

Detecting Human Movement by Differential Air Pressure Sensing in HVAC System Ductwork: An Exploration in Infrastructure Mediated Sensing

Shwetak N. Patel¹, Matthew S. Reynolds², and Gregory D. Abowd¹

¹College of Computing, School of Interactive Computing, & GVV Center
Georgia Institute of Technology
85 5th Street NW, Atlanta GA 30332-0280 USA
{shwetak, abowd}@cc.gatech.edu

²Department of Electrical and Computer Engineering
Duke University
Box 90291, Durham, NC 27708 USA
matt.reynolds@duke.edu

Abstract. We have developed an approach for whole-house gross movement and room transition detection through sensing at only one point in the home. We consider this system to be one member of an important new class of human activity monitoring approaches based on what we call infrastructure mediated sensing, or "home bus snooping." Our solution leverages the existing ductwork infrastructure of central heating, ventilation, and air conditioning (HVAC) systems found in many homes. Disruptions in airflow, caused by human inter-room movement, result in static pressure changes in the HVAC air handler unit. This is particularly apparent for room-to-room transitions and door open/close events involving full or partial blockage of doorways and thresholds. We detect and record this pressure variation from sensors mounted on the air filter and classify where certain movement events are occurring in the house, such as an adult walking through a particular doorway or the opening and closing of a particular door. In contrast to more complex distributed sensing approaches for motion detection in the home, our method requires the installation of only a single sensing unit (*i.e.*, an instrumented air filter) connected to an embedded or personal computer that performs the classification function. Preliminary results show we can classify unique transition events with up to 75-80% accuracy

1 Introduction and Motivation

The development of low-cost and easy-to-deploy sensing systems to support activity detection in the home has been an important trend in the pervasive computing community. Much of this research has centered on the deployment of a network of inexpensive sensors throughout the home, such as motion detectors or simple contact switches [23, 24, 26]. Although these solutions are cost-effective on an individual sensor basis, they are not without some important drawbacks that limit their desirability as research tools as well as their likelihood of eventual commercial success through broad consumer acceptance.

We have developed an approach that provides a whole-house solution for detecting gross movement and room transitions by sensing differential air pressure at a single point in the home. Our solution leverages the central heating, ventilation, and air conditioning (HVAC) systems found in many homes. The home forms a closed circuit for air circulation, where the HVAC system provides a centralized airflow source and therefore a convenient single monitoring point for the whole airflow circuit.

Disruptions in home airflow caused by human movement through the house, especially those caused by the blockage of doorways and thresholds, results in static pressure changes in the HVAC air handler unit when the HVAC is operating. Our system detects and records this pressure variation from differential sensors mounted on the air filter and classifies where exactly certain movement events are occurring in the house, such as an adult walking through a particular doorway or the opening and closing of a door. Preliminary results show we can classify unique transition events with up to 75-80% accuracy. We also show how we detect movement events when the HVAC is not operating.

The principal advantage of this approach, when compared to installing motion sensors throughout an entire house space, is that it requires the installation of only a single sensing unit (*i.e.*, an instrumented air filter) that connects to a computer. By observing the opening and closing of doors and the movement of people transitioning from room to room, the location and activity of people in the space can later be inferred. In addition, detecting a series of room transitions can be used for simple occupancy detection or to estimate a person's path in the house.

Because of the use of a single monitoring point on an existing home infrastructure (the HVAC air handler, in this example) to detect human activity throughout an entire house, we consider our system a member of an important new class of activity monitoring systems that we call ***infrastructure mediated sensing***. In the remainder of this paper, we further define this new category of sensing and explain the theory and implementation of the HVAC-facilitated motion detection.

2 Related Work

We distinguish between distributed direct sensing and a newly described category, infrastructure mediated sensing, which we informally call "home bus snooping" by analogy to computer network snooping. Distributed direct sensing involves the installation of a new sensing infrastructure into the home. This sensing infrastructure directly senses the presence, motion or activities of its residents through sensors that are physically located in each space where activity is occurring. Example systems include a new set of sensors and an associated sensor network (wired or wireless) to transfer the sensor data to a centralized monitoring system where sensor fusion and activity inference take place. In contrast, infrastructure mediated sensing leverages existing home infrastructure, such as the plumbing or electrical systems, to mediate the transduction of events. In these systems, infrastructure activity is used as a proxy for a human activity involving the infrastructure (see Figure 1). A primary goal of this second category of systems is to reduce economic, aesthetic, installation, and

maintenance barriers to adoption by reducing the cost and complexity of deploying and maintaining the activity sensing infrastructure.

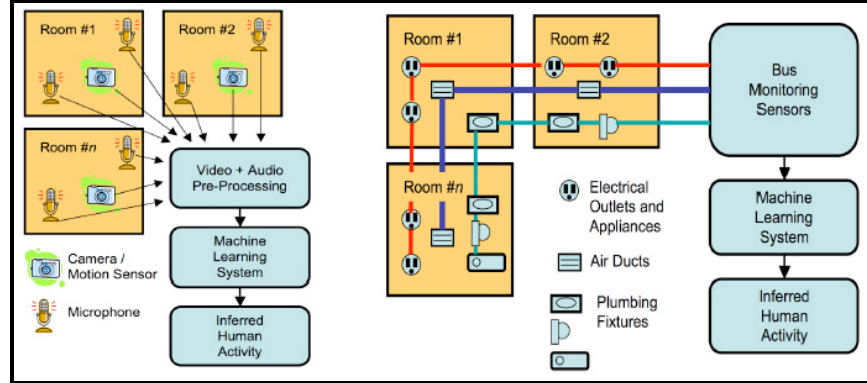


Figure 1: The distributed direct sensing (DDS) approach for activity detection and classification (left). The infrastructure mediated sensing approach for activity detection and classification (right).

2.1 Prior Work in Distributed Direct Sensing

Most of the existing literature in human activity sensing in the home falls into the distributed direct sensing category. In the pervasive computing research context, commonly used sensors for detecting human activity in the home include high-fidelity sensors such as visible light and IR cameras [25, 28] or microphones [4], as well as low-fidelity sensors such as passive infrared (PIR) motion detectors [27] and floor weight sensors [18]. High-fidelity distributed direct sensing has a long history of use in activity detection and classification research, primarily focused on computer vision or machine learning systems that capture the movement of people in spaces [13]. For example, Chen *et al.* installed microphones in a bathroom to sense activities such as showering, toileting, and hand washing [6]. The use of these high fidelity sensors in certain spaces often raises concerns about the balance between value-added services and acceptable surveillance, particularly in home settings [5, 9, 10]. Low-fidelity, distributed direct sensing work includes the use of a large collection of simple, low-cost sensors, such as motion detectors, pressure mats, break beam sensors, and contact switches, to determine activity and movement [23, 24, 26]. The principal advantages are lower per-sensor cost and reduced privacy concerns.

All distributed direct sensing approaches share the advantages and disadvantages of placing each sensor in close proximity to where human activity occurs. For example, commonly used cameras or PIR sensors require a clear line of sight to the desired room coverage area; the person being sensed will be able to see the camera or PIR sensor. Generally, cameras or PIR sensors are deployed in places that have adverse aesthetics, such as on walls, on ceilings, or above a door [7, 9]. The large number of sensors required for coverage of an entire building presents an inherent complexity hurdle. Installation and maintenance of (typically) tens of sensors in a

home, or hundreds to thousands of sensors in a larger building such as a hotel, hospital, or assisted living facility, results in high labor costs during installation, and an ongoing maintenance and sensor network management challenge during routine operation.

It is often difficult to balance the value of in-home sensing and the complexity of the sensing infrastructure. One example that illustrates this difficulty is the Digital Family Portrait system, a peace of mind application for communicating well-being information from an elderly person's home to a remote caregiver [21]. In the system's deployment study, movement data was gathered from a collection of strain sensors attached to the underside of the first floor of an elder's home. The installation of these sensors was difficult, time-consuming, and required direct access to the underside of the floor. Though the value of the application was proven, the complexity of the sensing limited the number of homes in which the system could be easily deployed.

2.2 Infrastructure Mediated Sensing

Some recent innovative work in the infrastructure mediated sensing category leverages the existing infrastructure in a home to collect signals at a single location. A few researchers have recently begun exploring the use of existing home infrastructure to detect human originated events [8, 19, 20]. A few microphones on the plumbing infrastructure in the basement of a home can infer basic activities, such as bathing or washing dishes, through acoustically-transduced signals [8]. A single plug-in sensor can classify events, such as the actuation of a light switch, through the analysis of noise, transduced along the power line, from the switching and operation of electrical devices [19]. These two approaches cover a complementary set of human activities, depending on whether a water- or power-related event precedes that activity.

Both of these approaches require human-initiated events, as identified through signals carried via the infrastructure of their corresponding resources, in order to provide human activity information. They ignore activities that do not include the use of the plumbing or electrical systems, such as movement and transitions between parts of the home. In the case of water event detection, there may be only a few water usage events per person per day, whereas with electrical event detection, there may be limited electrical actuations during the day when incoming sunlight illumination may result in reduced light switch use. This results in a relatively sparse activity dataset compared to a dataset obtained using a dense network of PIR motion sensors located throughout the home. Therefore, we were motivated to find an infrastructure mediated sensing technique that delivers movement information.

We contrast infrastructure mediated sensing with a "piggybacking" approach that simply reuses an existing sensing infrastructure in the home that may be present for other purposes. For example, ADT Security System's QuietCare [1] offers a peace of mind service that gathers activity data from the security system's PIR motion detectors. Although a promising approach, security motion sensors are typically only installed in a few locations in the home, primarily on the ground floors, resulting in a much sparser dataset than is needed for general activity recognition.




Figure 2: We instrument a standard HVAC air filter with pressure sensors that are able to detect airflow in both directions. The air filter is then installed in the HVAC's air handler unit.

3 Our Approach and System Details

We instrumented an HVAC's air filter with five pressure sensor units, each sensing in both directions (see Figure 2). The sensors do not interfere with the operation of the air filter or HVAC and instrumenting the air filter allows for easy installation in standard HVAC units. The sensors on the air filter capture the pressure differential across the filter in the air handler chamber. The magnitude of the pressure change across all the sensors is used to identify unique disruptions in airflow in the physical space. Machine learning techniques then classify these disruption signatures.




Figure 3: Cross section of a HVAC air handler unit.

3.1 Theory of Operation

The HVAC system's air handler is a device used to circulate conditioned air throughout a space. Typically, an air handler is a large, sealed metal box containing a blower, heating/cooling coils, filter, and dampers (see Figure 3). An air handler consists of a discharge, or supply, chamber where the conditioned air exits through ductwork, and is drawn back into the return chamber through a separate set of vents and ductwork. During its operation, a pressure differential, ΔP , is built up in the blower chamber, known as the total static pressure. The static pressure is a measure of resistance imposed on the HVAC's blower in the air handler. The static pressure is affected by a variety of factors that impede the airflow between the supply and return. These include the length of ducting, number of fittings used in the ductwork, closed air vents, or dirty air filters. When installing an HVAC unit, a technician usually takes care in properly balancing the static pressure to ensure its proper operation. This includes installing sufficient supply and return ductwork in the right locations. Technicians also install ductwork to various rooms to ensure effective coverage. Figure 4 shows a cross-sectional drawing of a home and example locations of the supply and return vents and the potential airflow paths.

When the HVAC is running, air flows from the supply vents to the return vents through the conditioned space (*i.e.*, a room). There is always some airflow from each supply vent to all the return vents. Depending on the location of the vents, the airflow paths and amount of airflow can vary. When there is disruption to the airflow, there is a change in the static pressure in the air handler as a result of the resistance in the airflow. Depending on the location of return vents, a disruption in airflow can cause a more persistent change in the overall static pressure, such as from a direct blockage of a return vent. In a home, one contributor to this airflow disruption is doorways, where airflow can either be disrupted by the closing or opening of a door or the partial blockage of an adult passing through the threshold. Sometimes, an individual may even feel the “resistance” from the airflow when trying to open a door. Also depending on the location in the house where this disruption is occurring, the “resistance” differs because of the airflow path. Another way to look at this phenomenon is using an electrical circuit analogy (see Figure 4).

When the HVAC is not in operation, the ductwork acts as a “wave guide.” Significant airflow produced in the space flows through the ductwork. Although small movements cannot generate enough airflow, the movements of large surfaces, such as doors, can produce detectable amounts of airflow through the air handler. Thus, there are opportunities to detect certain movement in the space with the HVAC both in operation and not in operation.

We use the air filter chamber as the sensing point for two important reasons. First, it is between the supply and return chambers and near the blower assembly, making it a good place for recording the static pressure changes. Second, the filtration unit typically has the easiest access to the air handler, potentially making it easy-to-deploy for installers and end-users. The static air pressure is determined by installing pressure sensors facing each direction on the air filter and calculating the differential (ΔP). A single differential pressure sensor would also be appropriate. However, using two pressure sensors makes their placement easier. This is because typical differential pressure sensors have the pressure ports on one side, which requires routing an air

tube to the other side. The sensors required for our approach are capable of measuring up to 2 bars of pressure and sensitive enough to measure small pressure changes down to .1 mbar. Figure 5 shows a graph of the change in static pressure as a door is opened and closed. There is an initial spike in the pressure followed by a flattening. After the door is reopened, the pressure returns to the previous state.

We placed multiple pressure sensors on the air filter to help estimate the location of the resulting pressure change. In standard ductwork, multiple ducts combine to feed larger trunks, which then attach to the supply and return chambers. Because multiple ducts feed into the chambers, pressure sensors closer to the ductwork that is contributing to the airflow disruptions will see greater initial change in pressure compared to the other sensors.




Figure 4: Diagram of airflow from return and supply ducting in a home (top). Electrical circuit diagram analogy of our sensing approach (bottom).




Figure 5: Examples of the pressure changes in the air handler as a result of an opening and closing of a door (left) and an adult walking through two different doorways (right).

3.2 Data Collection Hardware and Software

We used the Intersema MS5536 piezoresistive pressure sensor module for building our sensor units. The MS5536 modules are high resolution (.1 mbar), provide a stable output of up to 2 bars, and have a maximum rating of up to 5 bars, which is sufficient for many residential HVAC applications. The modules incorporate a temperature sensor for proper pressure compensation, a built-in 15-bit ADC, and also provide easy communication using SPI. To obtain pressure differentials, the MS5536 uses two sensors facing opposite directions. The pressure sensor modules are connected to an ATMEL microcontroller (see Figure 6). The microcontroller samples the pressure and temperature sensors on the MS5536 and calculates a temperature-compensated pressure value every 35 milliseconds. Intersema's temperature compensation formula was used in our calculations [12]. The pressure values are then transferred to a PC via a USB connection. Multiple sensor units are connected to a single PC using a USB hub. We chose to use individual units to give us some flexibility when experimenting with a variety of sensor placements on the air filter. The sensor units are small enough to attach easily to the air filter with zip ties. In a production version, the sensors would be mounted on a framed bracket that would just attach to the air filter. A fully deployable unit would have all the pressure sensors feeding into it a single microcontroller. A unit incorporating five differential pressure sensors costs about \$100 USD at low volumes.

The software used in our data collection is written in C++ and records the temperature-compensated pressure data, the raw pressure values, and the temperature from the sensors units. The application continually timestamps and records the pressure-related data from all the sensor units every 50 millisecond.




Figure 6: Block diagram of our pressure sensor unit.

3.3 Detecting Door Opening and Closing Events

We observed two important features that were characteristic of door opening and closing events. When a door is closed, there is first an initial abrupt change in static pressure (change in ΔP) followed by persistent change until the door is reopened (see Figure 5). After opening the door, the static pressure gradually drops to the previous state. We detect this phenomenon by first looking for a significant change in the static pressure by at least one of the five sensing units. We do this by comparing the average of the 5 previous pressure differential reading with the current. When there is a pressure change greater than 10 mbar, we record the subsequent pressure values for further processing until there are no more changes for a period of 4000 ms. All other sensors also record at the same time. The 10 mbar threshold is to avoid detecting any slight variations from the sensor or noise from the ADC. From the recorded data we next extract the initial pressure value, the initial maximum pressure change, and the resulting final stable pressure. These features are extracted for all 5 sensor units, producing a final feature vector of 15 components.

3.4 Detecting Movement of People through Doorways

A person passing through a doorway is a brief event, and the size of the individual can vary, decreasing the likelihood of detection. However, we still wanted to explore the feasibility of detecting those events. During our experimentation, we observed variations in the static pressure as individuals moved through various doorways. Unlike the door events, the changes in pressure are very short-lived. There is a slight change in the static pressure and then the pressure settles back to its original state. The effect is dependent on the location of the supply and return vents relative to the doorway and the ratio of the size of the person to the size of the doorway. From our observations, a ratio of 1:3 resulted in detectable airflow disruptions (>10 mbar).

We isolated these events by comparing the average of the 5 prior pressure differentials to the current. We recorded the pressure values when there was a change of more than 10 mbar by at least one sensor unit. All other sensor units also triggered to record at the same time. Values were gathered until the pressure stabilized. We use the maximum pressure change from each of the 5 sensor units as the feature vector.

3.5 Detecting Door Transition Events with HVAC Off

When the HVAC is not operating, there is no static pressure build-up in the air handler. Instead, the pressure is equal to the atmospheric pressure of approximately 1 bar. Any significant airflow generated in the conditioned space is guided through either the supply or return ducts and eventually reaches the sensor units on the filters. The sensitivity of the sensor units make it possible to detect airflow reaching the sensors. We can use the pressure values from both sides of the filter to help determine where the airflow originated. Similar to the previous approaches, we also use the multiple sensing points to help localize the origination of the induced airflow. Theoretically, it is also possible to detect airflow caused by people moving near an air vent and by other devices, such as a ceiling or desk fan. However, these events produce very small amounts of airflow and require more expensive, high-resolution and low-noise pressure sensors. In this case, we focus on just the movement of doors when the HVAC is not operating.

When the HVAC is off, we isolate door events by comparing the average of the 5 prior pressure differentials to the current. We then record the pressure values when there are any changes of more than 10 mbar by at least one sensor unit. All other sensor units are recorded at the same time. Values are gathered until the pressure stabilizes, and the feature vector of the maximum pressure change from each of the 5 sensor units is calculated.

3.6 Classifying Events

For our classification scheme, we used support vector machines (SVMs). SVMs perform classification by constructing an N-dimensional hyperplane that optimally separates the data into multiple categories. The separation is chosen to have the largest distance from the hyperplane to the nearest positive and negative examples. Thus, the classification is appropriate for testing data that is near, but not identical, to the training data as is the case for the feature vectors in our approach. In addition, SVMs can automatically determine the appropriate kernel type based on the data build characteristics, so kernels beyond linear functions can be factored in. For our experiments, we created three different SVM models for each of the three scenarios, using their respective feature vectors with each transition event labeled as the class. The open transition and the close transition for each door of interest were used as the classes in the learner. This was the case for both the HVAC in operation and not in operation. In the case of classifying human movement through a doorway, we do not differentiate between the directions of movement, thus the class labels were of the door where the movement occurred.

4 Feasibility Experiments

The goal of the feasibility experiments was to determine if and how often we could detect transition movements (*e.g.*, adults walking through doorways and the opening and closing of doors) and how accurately we could classify unique transition events.

In this section, we present the results from experiments in four different homes for the following three conditions: opening and closing of doors while the HVAC is in operation, adults moving through doorways while the HVAC is in operation, and the opening and closing of doors while the HVAC is not in operation.

4.1 Setup of Feasibility Experiments

We conducted experiments and observations in four different homes for a period ranging from 3 to 4 weeks (see Table 1). Home 1 and Home 2 were fairly large homes, with Home 1 having three separate central HVAC units, and Home 2 having two separate central HVAC units. We instrumented all three units in Home 1 and one unit in Home 2. Homes 3 and 4 were smaller apartments with a single, central HVAC system. Thus, we evaluated a total of six different spaces and HVAC units. For each HVAC unit, we installed an instrumented air filter (see Figure 2). The sensors were securely attached to prevent any movement from the airflow. The cables were run around the edge of the filter to prevent them from being drawn in to the fan assembly. Finally, the cables were connected to a laptop placed near the HVAC unit.

We used two techniques for obtaining labeled ground truth data. First, throughout the 3-4 week period we manually labeled numerous door close and open events and a person walking through doorways with the house in a closed and sealed state (windows and exterior doors closed). Second, we captured data for a longer time period using motion sensors placed at various locations in the house. Sensors on both sides of the top of the doorways (facing downwards) detected the direction of movement through the doorway. Although we were not able to accurately differentiate door movement and people movement, the motion sensors did allow us to determine if any transition events occurred at various times during the day. The large dataset allowed us to partition the data into sufficient training and test sets.

Table 1: Descriptions of the homes in which our system was tested. The deployment lasted approximately 3-4 weeks.

Home	Year Built	No. of HVAC Units Tested	Floors/ Total Size (Sq Ft)/ (Sq M)	Style/ No. of occupants	Bedrooms/ Bathrooms/ Total Rms./ Doorways considered	Deploy Length (weeks)
1	2003	3	3/4000/371	1 Family Home/3	4/4/13/20	4
2	2001	1	1/1600/149	1 Family Home/5	3/2/7/10	3
3	1997	1	1/700/58	1 Bed Apt/2	1/1/5/5	3
4	1986	1	1/500/46	1 Bed Studio/1	1/1/3/4	4

Table 2: Performance results of our manually-labeled experiments with the HVAC in operation. The accuracies are shown using 10-fold cross validation.

Home/ Floor	No. of Doorways Tested	No. of Door Instances/People Instances	Door Majority Classif. (%)	Door Classif. Accuracy (%)	People Majority Classif. (%)	People Classif. Accuracy (%)
1/1	5	375/375	21	84	23	72
1/2	4	300/300	18	61	18	42
1/3	11	825/600	9	77	12	61
2	10	750/400	8	73	10	63
3	5	375/375	20	74	20	70
4	4	300/300	26	81	25	76

Table 3: Confusion matrix of the classification results from the controlled experiments in Home 1/3 (HVAC in operation). D1 - D11 represent each doorway.

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D11
D1	<u>72</u>	0	0	0	1	0	0	1	0	0	1
D2	1	<u>57</u>	0	2	0	2	6	4	0	1	2
D3	0	1	<u>60</u>	1	0	1	3	2	5	2	0
D4	0	0	1	<u>57</u>	2	0	0	4	3	6	2
D5	4	0	1	4	<u>52</u>	5	0	6	2	0	1
D6	5	1	0	0	6	<u>53</u>	4	2	0	1	3
D7	0	2	3	3	0	1	<u>61</u>	0	3	2	0
D8	6	0	0	0	2	1	1	<u>55</u>	5	0	5
D9	1	0	4	0	1	5	2	0	<u>59</u>	2	1
D10	2	2	7	0	3	3	8	0	2	<u>43</u>	5
D11	0	1	0	0	0	2	0	0	0	2	<u>70</u>

Table 4: Performance results of our manually labeled door open/close events for when the HVAC is not in operation.

Home/ Floor	No. of Doorways Tested	No. of Door Event Instances	Door Majority Classif. (%)	Door Classif. Accuracy (%)
1/1	5	125	20	66
1/2	4	100	25	47
1/3	11	275	9	64
2	10	250	10	69
3	5	125	20	71
4	4	100	25	68

4.2 Manually-labeled Controlled Experiments

In these experiments, we wanted to test the feasibility of accurately classifying the various kinds of unique door or movement events in a quasi-controlled manner. For all four homes, we manually labeled sensor readings for each event using a remote handheld computer wirelessly connected to the data collection PC. We were able to accurately label the sensor readings for each of the five sensors after triggering the various events. We then used our feature extraction algorithms to construct the appropriate feature vectors to feed our classifier. For these experiments, all interior doors of interest were kept in the open position (90 degrees from the opening), while we were manually opening and closing each door. For the human movement experiments, the same individual triggered those events. We collected 25 instances for each of the doorway events three different times during the 3-4 week period (175 instances).

Table 2 shows the classification accuracies of all the spaces. We have also included an example confusion matrix (Table 3). It is clear that door transition events were more accurate than people transitions. However, the overall accuracy of classifying unique movement events was around 65%, which is still promising. Door events were classified correctly on an average of 75-80% of the time, suggesting that we can combine both of these events to provide good predictions on the location or movement of people through the space. Some of the low classification accuracies, such as from Floor 2 in Home 1, were attributed to the lack of door and doorways. That space was very open with the air vent a significant distance away from the interior doors. The results of the HVAC off experiment also showed some promising results (see Table 4). Although the accuracies are lower than with the HVAC in operation, there is still some predictive power. The higher performance came in smaller spaces where the vents tended to be closer to the doorway and in spaces where there were many vents, such as Homes 1 and 2.

4.3 Long-term Deployment

For the long-term deployment, we wanted to gather *in situ* or “more natural” data on the various events occurring in the home and provide some initial long-term *in situ* results for our sensing approach. For labeling, we used motion sensors placed at various doorways to determine any door movement or motion through the doorway and matched those events up with the corresponding sensor values from the HVAC. We conducted two analyses; one was the percentage of time we were able to determine particular events with our system and the second was to determine the classification accuracies of detecting unique events. Table 5 shows the number of events that were detected by our system, either as a door transition event or a human movement through the doorway, for each of the 4 homes. We present the results for two cases. One is with the HVAC in operation and the other is with the HVAC off.

Table 5: The percentage of events that our approach was able to detect. This is determined by comparing the number of detected events to the number of doorway events gathered by the motion sensors. These results include events detected with HVAC both on and off.

Home/ Floor	No. of Doorways Tested	No. of Total Motion Sensor Events	No. of Total Detected Events	HVAC On: Detected Events (%)	HVAC Off: Detected Events (%)
1/1	5	53	48	91	68
1/2	4	94	60	64	35
1/3	11	238	195	82	73
2	10	467	334	72	64
3	5	245	198	81	70
4	4	61	51	84	77

The results show that a larger percentage of events were detected with the HVAC in operation than with it in the off state. The reason for the lower percentage for the HVAC off case was because of the location of the return and supply vents. In some cases, the vents were not close enough to a door for the airflow to reach the sensing units, which we saw in our controlled experiment. The smaller spaces and the spaces with many doorways actually resulted in a higher number of detectable events. This is attributed to the greater number of vents and the likelihood that the doorways were near vents. The results with the HVAC in operation are promising, with almost 80% of the events being detected when compared to the motion sensors. Table 6 shows the results of classifying unique events in the house. We applied our SVM classification scheme to the entire *in situ* dataset for each of the 4 homes (6 spaces). This dataset included events from all three of the possible conditions (door open/close with HVAC on and off and human movement with HVAC on). The triggering of the motion sensor was used to provide the location label to the air pressure data collected by our sensing system. Because we did not know the type of event, we used the signal response to determine the event (*i.e.*, person or door).

We report the accuracy of our approach using 10-fold cross validation across the entire data set. Compared to the first controlled experiments, the overall accuracy on average is 15-20% lower. However, considering that we did not control the various other events occurring during that time, the results are still promising with classification accuracies between 60-70%. From these we can see that the status of other doors did not have a large impact on the classification accuracy of detecting door transitions with the HVAC off. The larger difference while the HVAC is in operation compared to the controlled experiment does indicate the door states have an impact on the pressure differentials, as expected. However, since we trained from a subset of the entire dataset, the learner seemed to incorporate the various door combinations. This is intuitive because people tend to be consistent with how they leave many of their doors throughout the day, while only actually using a few doors.

Table 6: The performance of using our learning approach to the data from the long-term deployment. The motion sensor data was used to label each event, so the dataset consists of *in situ* event instances. The accuracies are show using 10-fold cross validation

Home/ Floor	No. of Doorways Tested	No. of Doorway Transition Instances	Door Majority Classif. (%)	Door Classif. Accuracy (%)	People Majority Classif. (%)	People Classif. Accuracy (%)
1/1	5	48	26	65	28	61
1/2	4	60	26	53	26	42
1/3	11	195	14	72	17	63
2	10	334	19	62	12	65
3	5	198	28	72	23	71
4	4	51	34	78	38	81

5 Deployability: Prevalence of Central HVAC Systems

Although central home HVAC systems are not as prevalent in some geographic regions as plumbing or electrical infrastructure, our approach is still useful in the significant number of homes or buildings that do have central HVAC. Because central HVACs are more efficient than using a collection of window units [16], the upward trend in energy cost has driven the use of central HVAC systems to a growing number of homes. In 1997, 66% of the homes in the United States and Canada were reported to have central HVAC, and its prevalence is growing at a fast rate [3, 15, 22]. In addition, nearly all new homes built in the southern part of the U.S. and 80% in the rest of the U.S. and Canada have central HVAC installed during construction [15]. Europe and Australia show a similar trend, with approximately 55% homes using central HVAC [14, 11]. However, in some Asian counties such as Japan and Korea, central HVAC is not as common in homes because of the smaller dwelling sizes prevalent in those regions. If the home is very small, such as a small Japanese or Korean home, the deployment of distributed direct sensors may not be as arduous because of the smaller amount of floor space to cover. Regardless of the regional prevalence of central HVAC, the value of our approach becomes more apparent in larger homes or in assisted living facilities that have many rooms, precisely the settings where installing many distributed sensors is economically unattractive.

HVAC systems will probably increase in prevalence because they can provide more functionality than just heating and cooling. Recent EnergyStar reports have shown that running the HVAC for longer periods of time, but using alternate conditioning features, such an air-to-air exchanger, is more energy efficient [16]. This EnergyStar report also recommends that HVAC units incorporate whole house HEPA filtration. Construction codes, such as for hospitals and assistive care facilities, also have a minimum air movement requirements to ensure proper filtration [2, 17]. All of these factors increase the motivation for having the HVAC in operation, increasing the effectiveness of our sensing approach. If we take a standard 2-ton (24,000 BTU) HVAC unit and run the air handler’s fan continuously for an entire month it would cost about \$6 US (assuming an electricity price of \$0.05 US per 1 kW-h), which would need to be balanced against any value-added capability our sensing provides.

6 Discussion of Limitations and Potential Improvements

Our approach is certainly not without limitation. It does require a training phase and further research is still needed in coming up with a mechanism to ease the training process. Some possible directions are to use events generated from other calibrated systems (water line or power line) to feed the training of this system. Although this might not cover all possible training cases, it can be used to relieve some of the burden. Those systems can also provide continual feedback for verifying the training set. In addition, partial training may also be feasible for certain applications, where only certain doorways are first trained. Then, if there is any interest in observing other events, the training can occur after the fact and the other past events can be reviewed.

We considered only the amplitude of the static pressure change and using multiple pressure sensor units to determine unique movement and door events. Other possible approaches would look at the changes in the laminar airflow. Although we use the temperature values for calculating the temperature-compensated pressure values, we could use the temperature reading as an additional feature. Our current focus was on residential central HVAC systems, but our system can scale reasonably to larger units used in most commercial buildings. Further investigations are needed to explore those systems. Our feasibility experiments did not directly factor in the opening and closing windows and doors. Finally, our current approach does not directly address compound events—multiple simultaneous door and person movements—although these events occurred in the long-term studies. Modeling airflow variations and creating a new learning approach that incorporates that domain knowledge could address this.

7 Conclusion

We have developed an approach for whole-house gross movement and room transition detection through sensing at only one point in the home. We consider this system to be one member of an important new class of human activity monitoring approaches based on infrastructure mediated sensing, or "home bus snooping." Our solution leverages the existing ductwork infrastructure of central heating, ventilation, and air conditioning (HVAC) systems found in many homes. Disruptions in airflow caused by human inter-room movement result in static pressure changes in the HVAC air handler unit. This is particularly apparent for room-to-room transitions and door open/close events involving partial blockage of doorways and thresholds. We detect and record this pressure variation from sensors mounted on the air filter and classify where certain movement events are occurring in the house, such as an adult walking through a particular doorway or the opening and closing of a door. Although less precise, we also show the detection of movement when the HVAC is not operating. In contrast to more complex distributed sensing approaches for motion detection in the home, our method requires the installation of only a single sensing unit.

The combination of different types of infrastructure mediated sensors offers a number of attractive properties for deployment of useful applications in the home. For example, the combination of detecting human-initiated electrical [19] or water events [8] with our work on movement detection through airflow sensing enables a variety of

new approaches for integrating energy and environmental conservation with ordinary human activities in the home. A system could alert an individual that he or she should attend to an energy or environmental conservation task, such as turning off an unneeded light or a running faucet, when the system detects that he or she is near that part of the house. The combination of electrical event detection and airflow detected movement information can also provide important correlation data for energy conservation applications by relating a person's usage of the physical space with the usage of electrical devices. One could design an energy-efficient zoned HVAC unit that selectively heats or cools each zone on the basis of activity information passively sensed through the HVAC system itself, which would offer a tremendous installation and maintenance cost benefit over competing distributed sensing approaches.

Acknowledgments

This work was sponsored in part by the National Science Foundation Graduate Fellowship and the Intel Research Council. The authors would also like to thank the members of the Ubicomp Research Group at Georgia Tech, in particular Julie Kientz and Lana Yarosh.

References

1. ADT QuietCare. <http://www.adt.com/quietcare/>. 2008.
2. American Institute of Architects. Guidelines for Design and Construction of Hospital and Health Care Facilities. The American Institute of Architects Press, Washington D.C. 2001.
3. Barnes Reports. 2008 U.S. Plumbing & Heating & A/C Contractors Report. Oct 2007.
4. Bian, X., Abowd, G.D., and Rehg, J.M. Using Sound Source Localization in a Home Environment. In Proc. of the Pervasive 2005. pp 19-26. 2005.
5. Beckmann, C., Consolvo, S., and LaMarca, A. Some Assembly Required: Supporting End-User Sensor Installation in Domestic Ubiquitous Computing Environments. In the Proc. of Ubicomp 2004. pp. 383-399. 2004.
6. Chen, J., Kam, A.H., Zhang, J., Liu, N., Shue, L. Bathroom Activity Monitoring Based on Sound. In the Proc. of Pervasive 2005. pp. 47-61. 2005.
7. Chetty, M., Sung, J., Grinter, R.E. How Smart Homes Learn: The Evolution of the Networked Home and Household. In the Proc. of Ubicomp 2007. pp. 127-144. 2007.
8. Fogarty, J., Au, C., and Hudson, S.E. Sensing from the Basement: A Feasibility Study of Unobtrusive and Low-Cost Home Activity Recognition. In the Proc. of UIST 2006. pp. 91-100. 2006.
9. Hirsch, T., Forlizzi, J., Hyder, E., Goetz, J., Kurtz, C. and Stroback, J. The ELDer Project: Social, Emotional, and Environmental Factors in the Design of Eldercare Technologies. In the Proc. of the ACM Conference on Universal Usability. pp 72-79. 2000.
10. Iachello, G and Abowd, G.D. Privacy and Proportionality: Adapting Legal Evaluation Techniques to Inform Design in Ubiquitous Computing. In the Proc. of CHI 2005. pp. 91-100. 2005.

11. IBISWorld. AC and Heating Services in Australia-Industry Market Research Report. Aug 2007.
12. Intersema. <http://www.intersema.com/site/technical/ms5536.php>. 2008.
13. Koile, K., Tollmar, K., Demirdjian, D., Howard, S., and Trevor, D. Activity Zones for Context-Aware Computing. In the Proc. of UbiComp 2003. pp. 90-106. 2003.
14. Market and Bus. Development. UK Domestic Central Heating Market Development. Sep 2007.
15. Menzer, Mark. Heat Pump Status and Trends in North America. IEA Heat Pump Conference. http://www.ari.org/research/engineering_research/. May 31, 1999.
16. Nadel, S. Increasing Appliance Energy Savings by Looking Beyond the Current Energy Star. ACEEE 2004 Energy Star Appliance Partner Meeting. http://www.energystar.gov/ia/partners/downloads/Plenary_B_Steve_Nadel.pdf 2004.
17. Ninomura, P. and Bartley, J. New Ventilation Guidelines For Health Care Facilities. Air Conditioning and Refrigeration Journal. July-September Issue. 2002.
18. Orr, R.J. and Abowd, G.D. The Smart Floor: A Mechanism for Natural User Identification and Tracking. In Proc. of the Extended Abstracts of CHI 2000. pp. 275-276. 2000.
19. Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S., Abowd, G.D. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. In the Proc. of UbiComp 2007. pp. 271-288. 2007.
20. Patel, S.N., Truong, K.N., and Abowd, G.D. PowerLine Positioning: A Practical Sub-Room-Level Indoor Location System for Domestic Use. In the Proc. of UbiComp 2006. pp. 441-458. 2006.
21. Rowan, J. and Mynatt, E.D. Digital Family Portrait Field Trial: Support for Aging in Place. In Proc. of CHI 2005. pp 521-530. 2005.
22. Supplier Relations US, LLC. Ventilation, Heating, Air-Conditioning, and Commercial Refrigeration Equipment Manufacturing Industry in the U.S. and its Foreign Trade. August 2007.
23. Tapia, E.M., Intille, S.S., and Larson, K. Activity recognition in the home setting using simple and ubiquitous sensors. In Proc. of Pervasive 2004. pp. 158-175. 2004.
24. Tapia, E.M., Intille, S.S., Lopez, L., and Larson, K. The design of a portable kit of wireless sensors for naturalistic data collection. In Proc. of Pervasive 2006. pp. 117-134. 2006.
25. Vicon MX. <http://www.vicon.com/products/systems.html>. 2008.
26. Wilson, D.H. and Atkeson, C.G. Simultaneous Tracking and Activity Recognition (STAR) Using Many Anonymous, Binary Sensors. In Proc of Pervasive 2005. pp 62-79. 2005.
27. Wren, C.R. and Munguia-Tapia, E. Toward Scalable Activity Recognition for Sensor Networks. In the Proc. of the International Workshop in Location and Context-Awareness (LoCA 2006). Pp. 168-185. 2006.
28. Yang, Z and Bobick, A.F. Visual Integration from Multiple Cameras. In the Proc. of Application of Computer Vision, WACV/MOTIONS 2005. pp. 488-493. 2005.