

HydroSense: Infrastructure-Mediated Single-Point Sensing of Whole-Home Water Activity

Jon Froehlich¹, Eric Larson², Tim Campbell³, Conor Haggerty⁴, James Fogarty¹, Shwetak N. Patel^{1,2}

¹Computer Science & Engineering, ²Electrical Engineering,
³Mechanical Engineering, ⁴Community, Environment, and Planning
DUB Institute, University of Washington
Seattle, WA 98195

{ jfroehli@cs, eclarson@u, tcampbl@u, conorh@u, jfogarty@cs, shwetak@cs } .washington.edu

ABSTRACT

Recent work has examined infrastructure-mediated sensing as a practical, low-cost, and unobtrusive approach to sensing human activity in the physical world. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed by their manifestations in an environment's existing infrastructures (e.g., a home's water, electrical, and HVAC infrastructures). This paper presents HydroSense, a low-cost and easily-installed single-point sensor of pressure within a home's water infrastructure. HydroSense supports both identification of activity at individual water fixtures within a home (e.g., a particular toilet, a kitchen sink, a particular shower) as well as estimation of the amount of water being used at each fixture. We evaluate our approach using data collected in ten homes. Our algorithms successfully identify fixture events with 97.9% aggregate accuracy and can estimate water usage with error rates that are comparable to empirical studies of traditional utility-supplied water meters. Our results both validate our approach and provide a basis for future improvements.

Author Keywords

Infrastructure-mediated sensing, activity sensing.

ACM Classification Keywords

H5.2. Information Interfaces and Presentation: User Interfaces;
H1.2. Models and Principles: User/Machine Systems.

General Terms

Algorithms, Experimentation, Measurement

INTRODUCTION AND MOTIVATION

Effective methods for sensing and modeling human activity in the physical world are a cornerstone of ubiquitous computing research and practice. Many complementary approaches have been developed, including recent interest

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UbiComp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.
Copyright 2009 ACM 978-1-60558-431-7/09/09...\$10.00.

in *infrastructure-mediated sensing* [7, 14, 15, 16, 17, 20]. This approach is based on the idea that human activities (e.g., running a dishwasher, turning on a reading light, or walking through a doorway) can be sensed via their manifestations in an environment's existing infrastructures (a home's water [7], electrical [15, 16, 17, 20], and HVAC [14] infrastructures). Because of its practical, low-cost, and unobtrusive nature, infrastructure-mediated sensing offers significant promise as a general method.

This work focuses on infrastructure-mediated sensing of home water activity. Water is essential to many home activities (e.g., washing, cleaning, cooking, drinking, gardening) which are in turn central to important potential ubiquitous computing applications (e.g., helping elders live more independently, helping people monitor their own water usage to reduce waste). Previous work monitoring home water usage [6, 7, 9] required multiple sensing points, exposed piping, could not infer *both* fixture and flow, and received limited or no validation in actual homes.

This paper presents HydroSense, a low-cost, *single-point* solution for activity sensing mediated by a home's existing water infrastructure. HydroSense is based on continuous analysis of *pressure* within a home's water infrastructure. Specifically, we *identify individual water fixtures* (e.g., a particular toilet, a kitchen sink, a particular shower) within a home according to the unique *pressure waves* that propagate to the sensor when valves are opened or closed. We also estimate the *amount of water being used* at a fixture based on the magnitude of the resulting *pressure drop* within the water infrastructure. Our work represents a significant advance over prior research in several regards:

First, HydroSense can be *easily installed* at any accessible location within a home's existing water infrastructure. Typical installations will be at an exterior hose bib, utility sink spigot, or water heater drain valve. If unavailable or not easily accessed (e.g., in an apartment unit), HydroSense can also be installed at the water connection point for a dishwasher, clothes washer, or toilet. All of these are simple screw-on installation points, with no need for a plumber.

Second, HydroSense's analysis of *pressure* provides the unique capability of sensing both the *individual fixture* at




Figure 1: An illustrative schematic of a basic plumbing layout in a two-bathroom home. HydroSense can be easily installed at any accessible location in a home’s water infrastructure, with typical installations at an exterior hose bib (shown above), a utility sink spigot, or a water heater drain valve. By continuously sensing water *pressure* at this single installation point, HydroSense can both *identify individual fixtures* at which water is being used as well as estimate the *amount of water being used*.

which water is currently being used as well as an estimate of the *amount of water being used*. HydroSense is the first practical approach to enabling applications that require both. Our sensing of pressure is also less susceptible to ambient noise, as has been encountered in previous microphone-based infrastructure-mediated systems.

Third, we evaluate HydroSense in ten very diverse homes, thus providing a more *robust evaluation* than any previous work on water-related home activity sensing. We demonstrate reliable segmentation of valve pressure *events* from the surrounding sensor stream, show reliable classification of *valve open* and *valve close* events, show the successful identification of *individual fixtures* with 97.9% aggregate accuracy, and show that an appropriately located and calibrated system can estimate water usage with error rates comparable to empirical studies of traditional utility-supplied water meters. In addition, we present initial forward-looking analyses of compound event detection, a comparison of sensing at different locations, and a first look at the temporal stability of pressure event signatures. Our evaluation both validates the feasibility of our approach and provides a basis for future analyses and improvements.

Figure 1 illustrates a typical plumbing arrangement in a two-bathroom home (discussed in greater detail in a later section). Figure 2 shows an actual annotated signal captured by our sensor. The signal is a kitchen faucet fixture being turned on, captured by our sensor at an exterior water bib. The remainder of this paper first discusses the theory behind our approach, presents our sensor implementation, and summarizes our in-home data collections. We then present our analyses of individual fixture identification and




Figure 2: An actual event generated when a kitchen faucet is turned on, detected by our sensor at an exterior water bib. Our fixture identification is based in segmenting the event (indicated by the highlight) and then classifying it according to its shape (see also Figure 3 and Figure 6).

water flow estimation, follow by a discussion of some important directions for future work.

RELATED WORK

Prior work has demonstrated at least three approaches to the fundamental challenge of sensing human activity in the physical world: mobile and wearable sensing, distributed direct environmental sensing, and infrastructure-mediated environmental sensing. Promising mobile and wearable activity sensing methods include accelerometer-based activity recognition [2, 10] and the detection of interaction with tagged objects via a wearable RFID reader [18]. There are many compelling applications of mobile and wearable methods, but they share a common need for a person to be willing to wear or carry the necessary device.

Environmental sensing systems take a complementary perspective, instrumenting an environment to detect activity within it. In the home, distributed direct sensing can be

based on computer vision [4], microphones within the living environment [5], many simple sensors throughout the home (e.g., reed switches on cabinet doors, accelerometer-based object manipulation sensors, infrared motion and break-beam sensors) [11, 12, 21], or a smaller number of targeted direct sensors (e.g., strain sensors under floorboards at strategic locations) [19]. Direct sensing provides valuable insight into home activities, but comes with practical costs. Installation and maintenance can be cost-prohibitive, direct sensing can create privacy concerns and a feeling of stigmatization (especially with cameras or microphones), and a littering of sensors throughout a home can be problematic with children and pets [3, 8].

Recent work has therefore examined whole-home activity sensing using just a handful of inexpensive sensors at strategic locations in a home's existing infrastructure (e.g., a home's water [7], electrical [15, 16, 17, 20], and HVAC [14] infrastructures). The central idea in infrastructure-mediated sensing is to recognize human activities, such as running a dishwasher, turning on a reading light, or walking through a doorway, according to their manifestations in these existing home infrastructures.

Infrastructure-mediated sensing helps to address many practical obstacles to everyday deployment of home activity sensing, but is inherently limited by what information can be practically and reliably extracted from a home's infrastructure. The limitations of existing approaches are especially salient for water. Fogarty *et al.* used microphones pressed against the exterior of a single home's major water pipes (cold water inlet, hot water inlet, waste water exit) to demonstrate recognition based on patterns of water use (e.g., the series of fill cycles associated with a dishwasher) [7]. However, Fogarty *et al.* found they could not reliably differentiate among multiple instances of similar fixtures (e.g., multiple sinks or toilets within a home), could not reliably identify concurrent activities (e.g., a toilet flush while a person is showering), and did not attempt to estimate the volume of water being used. They also reported difficulties with ambient noise and audio-based sensors (e.g., an air conditioning unit in close proximity to a sensor placed on a home's hot water heater). By addressing these shortcomings, HydroSense significantly advances the state-of-the-art for water-mediated home activity sensing.

With regards to water flow estimation, we are not aware of any prior work using a single sensor to estimate water flow rate to individual fixtures throughout a home. Industrial applications (e.g., irrigation systems, pharmaceutical manufacturing) have motivated sensors for high-granularity flow rate monitoring, but these existing approaches are either prohibitively expensive for residential use (2000 USD to 8000 USD for a single ultrasonic or laser Doppler velocimetry sensor) or require the professional installation of inline flow sensors (a plumber cutting into existing pipes to install an inline flow sensor for each fixture of interest). Evans *et al.* show in a laboratory environment that accelerometers mounted on the exterior of water pipes have

a strong deterministic relationship to water flow rate [6], but this is highly sensitive to pipe diameter, material, and configuration. Kim *et al.* propose using a home's existing aggregate water flow meter together with a network of accelerometers on pipes to infer flow rates throughout a home [9]. All of the above approaches require placement of multiple sensors along water pipe pathways that are uniquely associated with each fixture of interest (i.e., they are distributed direct sensing methods). Furthermore, both Evans *et al.* and Kim *et al.* require exposed piping and neither has been validated in actual home environments.

BACKGROUND AND THEORY OF OPERATION

In this section, we provide background on residential water supply systems and in-home plumbing. We also introduce the basic theory of operation, motivating our approach.

Households obtain water from one of two sources: public water supply or a private well. Public water is distributed by local utilities, relying on gravity and pumping stations to push water through major distribution pipes. Residences are connected to a water main by a smaller service line, where the water meter is typically found. Homes with private wells use a pump to draw the water out of the ground and into a small tank within the home, where it is stored under pressure. Private wells are typically unmetered.

Figure 1 depicts a typical in-home plumbing system. Cold water enters through the service line, typically at 50-100 pounds per square inch (psi)* depending on such factors as the elevation and proximity to a reservoir or pumping station. Many homes have a pressure regulator that protects the home from transients (or pressure spikes) from the main and also reduces the incoming water pressure to a level safe for household fixtures.

After the regulator, there are two basic layouts found in typical residential piping, *series plumbed* and *branched*. Almost all multi-fixture homes have a combination. The cold water supply branches to the individual water fixtures (e.g., toilets/sinks/showers) and into the water heater. A traditional water heater heats water in an insulated tank using electric coils or gas. When hot water is used, the pressure from the cold water supply line pushes hot water out of the tank and refills it with cold water. Every hot water tank has a pressure relief valve and a drain valve. Many homes also have a thermal expansion tank connected to the water heater, providing space to store excess water as it expands during heating. Some homes instead use tankless heaters, which provide hot water on demand by circulating it through burners or electric coils. Both approaches connect cold and hot regions of a home's water system, and the pressure fluctuations leveraged in our approach are propagated through both types of water heaters.

Identifying Water Fixtures

The plumbing system forms a closed loop pressure system, with water held at a stable pressure throughout the piping

* 1 psi \approx 69 mbar.




Figure 3: An actual kitchen faucet valve close event, as detected by our sensor at an exterior water bib.

when no water is flowing. Homes with a pressure regulator have stable pressure unless the supply pressure drops below the regulator's set point. Homes without a regulator may experience occasional minor changes in water pressure depending on neighborhood water demand.

The instant a valve is opened or closed (be it a bathroom faucet or a mechanical valve in a dishwasher), a pressure change occurs and a pressure wave is generated in the plumbing system (Figure 2 and Figure 3). Transient pressure wave phenomenon results from the rapid change of water velocity in a pipeline (similar to electrical transients over power lines). This is often referred to as a *surge* or *water hammer* and can create a loud hammering noise as the shockwave travels through pipes. The magnitude of the surge is independent of and much greater than the operating pressure. The transient can have a positive or negative rate of change depending on whether a valve is being opened or closed. Appliances such as dishwashers or clothes washers control their valves mechanically and thus often create the most pronounced water hammer. An abrupt change in flow can create dangerously high transients that exceed safe operating limits for residential pipes. A thermal expansion tank and water hammer arresters offer partial dampening of these transients. Most valves manifest as a water hammer impulse that is harmless but can be detected by a pressure sensor installed on the plumbing system. Water hammer typically lasts several seconds, as the pressure wave oscillates back and forth through the pipes. We can detect this water hammer effect anywhere along the plumbing infrastructure (even with dampeners installed), thus enabling single-point sensing.

The unique transient or water hammer signature that we sense for a particular fixture depends on the *valve type* and its *location* in the home pipe network. This latter point provides great discriminative power, allowing us to distinguish between two fixtures of the exact same model (e.g., two of the same toilets in the house) and even between two valves in the same fixture (e.g., the hot and cold water valves in a sink fixture) because their pressure wave impulses traverse different paths through the pipe infrastructure before reaching our sensor. Note that the magnitude of the pressure drop and resulting shockwave is dependent on the relative location of our sensor to the source of the event and the speed that the valve is opened or closed, but the shape of the signature does not change.

Estimating Flow

Changes in pressure and the rate of transient onset allow us to accurately detect valve open and valve close events as well as to estimate flow. This is analogous to an electrical circuit, where knowing the resistance (i.e., pipe restrictions) and the change in voltage (i.e., pressure) allows one to determine the current (i.e., flow).

Flow rate is related to pressure change via Poiseuille's Law, which states that the volumetric flow rate of fluid in a pipe Q is dependent on the radius of the pipe r , the length of the pipe L , the viscosity of the fluid μ and the pressure drop ΔP :

$$Q = \frac{\Delta P \pi r^4}{8 \mu L}$$

This can be simplified by the fluid resistance formulation, which states that the resistance of flow is proportional to the drop in pressure divided by the volumetric flow rate.

$$R_f = \frac{\Delta P}{Q} \equiv \frac{8 \mu L}{\pi r^4}$$

Thus, we can use fluid resistance to abstract some of the variable complexity from Poiseuille's law, resulting in:

$$Q = \frac{\Delta P}{R_f}$$

HydroSense measures the change in pressure ΔP . In order to compute Q , we must estimate the remaining unknown R_f . R_f is bounded by two factors: (1) water viscosity, which can easily be calculated according to temperature and (2) the radius of residential pipes, which are either 1/4" or 3/8" in diameter. This leaves L , the length of the pipe, as the main unknown. L will change depending on the water fixture being used, as each path from intake to fixture is different.

These equations are not comprehensive. They do not account for the smoothness of the inner pipe surface, the number of bends, valves, or constrictions in pipes, nor pipe orientation (e.g., the forces of gravity and changes in barometric pressure). However, we have found these effects can be treated as negligible for home pipe networks. We simply estimate R_f for each home by sampling flow rate at strategic locations (varying distances from the supply inlet).

PROTOTYPE SENSOR DESIGN

Our prototype HydroSense sensor implementation consists of a customized stainless steel pressure sensor, an analog-to-digital converter (ADC) and microcontroller, and a Bluetooth wireless radio (see Figure 4). We built two different HydroSense prototypes: one with a pressure range of 0-50 psi and the other 0-100 psi. The higher dynamic range is useful for homes with high supply pressure or without a pressure regulator. The pressure sensor is a P1600 series manufactured by Pace Scientific™. It comes standard with a built-in 1/4" NPT male connector, which we fitted with a 3/4" brass adaptor and Teflon tape. This allows us to easily install our sensor at any ordinary water spigot or outlet. The sensor has an operating temperature of -40 to




Figure 4: Our prototype sensor implementation. The sensor twists on to a fixture and communicates wirelessly.

257°F and a pressure response time of less than 0.5 milliseconds. The theoretical maximum sampling rate is therefore 2 kHz, but we found 1 kHz more than sufficient.

The pressure sensor's output is ratiometric to the 5 VDC supply voltage (the output voltage is a ratio of the supply). The sensor is connected to a 16-bit Texas Instruments ADS8344 ADC and AVR microcontroller, with a resolution of approximately 0.001 psi for the 50 psi sensor and 0.002 psi for the 100 psi sensor. The microcontroller is connected to a Class 1 Bluetooth radio implementing the serial port profile. It can reliably sample and stream pressure data over the Bluetooth channel. We use a 5V low-drop power regulator and the entire unit operates on a single 9V battery.

The pressure sensor has a mechanical shock rating of over 100g, making it insensitive to pipe vibration occasionally caused by some water hammer events. Although the pressure sensor comes calibrated and tested for linearity from the factory, we confirmed the output of our entire sensor system using known pressure loads. Ten samples were taken with our sensor connected to a pressure-regulated water compressor. All measurements were well within pressure sensor's tolerance of 0.25% at 25° C. The entire unit is weatherproof and can be installed in damp locations. Our current implementation does not offer a pass-through solution (i.e., allowing the installation fixture to be used as normal), but this modification is trivial.

IN-HOME DATA COLLECTION

In order to validate our general approach, our sensor implementation, and our algorithms, we collected labeled data in ten homes, in four cities, of varying, style, age, and diversity of plumbing systems (see Figure 5).

For each home, we first measured the baseline static water pressure and then installed the appropriate HydroSense unit (0-50 or 0-100 psi) on an available water hose bib, utility sink faucet, or water heater drain valve. Each collection session was conducted by a pair of researchers: one would record the sensed pressure signatures to a laptop while the other activated the home's water fixtures. The pressure signatures were recorded using a graphical logging tool, which also provided real-time feedback of the pressure data

ID / Water Supply	Style / Built / Remodel	Size / Floors / Fixtures	Exp. Tank/ Regulator / Recirc. Pump	Water Heater / Plumbing/ Static PSI	Sensor Install Point
H1 Public Utility	Single-Family 2002	3200 sqft 2 flr + bas 12 fixture	Yes Yes No	Tank PVC 46 PSI	Hose Bib
H2 Public Utility	Multi-Family 1909/96	2160 sqft 2 flr + bas 5 fixtures	No No No	Tankless Copper 46 PSI	Hose Bib
H3 Public Utility	Single-Family 2003	4000 sqft 2 flr + bas 6 fixtures	Yes Yes No	Tank Copper 41 PSI	Hose Bib
H4 Public Utility	Single-Family 1921	1630 sqft 1 flr + bas 4 fixtures	No No No	Tank Galvan. 43 PSI	Hose Bib
H5 Public Utility	Single-Family 1913	2000 sqft 2 flr + bas 5 fixtures	No No No	Tank Copper 55 PSI	Hose Bib
H6 Public Utility	Single-Family 1974/85	3100 sqft 2 flr 8 fixtures	Yes Yes Yes	Tank Galvan. 46 PSI	Hose Bib
H7 Public Utility	Aptmnt 1927	746 sqft 1 flr 5 fixtures	No Yes No	Tank Cop+Gal 33 PSI	Water Heater
H8 Public Utility	Single-Family 1922 / 2006	3650 sqft 2 flr + bas 3 fixtures	Yes Yes Yes	Tank Copper 75 PSI	Utility Sink Faucet
H9 Public Utility	Single-Family 1904 / 95 est.	1790 sqft 2 flr + bas 4 fixtures	No No No	Tank Copper 72 PSI	Hose Bib + Water Heater
H10 Private Well	Resort Cabin 1950/80	900 sqft 1 flr 4 fixtures	No No No	Tank Galvan. 65 PSI	Hose Big

Figure 5: A summary of the homes in which we collected data, including the style, size (1 sqft ≈ .093 sqm), age of the home, how many fixtures we tested, characteristics of the plumbing system, and where we installed our sensor.

via a scrolling time-series line graph. We conducted five trials per valve on each fixture (e.g., five trials for hot water and five trials for cold water). For each trial, a valve was opened completely for at least five seconds and then closed. For the toilet trials, the toilet flush and full fill cycle were logged. Note that for the faucet experiments, we did not collect data on partially opened valves nor the speed with which they were opened. We return to this issue in the discussion section.

For four of the ten houses (H1, H4, H5, and H7), we also collected flow rate information for the faucet (kitchen and bathroom) and shower fixtures. In addition to logging sensed pressure, we measured the amount of time it took to fill a calibrated bucket to one gallon (a method preferred by water utilities for accurately measuring flow). This was repeated for five trials for each valve.

This in-home data collection yielded a total of 706 fixture trials and 155 flow rate trials across 84 fixtures.




Figure 6: Several actual sensor streams from our in-home data collections. Each stream corresponds to a water valve being opened, remaining open for some amount of time, and then closed. We separately segment the valve open and valve close events from the sensor stream, as indicated by the highlighted regions of the streams. We estimate water flow to the valve based on the stabilized pressure drop while the valve remains open (the difference in pressure before the valve open versus after it).

ANALYSIS OF FIXTURE EVENT IDENTIFICATION

Given our collected data, we now pursue a three-step approach to examine the feasibility of identifying individual fixture events according to the unique transient pressure waves that propagate to our sensor. Recall that each *valve event* corresponds to water hammer when a valve is either opened or closed. We first *segment* each individual valve event from the stream, identifying its beginning and end to enable further analysis. We then *classify* each valve event as either a *valve open* or a *valve close* event. Finally, we *classify* the valve event according to the *individual fixture* that generated it. This section explicitly considers only events that occur in isolation, deferring our discussion and analysis of compound events until a later section.

Valve Event Segmentation

Before analyzing the characteristics of a valve event, we first segment it (i.e., isolate it) from the surrounding sensor stream. Segmentation must be effective for many different types of events, and so it is important to consider only features that are likely to be most typical of all valve events. Our approach is illustrated in Figure 6 (and in Figure 2 and Figure 3). The raw signal is smoothed using a low-pass linear phase finite impulse response filter. The smoothed signal and its derivative are then analyzed in a sliding window of 1000 samples (one second of sensed pressure).

The beginning of a valve event corresponds to one of two conditions. The most common is when the derivative of the smoothed signal exceeds a specified threshold relative to static pressure, indicating a rapid change (approximately 2 psi/sec for a home with 45 psi static pressure, scaled by the home’s actual static pressure). The less common second condition is when the difference between the maximum and minimum values in the sliding window exceeds a threshold relative to static pressure, indicating a slow but substantial change (approximately 1 psi for a home with 45 psi static

pressure, scaled by the actual static pressure). After the beginning of a valve event is detected via either method, the next change in the sign of the derivative represents the extreme of this valve event relative to the preceding static pressure (which may be a maximum or a minimum).

The end of a segmented valve event is typically detected as the first point at which an extreme of a fluctuation (a change in the sign of the derivative, dP/dt) is less than 5% of the magnitude of the first extreme following the beginning of the event. It is also possible for an event to be ended by a rapid increase in the magnitude of fluctuation. This corresponds to the occurrence of a compound event, as we will discuss in greater detail in a later section.

Applying this method to our collected in-home data yielded appropriate segmentations of 100% of our valve events from their surrounding sensor stream.

Classifying Valve Open and Valve Close Events

After segmenting each valve event, we classify it as either a *valve open* or a *valve close* event. We apply a hierarchical classifier that first considers the difference in the smoothed pressure at the beginning and the end of the segmented event. If the magnitude of this difference exceeds a threshold (approximately 2 psi for a home with 45 psi static pressure, scaled by the actual static pressure), the event can be immediately classified (a pressure decrease corresponds to a *valve open* and a pressure increase to a *valve close*). Otherwise, the event is classified according to the average value of the derivative between its beginning and its first extreme. A *valve open* creates an initial pressure decrease (a positive average derivative), while *valve close* events create an initial pressure increase (a negative average derivative).

Applying this method to the segmented valve events from our collected in-home data yields 100% correct classification of *valve open* and *valve close* events.

Fixture Classification

We associate *valve open* and *valve close* events with individual fixtures in a home using a template-based hierarchical classifier. When classifying an unknown event, we first filter potential templates according to four complementary distance metrics.

The first distance metric we use is a *matched filter*. Very common in signal detection theory, the matched filter is the optimal detection mechanism in the presence of additive white noise. Its primary limitation is that the signals we want to differentiate are not orthogonal. Making them orthogonal would require specific knowledge of the source of each event, exactly the information we want to infer.

Our second distance metric is a *matched derivative filter*. We include this because the derivatives of our events always resemble exponentially decreasing sinusoids. It is therefore reasonable to believe the derivatives are more orthogonal than the original pressure signals, and that this filter might provide value distinct from the above filter.

The third distance metric is based on the *real Cepstrum*, which is the inverse Fourier transform of the natural log of the magnitude of an event’s Fourier transform. This approach attempts to approximate the original version of a signal that has been run through an unknown filter (the valve event we are trying to classify has been transformed by propagation through an unknown path in a home’s water pipes). It can be shown that the lower coefficients of the Cepstrum result largely from the transfer function (an event’s propagation through a home’s pipes) and the higher coefficients largely from the source (the original valve pressure event) [13]. We are interested primarily in the transfer function (in part because it allows differentiating among multiple instances of identical fixtures in a home), and so we truncate our Cepstrum to the lower coefficients. The resulting space is highly orthogonalized, yielding a third effective and complementary matched filter.

Finally, our fourth distance metric is simple *mean squared error*, computed by truncating the longer of two events.

Similarity thresholds used to filter potential templates based on these distance metrics are learned from training data (filtering templates whose similarity to the unknown event are less than the minimum within-class similarity in the training data). If no template passes all four filters, the unknown event is not classified (an application might ignore the event, prompt a person to label an unrecognized fixture, or consider the possibility that the new event indicates the presence of a leak). If templates corresponding to multiple fixtures pass all filters, we choose among them using a nearest-neighbor classifier defined by the best performing distance metric, the matched derivative filter.

Fixture Classification Evaluation

We evaluate fixture classification using an experimental design selected to demonstrate robustness of learned model parameters across the multiple homes in our collected data.

Home	Fixture Open Identification	Fixture Close Identification
H1 (12 valves)	100%	100%
H2 (8 valves)	96.4%	100%
H3 (6 valves)	100%	100%
H4 (5 valves)	96.2%	100%
H5 (9 valves)	100%	100%
H6 (8 valves)	100%	90.0%
H7 (8 valves)	100%	100%
H8 (6 valves)	100%	97.1%
H9 (7 valves)	97.1%	97.1%
H10 (7 valves)	97.1%	77.1%
Aggregate	98.9%	96.8%
		97.9%

Figure 7: In a cross-validation test of the robustness of learned model parameters across multiple homes, our template-based classification enables identification of the individual fixtures associated with valve open and valve close events with aggregate 97.9% accuracy.

Fixture Type (number of fixtures)	Fixture Open Identification	Fixture Close Identification
Sinks (27 in 10 homes)	98.1%	95.1%
Toilets (14 in 10 homes)	98.7%	97.5%
Showers (8 in 8 homes)	95.5%	89.4%
Bathtubs (3 in 3 homes)	100%	100%
Clothes Washer (2 in 2 homes)	100%	100%
Dishwasher (1 in 1 home)	100%	N/A

Figure 8: A different view of the results, showing accuracy of identification of individual fixtures by fixture type.

Specifically, we conduct a cross-validation experiment that folds our data according to the home in which it was collected. There are ten trials in the cross-validation, with each trial using data from one home as the test data and data from the other nine homes as the training data. After learning model parameters from the test data (the four similarity filter thresholds), we classify each event in the test home using a leave-one-out method. Each test home event is classified using the other events as templates together with the model parameters learned in training.

Figure 7 presents the results of this evaluation. The figure shows the accuracy of fixture-level identification of *valve open* and *valve close* events within each home (and thus each test fold of the cross-validation), as well as the aggregate 97.9% accuracy of fixture-level classification. The relatively poor performance in identifying fixture close events in H10 (77.1% vs. > 90% for other homes) was due to noise from the eleven cabins that share the same supply line at the resort. Our sensor was picking up water events from a portion of these cabins during data collection, because the cabin was not separately metered. More work is needed to disambiguate signals in a single-meter multi-unit domain (e.g., a duplex, small apartment building), but these results indicate a single sensor may be sufficient to sense across more than one housing unit on a shared supply line.

Figure 8 presents a different view on the same data, showing the accuracy of fixture-level classification for different types of fixtures across homes. Our overall fixture-level classification across all homes is above 90%, including a number of cases where classification accuracy is 100%. All of these results are equal to or better than prior results by Fogarty *et al.* using microphone-based sensors [7]. Of particular note is our ability to reliably distinguish among different sinks within a home, as Fogarty *et al.* found that their microphone-based sensors did not capture enough information to reliably make this distinction. Our dataset contains only a few instances of clothes washer or dishwasher use, in part due to time constraints during data collection and in part because Fogarty *et al.* found these fixtures can be easily recognized by their structured cycles of water usage (an approach that can be combined with ours). However, we note that our approach is independent of the number of fill cycles (important if a dishwasher is sometimes run with a pre-rinse cycle) and allows recognition as soon as these appliances first use water (in contrast to being able to recognize them only after their pattern of fill cycles becomes apparent).

ANALYSIS OF FLOW ESTIMATION

As previously discussed, the volumetric flow rate Q is proportional to the change in pressure ΔP divided by a resistance variable R_f ($Q = \Delta P / R_f$). We calculate the change in pressure ΔP by measuring the difference between the pressure at the onset of a detected *valve open* event to the stabilized pressure at the end of the segmented *valve open* pressure wave impulse. The resistance variable R_f cannot be directly measured, so we instead learn it empirically by capturing ground truth flow rate information together with the corresponding change in pressure for each valve. This section considers two scenarios with regard to learning R_f . In the first, we assume a single calibration of flow for every valve of interest. In the second, we attempt to use information from the calibration of some valves to estimate R_f at valves that have not been calibrated.

Individually Calibrated Valves

It is not unreasonable to imagine that the process of installing a system like HydroSense might include a single calibration of each fixture in a home. In such a scenario, each valve in the home would be labeled with a known R_f value which could be combined with the sensed pressure change ΔP to estimate water flow at those valves.

We examined the accuracy of the flow estimation that might be obtained in this scenario using a cross-validation experiment to analyze the five calibrated bucket trials collected for each of the faucet and shower fixtures in H1, H4, H5, and H7 (as previously discussed in our in-home data collection section). Each trial in the cross-validation used a single calibrated bucket trial to infer a resistance variable R_f for the valve. The inferred value of R_f was then used to estimate flow in the other four trials according to the measured change in pressure ΔP . We then noted the difference between these estimated flow rates (based on the

Home	Avg Error (GPM)	Stdev Error (GPM)	Avg Error (%)	Stdev Error (%)
H1 (7 valves)	0.17	0.13	7.3	6.7
H4 (6 valves)	0.19	0.17	5.6	5.3
H5 (8 valves)	0.13	0.11	4.5	5.5
H7 (8 valves)	0.67	1.47	22.2	46.0

Figure 9: In homes H1, H4, and H5, our system is able to estimate flow at individual fixtures throughout the home with error rates comparable to that found in empirical studies of traditional utility-supplied water meters. In H7, placing the sensor on a hot water heater appears to result in a confounding of supply water main pressure with gravitational pressure due to the water in the tank.

inferred R_f) and their corresponding actual flow rates (obtained through the calibrated bucket trials). The results of this experiment are shown in Figure 9.

Three of four houses tested (H1, H4, H5) have error rates below 8% (or approximately 0.16 GPM**), comparable to 10% error rates found in empirical studies of traditional utility-supplied water meters [1]. The fourth house (H7), however, had an error rate above 20%. We believe this is due to the installation location of the sensor. Whereas the first three homes had HydroSense installed on an exterior water bib, H7's installation used the hot water heater drain valve. This results in two confounding pressure sources (the supply water main and the gravitational pressure of the water in the tank). As previously discussed, our simple pressure models currently assume a straight pipe. It is likely this situation requires a different model of R_f , and it seems that cold water valves in H7 were particularly affected. Indeed, removing H7's four cold water valves from our analysis results in a dramatically improved average error of 0.15 GPM (SD=0.18), or 4.5% (SD=3.8%). Because our dataset includes only one home with both hot water heater installation and flow rate information, future work is needed to investigate the feasibility of measuring cold water flow using a sensor installed at the hot water heater drain.

Generalizing to Uncalibrated Valves

In a scenario where only some of the valves in a home have been calibrated, it is reasonable to attempt to build a model of fluid resistance for the entire home from that subset of valves. The key idea here is that, although the pathway to each valve in the home is unique, those paths also share a fair amount of spatial overlap in the length and overall layout of the piping. For example, the toilet and sink in a particular bathroom share the same branch.

To examine this approach, we separated our calibrated bucket trials data into two datasets: a model and a test. The model was initially populated by a single randomly selected trial which was then used to infer a baseline R_f value. We applied this R_f value to calculate a flow estimate for each trial in the test dataset, comparing each to the corresponding actual flow. We next added a second random trial to the model (and removed it from the test dataset), then used the

** 1 GPM \approx .06 liters per second .




Figure 10: Because valves share a fair amount of spatial overlap in the length and overall layout of their piping, it is possible to learn to generalize calibrations across valves.

model to create a linear regression ($Q = R_f * \Delta P + b$). This regression was used to calculate flow estimates for the remaining trials in the test set. This process was repeated until all trials had been sampled. To avoid a particularly fortunate or unfortunate random sampling, we repeated this process five times for each home and averaged the results. Figure 10 presents the results (note we exclude the cold water valves from H7, consistent with our prior analysis).

After sampling five trials, the average error dropped 74% to 0.27 GPM across the four homes and within 0.11 GPM of the more comprehensive R_f data from the previous analysis. This initial result indicates considerable potential for learning to generalize calibrations across valves in a home.

DISCUSSION

The initial results presented in this paper show significant promise for single-point sensing of whole-home water activity via continuous monitoring of water pressure. We have presented a reliable method for segmenting valve pressure events from their surrounding sensor stream and for determining whether a segmented event corresponds to a valve being opened or closed. Using data collected in ten homes, we have shown 97.9% aggregate accuracy for identifying the individual fixture associated with a valve event. Analyzing flow data collected in four of those homes, we have shown that an appropriately located and calibrated system can estimate water usage with error rates comparable to empirical studies of traditional utility-supplied water meters. Our ability to identify activity at individual fixtures using a single sensor is itself an important advance, and we are not aware of prior work even attempting single-sensor estimation of the amount of water being used at fixtures throughout a home.

Although our analysis focused on identifying fixture events occurring in isolation, it is clearly important to consider the case where multiple events overlap (see Figure 11). In Fogarty *et al.*'s prior work with microphone-based sensing, they note an inability to even detect this situation [7]. As an initial investigation, we collected six compound events in H1 (two each of shower/sink, toilet/sink, and shower/toilet/sink overlaps). Our event segmentation algorithm correctly segments these overlapping events




Figure 11: Our existing event segmentation correctly identifies and segments overlapping events (shown here for overlapping shower, toilet, and faucet valve open events), and classification of such compound events is an important direction for future work. Prior work was unable to even identify the occurrence of compound events.

(ending the ongoing event when it detects a rapid increase in the magnitude of fluctuation corresponding to the beginning of another event). Preliminary experiments suggest the magnitude and shape of the events is indeed altered by the overlap. Some aspects of the frequency domain signature remain constant (i.e., high energy harmonics) but we do not have enough of these events to convincingly evaluate the effectiveness of a classification procedure. Furthermore, events that occur at the exact same instant cannot be distinguished as separate events with our current segmentation algorithm. In any case, classification of compound events is a highly important direction for building upon our current results. Bathroom activity, for example, is dominated by the usage of multiple fixtures.

An additional limitation relates to how closely our controlled experiments represent naturalistic usage of fixtures (e.g., how often do people partially open valves vs. open them full stop, how is the signal affected by the speed of valve opening or closure). These concerns primarily involve fixtures with manually controlled valves (e.g., bathtubs/showers/faucets vs. dishwashers/toilets/laundry machines). For partially opened valve events, the primary effect is a change in magnitude of the transient, not the underlying shape. We have begun longitudinal deployments in multiple households collecting labeled, naturalistic water usage data to investigate these issues further.

We have found that reliable estimation of flow is sensitive to calibration, and we have noted that our segmentation and identification algorithms include threshold parameters that worked well in the homes we studied but are not necessarily ideal. We are interested in developing techniques for automatically calibrating our methods over the course of extended usage. For example, flow estimation could potentially be automatically calibrated through occasional knowledge of whole-home aggregate water usage. Continuing deployments of wireless utility meters make this an increasingly viable approach. We also have initial evidence that system behavior is stable over time, based on a second dataset collected in H1 five weeks after our original collection. We applied our fixture classification methods to this dataset using templates from the opposite

dataset (classifying unknown events using templates collected 5 weeks apart), finding no degradation in fixture identification performance. This preliminary analysis suggests system behavior might be stable enough to apply a variety of machine learning methods for auto-calibration.

Our in-home data collection included installations at several different types of fixtures (hose bibs, utility sink faucets, water heater drain valves) with generally good results. We conducted two identical collections in H9, one using a hose bib and one using the hot water heater drain valve, under the expectation performance would be nearly identical. Figure 7 reports performance for the hose bib, but we were surprised to find performance fell to 88.6% for open events (from 97.1%) and 77.1% for close events (also from 97.1%) when we moved the sensor to the water heater drain valve (only individual fixture classification was affected, not segmentation or the determination of whether events are open or close events). On one hand, this highlights an opportunity to further explore the role of sensor placement. On the other hand, there are many other examples in Figure 7 where our current approach differed in its ability to identify the fixture associated with valve open and close events. Although these obviously come in pairs, our current approach classifies them individually. We believe there is a significant opportunity to examine methods for jointly classifying pairs of valve open and valve close events. Similarly, we currently estimate flow independent of fixture identification, but the two are clearly related and an improved method could consider them simultaneously.

CONCLUSION

We have presented a new approach to single-point infrastructure-mediated sensing of whole-home water activity. Our initial results both validate the effectiveness of our approach and provide a basis for future analyses and improvements. The infrastructure-mediated sensing strategy shows significant promise as a practical, low-cost, and unobtrusive approach to the broad deployment of sensing-based ubiquitous computing applications.

REFERENCES

- Arregui, F.J., Palau, C.V., Gascon, L. and Peris, O. (2003). *Evaluation of Domestic Water Meter Accuracy: A Case Study*. E. Cabrera and E. Cabrera, eds. 343-352.
- Bao, L. and Intille, S.S. (2004). Activity Recognition from User-Annotated Acceleration Data. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2004), 1-17.
- Beckmann, C., Consolvo, S. and LaMarca, A. (2004). Some Assembly Required: Supporting End-User Sensor Installation in Domestic Ubiquitous Computing Environments. *Proceedings of the International Conference on Ubiquitous Computing* (UbiComp 2004), 107-124.
- Brumitt, B., Meyers, B., Krumm, J., Kern, A. and Shafer, S. (2000). EasyLiving: Technologies for Intelligent Environments. *Proceedings of the International Symposium on Handheld and Ubiquitous Computing* (HUC 2000), 12-29.
- Chen, J., Kam, A.H., Zhang, J., Liu, N. and Shue, L. (2005). Bathroom Activity Monitoring Based on Sound. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2005), 47-61.
- Evans, R., Blotter, J. and Stephens, A. (2004). Flow Rate Measurements Using Flow-Induced Pipe Vibration. *Journal of Fluids Engineering* **126**(2), 280-285.
- Fogarty, J., Au, C. and Hudson, S.E. (2006). Sensing from the Basement: A Feasibility Study of Unobtrusive and Low-Cost Home Activity Recognition. *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST 2006), 91-100.
- Hirsch, T., Forlizzi, J., Hyder, E., Goetz, J., Kurtz, C. and Stroback, J. (2000). The ELDER Project: Social and Emotional Factors in the Design of Eldercare Technologies. *Proceedings of the ACM Conference on Universal Usability* (CUU 2000), 72-29.
- Kim, Y., Schmid, T., Charbiwala, Z.M., Friedman, J. and Srivastava, M.B. (2008). NAWMS: Non-Intrusive Autonomous Water Monitoring System. *Proceedings of the ACM Conference on Embedded Network Sensor Systems* (SenSys 2008), 309-322.
- Lester, J., Choudhury, T., Kern, N., Borriello, G. and Hannaford, B. (2005). A Hybrid Discriminative / Generative Approach for Modeling Human Activities. *International Joint Conference on Artificial Intelligence* (IJCAI 2005), 766-772.
- Munguia Tapia, E., Intille, S.S. and Larson, K. (2004). Activity Recognition in the Home Using Simple and Ubiquitous Sensors. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2004), 158-175.
- Munguia Tapia, E., Intille, S.S., Lopez, L. and Larson, K. (2006). The Design of a Portable Kit of Wireless Sensors for Naturalistic Data Collection. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2006), 117-134.
- Oppenheim, A. and Schaffer, R. (2004). From Frequency to Quefrency: A History of the Cepstrum. *IEEE Signal Processing Magazine* **21**(5), 95-106.
- Patel, S.N., Reynolds, M.S. and Abowd, G.D. (2008). Detecting Human Movement by Differential Air Pressure Sensing in HVAC System Ductwork: An Exploration in Infrastructure Mediated Sensing. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2008), 1-18.
- Patel, S.N., Robertson, T., Kientz, J.A., Reynolds, M.S. and Abowd, G.D. (2007). At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. *Proceedings of the International Conference on Ubiquitous Computing* (UbiComp 2007), 271-288.
- Patel, S.N., Stuntebeck, E.P. and Robertson, T. (2009). PL-Tags: Detecting Batteryless Tags through the Power Lines in a Building. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2009), 256-273.
- Patel, S.N., Truong, K.N. and Abowd, G.D. (2006). PowerLine Positioning: A Practical Sub-Room-Level Indoor Location System for Domestic Use. *Proceedings of the International Conference on Ubiquitous Computing* (UbiComp 2006), 441-458.
- Philipose, M., Fishkin, K.P., Perkowitz, M., Patterson, D.J., Fox, D., Kautz, H. and Hahnel, D. (2004). Inferring Activities from Interactions with Objects. *IEEE Pervasive Computing*, **3**(4), 50-57.
- Rowan, J. and Mynatt, E.D. (2005). Digital Family Portrait Field Trial: Support for Aging in Place. *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI 2005), 512-530.
- Stuntebeck, E.P., Patel, S.N., Robertson, T., Reynolds, M.S. and Abowd, G.D. (2008). Wideband Powerline Positioning for Indoor Localization. *Proceedings of the International Conference on Ubiquitous Computing* (UbiComp 2008), 94-103.
- Wilson, D. and Atkeson, C.G. (2005). Simultaneous Tracking & Activity Recognition (STAR) Using Many Anonymous, Binary Sensors. *Proceedings of the International Conference on Pervasive Computing* (Pervasive 2005), 62-79.