. An rfid based system for monitoring free weight exercises. In *SenSys*, 2008.
- [10] J. Demsar, B. Zupan, G. Leban, and T. Curk. Orange: From experimental machine learning to interactive data mining. In *ECML/PKDD*, 2004.
- [11] H. V. der Loos, H. Kobayashi, G. Liu, Y. Tai, J. Ford, J. Norman, T. Tabata, and T. Osada. Unobtrusive vital signs monitoring from a multisensor bed sheet. In *RESNA*, 2001.
- [12] M. Ester, H. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, 1996.
- [13] F. Giganti, G. Ficca, S. Gori, and P. Salzarulo. Body movements during night sleep and their relationship with sleep stages are further modified in very old subjects. *Brain Research Bulletin*, 2008.
- [14] W. Karlen. *Adaptive Wake and Sleep Detection for Wearable Systems*. PhD thesis, Ecole Polytechnique Federale de Lausanne (EPFL), 2009.
- [15] K. Krauchi, C. Cajochen, and A. Wirz-Justice. Waking up properly: Is there a role of thermoregulation in sleep inertia? *Journal of Sleep Research*, 2004.
- [16] D. Mack, M. Alwan, B. Turner, P. Suratt, and R. Felder. A passive and portable system for monitoring heart rate and detecting sleep apnea and arousals: Preliminary validation. *Distributed Diagnosis and Home Healthcare*, 2006.
- [17] D. Mack, S. Kell, M. Alwan, B. Turner, and R. Felder. Non-invasive analysis of physiological signals (naps): A vibration sensor that passively detects heart and respiration rates as part of a sensor suite for medical monitoring. In *Summer Bioengineering Conference*, 2003.
- [18] A. Oksenberg and D. Silverberg. The effect of body posture on sleep-related breathing disorders: Facts and therapeutic implications. *Sleep Medicine Reviews*, 1998.
- [19] Y.-T. Peng, C.-Y. Lin, and M.-T. Sun. Multimodality sensors for sleep quality monitoring and logging. In *ICDEW*, 2006.
- [20] S. Redmond and C. Heneghan. Cardiorespiratory-based sleep staging in subjects with obstructive sleep apnea. *IEEE Transactions on Biomedical Engineering*, 2006.
- [21] A. Sadeh and C. Acebo. The role of actigraphy in sleep medicine. *Sleep Medicine Reviews*, 2002.
- [22] A. P. Sample, D. J. Yeager, P. S. Powladge, A. V. Mamishev, and J. R. Smith. Design of an rfid-based battery-free programmable sensing platform. In *IEEE Transactions on Instrumentation and Measurement*, 2008.
- [23] Z. Shinar, S. Akselrod, Y. Dagan, and A. Baharav. Autonomic changes during wake-sleep transition: A heart rate variability based approach. *Autonomic Neuroscience*, 2006.
- [24] A. Wood, J. Stankovic, G. Virone, L. Selavo, Z. He, Q. Cao, T. Doan, Y. Wu, L. Fang, and R. Stoleru. Context-aware wireless sensor networks for assisted-living and residential monitoring. *IEEE Network*, 2008.
- [25] Y. Yang, J. Shin, S. Jang, H. Lee, and Y. Yoon. Research of body movement during sleep with an infrared triangulation distance sensor, wavelets and neuro-fuzzy reasoning. In *IFMBE*, 2006.
- [26] D. Yeager, J. Holleman, R. Prasad, J. Smith, and B. Otis. Neuralwisp: A wirelessly powered neural interface with 1-m range. *IEEE Transactions on Biomedical Circuits and Systems*, To Appear.