

Non-intrusive load monitoring for water (WaterNILM)

C. Schantz, B. Sennett, J. Donnal, M. Gillman & S. Leeb
Massachusetts Institute of Technology, USA

Abstract

Better water consumption decisions benefit from detailed use information. Easily installed non-intrusive vibration sensors provide a “no-fuss” retrofit solution for detecting the operation of water consuming appliances. The sensors measure pipe vibration, which are revealed to be a rich source of information for identifying loads. Vibration is processed to extract power spectral density based features which are classified with a clustering algorithm trained during install. The results can be used to track load operating schedule from the vibration data collected from as little as one pipe in a home. Mechanics governing the observed signals, and signal processing to extract operating information, are discussed in this paper. Field data from three different homes demonstrates the accuracy of this approach.

Keywords: pipe vibration, smart water meter, consumption tracking, water utility monitoring, load classification, accelerometers, pattern matching, supervised learning, machine learning.

1 Introduction

The value of information feedback to consumers has been demonstrated by Darby [1] and others. Detailed resource use tracking can find waste (Stern [2]). Utility bills often contain very coarse consumption information. A resource constrained future will demand effective and actionable information at the lowest acquisition cost. A low-sensor count system, the nonintrusive load monitor for water (WaterNilm) can provide water consumption information. Pipe vibration signatures indicate the operation of water-consuming load(s). Data sets from three field sites demonstrate WaterNilm’s potential in domestic and commercial environments. Pipe vibration sensors can be low cost and require no specialist



labor for installation, simplifying the requirements for retrofit of existing home pipe networks. WaterNilm does not use pipe vibration to infer flow rate information. Where flow rate information is desired, a new approach to piggybacking magnetic sensors on the home's water meter, described in Schantz *et al.* [3], has demonstrated high resolution flow rate sensing. Flow rate information is not required for WaterNilm's functionality, but provides a complementary capability.

1.1 Similar work

Tracking of water consumption for conservation efforts has been a subject of considerable research. Early studies employed direct sensing methods, such as temperature sensors and inline flow meters (Weihl and Kempton [4]). A method to determine water use by fixture category using flow measurements from an inline flow meter called "flow trace analysis," introduced in Dziegielewski *et al.* [5], has become popular for large studies. The technique uses an assumption of consistent operation of the target fixtures. The method has challenges dealing with multiple simultaneous flows and variability in some loads. Flow trace analysis also cannot distinguish which of many identical fixtures is operating.

Investigators have explored the use of one or more microphones to identify loads acoustically [6]. These methods are susceptible to background acoustic noise. Another single sensor approach to both load disaggregation and flow rate estimation is HydroSense (Froehlich *et al.* [7]). This approach uses a pressure sensor, usually installed at a garden hose tap, to record and classify transient pressure signals during valve turn on and turn off events. The method relies on consistent valve operation and the unique characteristics of the propagation path between the valve and sensor to impart identifying information into the pressure transient signal. The method works best for automatically activated valves. Manual valves or valves that open to different flow rates create challenges for the pressure transient technique. Events occurring during flow to another fixture are adversely affected, creating challenges in identifying overlapping loads. In addition, only the opening and closing events provide classifiable information, precluding the correct handling of simultaneous events.

WaterNilm's use of quasi-steady-state pipe vibration signatures provides a continual source of fixture identity during flow, and can correctly classify manually operated and variable flow valves and valves in combination. Because WaterNilm does not rely on transient events, simultaneous events do not present a problem. The clamp-on sensors do not require a wet connection to the water supply or the sacrifice of a dedicated garden hose tap in a home. This removes sanitation and leakage concerns of pressure sensor based methods.

2 Pipe vibration sources and transmission

The geometry of the pipe network affects the signals available for pipe vibration monitoring. Exploiting the vibro-acoustic behavior that occurs in pipe networks can aid in classification algorithm design, feature selection, sensor placement,



and set realistic expectations for identification accuracy. Intra-building water supply networks often consist of segments of metal pipes connected by bends and tee junctions. The pipes are suspended by intermittent pipe supports of many types. The geometric configuration of the pipe network including factors such as internal pipe wall burs, solder beads, or pipe deposits influence fluid flow and pipe vibration response. Sources of fluid generated vibration include the random fluctuating pressure field produced by turbulent and separated flow, and the formation and collapse of vapor bubbles during cavitation. There are also external sources of vibration from operating appliances and building occupants.

2.1 Fluid generated vibration

Water flow in household pipe networks is controlled by flow control valves in the home's water fixtures or appliances. The majority of flow rates serving domestic water loads are high enough to achieve turbulence within the pipe network. Single family homes often have pipe diameters small enough to begin turbulent transition by 0.5 GPM flow rates, while even low flow versions of common fixtures like shower heads draw 1.5 GPM at design pressure. Turbulent flow in a pipe creates a broadband fluctuating pressure field against the inner wall of the pipe (Norton and Karczub [8]). The spectral characteristics of turbulent pressure fluctuations have been modeled by a number of researchers using semi-empirical methods (see Hwang *et al.* [9] for a review). The models show spectral energy peaks in the 100 Hz to 1 kHz range, when applied to domestic pipe sizes. The spectrum remains strong out to the 250 Hz to 2.5 kHz range, before decreasing with the first power of frequency.

Regions of separated flow can also generate significant vibration. If more than one branch is flowing in normal pipe network, there will be at least one tee junction with splitting flow. In general, an incompressible fluid will experience a sharp velocity change within the junction during splitting flow accompanied by high turbulence intensity and flow separation (Vasava [10]). The magnitude of the velocity change is dependent on the proportion of flow in each exit branch and port diameter, with equal flow in each branch causing the most turbulence. The vibration at tee junctions with splitting flow is energetic and can significantly influence vibration signatures during multi load operation.

Cavitation may also occur in home plumbing. Cavitation noise occurs from the impulse like formation and collapse of vapor bubbles which form in low pressure regions of the flow, creating a broadband vibration source. Cavitation typically occurs in the contraction portions of faucets and other valves where the local pressure falls below the vapor pressure of the fluid. Cavitation generated noise is an unreliable source of indentifying information because its intensity and peak frequency changes dramatically over the possible settings of a valve. The spectral peak of cavitation noise is dependent on bubble collapse time scale. The literature reports experimental spectral peaks above 1.5 kHz to 10 kHz (Brennen [11]). Because minimum pressures are reduced as a valve is opened, bubbles are smaller and collapse contains higher frequency content. Cavitation noise in homes is likely to lie on the high end of the range seen in the literature, and can therefore be removed by low pass filtering.



2.2 Transmission and propagation

The vibratory forces acting on a pipe segment are influenced by both locally generated vibration sources and traveling waves propagating throughout the pipe network. These waves can travel in the fluid and pipe wall material. Pressure wave propagation in circular metal pipes is a key mechanism of information transfer. Metal pipes in home plumbing are well approximated as a waveguide. Only acoustic modes that satisfy boundary conditions at the centerline and the walls will propagate. The majority of the acoustic energy generated at flow discontinuities occurs via the higher order modes in the region of the source (Norton and Karczub [8]). For the turbulence dominated frequency range and pipe diameters found in home pipe networks, only the lowest order longitudinal axisymmetric mode i.e. plane wave mode, may propagate. Because higher order modes dissipate strongly with distance, they do not serve as a useful source of identifying information.

Pipe walls are not perfectly rigid, however, and many forms of coupling allow internal pressure waves to cause wall vibration. This means vibration sensors should be placed outside of the dissipation region of higher order modes near major flow disturbances. The slowest dissipating higher order mode will dissipate approximately 100 dB every 3 pipe diameters assuming rigid pipe. This mode has the lowest cut-on frequency i.e. frequency above which the mode is non-dissipative, at ~ 33 kHz in a 1 inch pipe. Cut-on frequency is inversely proportional to pipe diameter. Loss mechanisms like viscosity, acoustic radiation, and turbulence in the flow will also affect the propagation of all pressure waves. The losses depend on the details but are estimated to be 15 dB of attenuation every 1000 pipe diameters, according to White and Sawley [12] for industrial steel piping.

Wiggert and Tijsseling [13] identify three forms of coupling between fluid and pipe structure: poison coupling, friction coupling, and junction coupling. Poison coupling has minor effect outside of high pressure fluid transients like water hammer capable of stretching the pipe walls. Friction coupling concerns the transfer of momentum between the fluid and structure through flow friction. For plane waves the primary vibration coupling mechanism is junction (or geometric) coupling. An example is a plane wave incident on a 90 degree elbow in the pipe. As long as this elbow is free to translate, coupling will occur, allowing information contained in the internal acoustic pressure field to be transformed into externally measurable pipe vibrations. Selecting a sensor location on a pipe span bordered by a strong source of junction coupling like an elbow fitting is recommended, provided t-junctions are avoided due to strong vibration during splitting flow.

The copper pipe sizes typically usually used in homes have wall thicknesses greater than 10% of the pipe radii, and their cross section is not deformed in the low frequency region. This permits use of a quasi one-dimensional beam model to describe their vibration (Jong [14]). Jong also states that the distribution of energy between the vibration modes and between the fluid and structure depend strongly on the boundary conditions of the pipe. This dependence provides the



mathematical basis for the reason why piping system layout and support geometry affects the vibro-acoustic energy flow in the pipe network, and is responsible for the identifying information in the pipe vibration signature.

3 Feature generation and clustering

The supervised learning approach requires that vibration signatures be learned in a commissioning step during system installation. During normal operation pipe vibration time series is recorded and processed to extract identifying features for classification to identify the responsible flow configuration.

3.1 Signal processing chain

We chose a low cost MEMS accelerometer (part number ADXL 203) for its bandwidth, low noise, high sensitivity, and analog output. The sensor is rated to 2.5 kHz bandwidth, with a resonance in the silicon mass at 5.5 kHz, which is followed by a drop in sensitivity. The accelerometer was sampled at rates between 12 kHz and 16 kHz to take advantage of the sensor's sensitivity roll-off and avoid the need for anti-aliasing filters. Digital low pass filtering was employed to attenuate the cavitation dominated frequencies above 2 kHz.

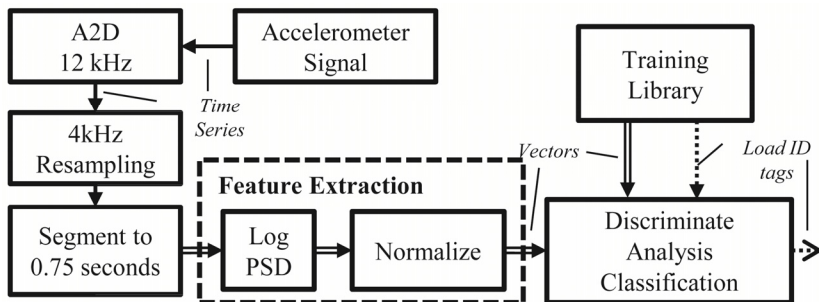


Figure 1: Block diagram of algorithm for vibration signature classification.

Pipe vibration signals can be modeled as a stochastic input modified primarily by the modal characteristics of the actively flowing branches of the pipe network. The intensity of the vibration input will vary with flow rate, but the modal properties of the pipe segments will remain relatively fixed. This makes frequency domain processing highly attractive, and power spectral density (PSD) estimates were chosen as classification features. A PSD is calculated for each segment of the vibration time series. Segment duration is a tunable parameter and 0.75 seconds was chosen as a default, as it was observed to perform well in cross validation. Log PSDs values are estimated with the Welch method and then the vector of values is normalized to unit length. The frequency bin width of the PSD estimate is another tunable parameter, and heuristically should be small enough to distinguish between the identifying modes in the pipe network without

resulting in excessive estimate variance. The vibration signal processing chain is shown schematically in figure 1.

3.2 Subclustering for variable flow valves

The performance of the features as described can be poor against large flow rate changes. The classification algorithms rely on clustering of feature vectors in a feature space and the ability to distinguish clear boundaries between clusters. Assigning the feature vectors associated with the full flow range of a variable flow valve to a single cluster during training leads to diffuse clusters and difficulty in determining effective classification boundaries. For this reason, the training step may employ sub clustering to divide the often diffuse clusters of these valves into multiple sub clusters. This improves compactness and generally results in better classification boundaries. The k-means algorithm is used for this purpose. The number of sub-clusters chosen for each variable valve is a tuneable parameter. Results exploring this choice are presented in section 4.3.

3.3 Segment classification

Many classification algorithms exist (Bishop [15]). Discriminant analysis was chosen, employing MATLAB's statistics toolbox functions, [16]. Both linear and quadratic discriminate analyses were tested. To reduce the training time required to generate the covariance matrices, the naive Bayes assumption of independent features is used to generate a diagonal covariance matrix. The classifier tags each vibration segment with the most similar flow configuration from the training set.

4 Field tests and identification accuracy

Field tests were performed at three single family residences to gather data for cross validation. The sites are denoted H1, H2, and H3. Each site had access to the point of water pipe ingress and metering, located close to the floor, followed by a vertical pipe segment rising to the basement ceiling. Pipe vibration was recorded at two locations and all corresponding flow configurations were documented. Accelerometer one (A1) was always installed downstream of the water meter on the vertical pipe, prior to the first t-junction. Accelerometer two (A2) was installed on the vertical output pipe of the home's hot water heater. The sensors were mounted with a pipe clamp and plastic mounting block. Photographs of installed sensors are shown in figure 2(a)–(c). Horizontal axes, i.e. normal to the pipe direction, of the accelerometers were sampled. At H3, flow rate was monitored by a prototype magnetic sensor attached to the home's water meter, described in Schantz *et al.* [3]. A laptop with USB data acquisition box recorded all signals.

4.1 Site 1 results and training length requirements

The objective of the H1 test was to gather proof of principle data. A gas furnace was operating intermittently the vicinity of A2, but data was gathered while the






Figure 2: (a) A1 installed on vertical pipe above water meter. (b) A2 installed on hot water heater output. (c) Accelerometer mounting detail. Sensors used modified ADXL 203 evaluation board and plastic carrier block secured with pipe clamp.

disturbance was not active. Ten or more seconds of steady state flow (discounting transients from valve actuation) were recorded for the nineteen flow configurations described in table 1. One fixture tested was kitchen faucet with side sprayer accessory. The fixture was tested at the hard stop settings bounding the left-right motion of the control lever, with and without the sprayer diverter active, accounting for four “single valve” flow configurations. Another complex fixture was the home’s bathtub tap and shower head diverter. The H1 data set includes a range of flow rates and mixture ratio settings for the bathtub tap. The shower head was always recorded at the fully hot valve setting.

H1 data allows exploration of the tradeoff between training data length and identification accuracy. A practical method must not require excessive time

Table 1: H1 Data load description.

Type:	Count	Description
Binary Valve	3	Washing machine, Toilet 1 st fl., Toilet 2 nd fl.
Throttling Valve	6	2 nd fl. H. and C. sinks, 1 st fl. H. and C. sink, Base. H. and C. sink
Kitchen Faucet	4	H. tap, C. tap, H. sprayer, C. sprayer.
Mixed hot and cold	2	Bathtub tap, shower head.
2 concurrent loads	4	Shower+(1 st fl. Toilet, 2 nd fl. Toilet, 2 nd fl. C. sink, C. kitchen tap)
3 concurrent loads	2	Shower + Toilets, Shower + 2 nd fl. C. sink + 2 nd fl. Toilet.

commitment during commissioning. The fourteen bolded flow configurations in table 1 were recorded for 20 seconds or greater and were selected to investigate training length duration. PSD estimate segment length was set to 0.75 seconds. The average misclassification rates (MCR) vs. training data duration of linear discriminate analysis (LDA) and quadratic discriminate analysis (QDA) is plotted in Fig 4. The MCR is generated from the average of 100 trials using randomly selected data of the given duration from the total available signal associated each flow configuration. The results show that training data length does not need to be excessive, approximately 10 seconds provides good accuracy.

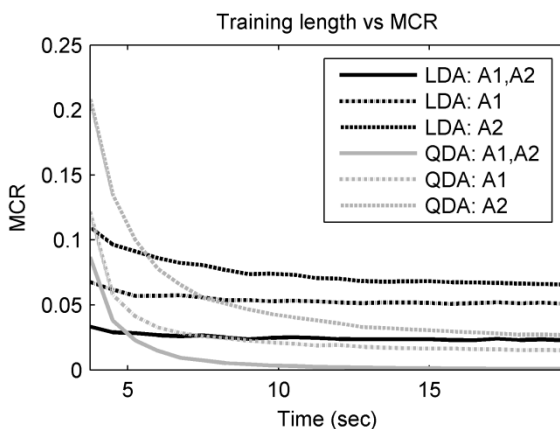


Figure 3: MCR vs. duration of training data. Durations with 10 segments (7.5 seconds) or fewer used “leave one out” cross validation. Durations of more than 10 segments used 10 fold cross validation.

4.2 Site 2 results and combination signatures

H2 is a two story single family home. The objective of H2 testing was to acquire a data set containing many simultaneous load flow configurations, and to study the effect of external vibration disturbances on pipe vibration. Thirty one flow configurations were recorded for 10 seconds or more, described in table 2. Additionally, four of these flow configurations were recorded again during the operation of a furnace nearby A1. Pipe vibration was also recorded without internal flow during furnace operation. Comparison of PSDs show the furnace disturbance had negligible effect on vibration at A1 compared to internal pipe flow.

H2’s data set allows an exploration of the possibility of creating artificial vibration signatures for simultaneous valve flow from signatures acquired during individual valve training. If a method could be found to achieve this, it would reduce the effort required for system commissioning. We determined, however, that linear superposition of PSDs will not succeed. This outcome is consistent

Table 2: H2 Data load description.

Type:	Count	Description
Binary valve	2	Base. Toilet, 1 st fl. Toilet.
Throttling valve	6	1 st fl. H. & C. sink, Base. H. & C. sink, Front hose tap, Rear hose tap
Kitchen faucet	4	H. tap, C. tap, H. sprayer, C. sprayer.
Mixed hot and cold	2	Bathtub tap, shower head.
Load concurrent with shower head	6	Base. Toilet, 1 st fl. C. sink, 1 st fl. H. sink, 1 st fl. Toilet, C. kitchen tap, H. kitchen tap
Load concurrent with 1 st fl. Toilet	5	1 st fl. C. sink, 1 st fl. H. sink, C. kitchen tap, H. kitchen tap, Bathtub tap
Loads concurrent with base. C. Sink	4	Base. Toilet, Bathtub tap, 1 st fl. Toilet, 1 st fl H. sink.
3 concurrent loads	2	Shower + 1 st fl. Toilet + (1 st fl. C. sink, or 1 st fl. H. sink)

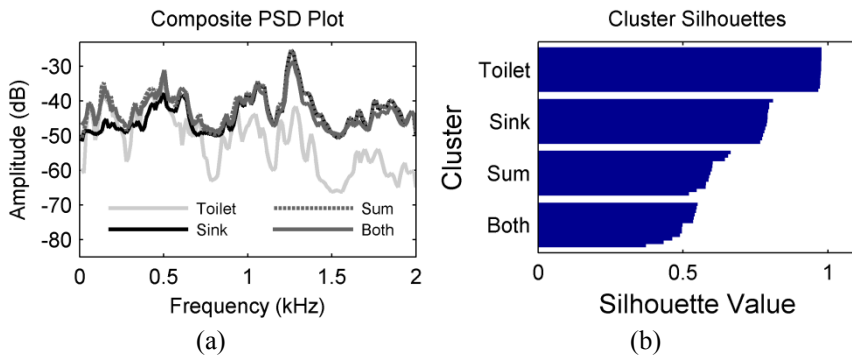


Figure 4: (a) PSD of separate and combined flow configurations. (b) Silhouette plot. Positive values indicate a point is closer to its own cluster.

with the nonlinear nature of fluid mechanics, especially at t junctions. Figure 4(a) shows the PSD of vibration at A1 for two single valve flow configurations; cold water flow at a bathroom sink and flow to refill the toilet reservoir in the same bathroom. Also shown is a linear combination of the sink PSD and toilet PSD, labeled 'Sum'. The actual PSD recorded with both valves flowing simultaneously is labeled 'Both'. The procedure to record the sample of faucet flow and dual valve flow first activated the faucet valve for a period of approximately 10 seconds before pressing the flush lever of the toilet. The setting of the faucet valve remained constant throughout the test to remove operator variability, and the toilet was selected as the partner load for repeatability.

While the PSDs labeled 'Sum' and 'Both' seem similar, the silhouette plot in figure 4(b), which displays a measure of cluster cohesiveness, shows that the

‘Sum’ and ‘Both’ signatures are distinct. Based on these results we determine that concurrent load operation must be trained explicitly. Despite the large increase in the number of clusters for the classifier to consider, classification accuracy remains above 90%. Table 3 shows the general MCR for all 31 flow configurations with 10 seconds of recorded vibration or greater, the majority of which involve combination load flow.

Table 3: Classification performance on H2 data.

Misclassification rate	A1 features	A2 features	A1 and A2 features
LDA	0.0945	0.0889	0.0705
QDA	0.0632	0.0508	0.0384

4.3 Site 3 challenge case and subclustering

H3 represents a challenge case. With three full bathrooms (one containing two faucets), one half bathroom and two kitchen sinks, H3 contains more water fixtures than H1 or H2. Flow rate was also recorded with the non-intrusive magnetic water meter monitor. To investigate the effect of flow rate on classification accuracy, testing procedure was modified to include more flow rate settings of variable flow valves. Variable flow valves were tested with at least four flow rates spanning the full range and always included a test at maximum opening. The mixed temperature tests of the single lever controlled faucets were also measured at four flow settings each. The flow configurations recorded are summarized in table 4.

Table 4: H3 Data load description.

Type	Count	Description
Binary valve	6	Four Toilets and two bidets.
Throttling valve	10	H. and C. sinks in basement, half bathroom, and two full bathrooms. Front hose tap, garage hose tap.
Two kitchen faucets	4	H. tap, C. tap in each faucet.
Two master faucets	4	H. tap, C. tap in each faucet.
Two kitchen faucets	2	Middle setting for H. and C. flow in each faucet.
Two master faucets	2	Middle setting for H. and C. flow in each faucet.
Mixed hot and cold.	6	Three Shower heads and three bathtub taps.
Loads concurrent with front hose tap	6	Half bath toilet, Full bath toilet, C. kitchen tap, H. kitchen tap, C. kitchen tap, H. kitchen tap
Loads concurrent with master shower	2	Master toilet, master bidet, hot and cold flow in two handle faucet.
Loads concurrent with bathroom shower	4	Bathroom faucet operated cold only, hot only, or mixed temperature flow. Bathroom toilet.

The MCRs for the H3 data set, shown in table 5 are higher than H1 and H2. This is largely the result of the extensive variation of flow rates recorded at the home's variable flow valves. The subclustering technique described in section 3.2 is applied to the H3 data set with good results. For cross validation we modify the definition of a correct match to include segments classified to any sub cluster of the correct parent cluster. The improvement of two and three group subclustering is evident. Subclustering may create a cluster with fewer than 10 seconds of associated data. Such clusters were removed completely before calculation of the MCR. One tradeoff of subclustering is increased training time of variable flow valves.

Table 5: Classification performance on H3 data.

MCR	No sub-clustering		2 group subclustering		3 group subclustering	
	LDA	QDA	LDA	QDA	LDA	QDA
A1 features	0.3612	0.3532	0.2264	0.2230	0.1510	0.1667
A2 features	0.2848	0.2464	0.1596	0.1513	0.1084	0.1211
A1 and A2 features	0.2282	0.1768	0.1130	0.0821	0.0514	0.0426

The identifying information in pipe vibration signatures are largely turbulence driven. The presence of flow rate information in the H3 data set allows for second evaluation of MCR and subclustering performance. The results of table 6 were computed identically to those in table 5 save for the removal of all segments with flow rate below 0.4 gpm, chosen for the beginning of turbulent flow in the smallest copper pipe used in the plumbing of H3. Notably, this filtering completely removed the two bidet only flow configurations.

Table 6: Classification performance on H3 data without low flow segments.

MCR	No subclustering		2 group subclustering		3 group subclustering	
	LDA	QDA	LDA	QDA	LDA	QDA
A1 features	0.3225	0.3076	0.2091	0.2120	0.1480	0.1623
A2 features	0.2424	0.2041	0.1220	0.1207	0.0739	0.0878
A1 and A2 features	0.1815	0.1310	0.0840	0.0588	0.0344	0.0350

Acknowledgements

This research was supported by the Grainger Foundation, the Kuwait-MIT research program, the ONR Structural Acoustics Program, and the MIT Energy Initiative.

References

- [1] Darby, S. The effectiveness of feedback on energy consumption. *A review for DEFRA of the literature on metering, billing and direct displays*, 2006.



- [2] Stern, P. C., Information, incentives, and proenvironmental consumer behavior. *Journal of Consumer Policy*, **22(4)**, pp. 461-478, 1999.
- [3] Schantz, C., Donnal, J., Leeb, S., Marimuthu, P. N., Habib, S., WaterWolf: Water Watch on Load Flow
- [4] Wehl, J. S., Kempton, W., Residential hot water energy analysis: instruments and algorithms, *Energy and Buildings*, **8(3)**, pp. 197-204, 1985.
- [5] Dziegielewski, B., Opitz, E. M., Kiefer, J. C., Baumann, D. D., & American Water Works Association. *Evaluating urban water conservation programs; a procedures manual*, AWWA, 1993.
- [6] J. Fogerty, J., Au, C., Hudson, S.E., Sensing from the basement: a feasibility study of unobtrusive and low-cost home activity recognition. *Proc. of the ACM Symposium on User Interface Software and Technology*, UIST, pp. 91-100, 2006.
- [7] Froehlich, J., Larson, E., Saba, E., Campbell, T., Atlas, L., Fogarty, J., Patel, S, A Longitudinal Study of Pressure Sensing to Infer Real-World Water Usage Events in the Home. *Proc. of the Ninth International Conference on Pervasive Computing*, San Francisco, California, June 12-15, 2011.
- [8] Norton, M., & Karczub, D., *Fundamentals of Noise and Vibration Analysis for Engineers Second Edition* Cambridge University Press, Cambridge, UK, 2003.
- [9] Hwang, Y. F., Bonness, W. K., Hambric, S. A., Comparison of semi-empirical models for turbulent boundary layer wall pressure spectra *Journal of Sound and Vibration* **319(1)** pp. 199-217, 2009.
- [10] Vasava, P. R., *Fluid flow in T-junction of pipes*. Diss. Lappeenranta University Of Technology, 2007.
- [11] Brennen, C. E., *Cavitation and bubble dynamics*. No. 44. Oxford University Press, 1995.
- [12] White, P. H., & Sawley, R. J. Energy transmission in piping systems and its relation to noise control. *Journal of Engineering for Industry*, **94(2)**, 746-751, 1972.
- [13] Wiggert, D. C., & Tijsseling, A. S., Fluid transients and fluid-structure interaction in flexible liquid-filled piping, *Applied Mechanics Reviews* **54(5)** pp. 455-481, 2001.
- [14] Jong, C. A. F. D., *Analysis of pulsations and vibrations in fluid-filled pipe systems*, Ph.D Thesis, Eindhoven University of Technology, Eindhoven, The Netherlands , 1994.
- [15] Bishop, C. M., *Pattern recognition and machine learning Vol. 1*. Springer, New York, 2006.
- [16] Matlab Statistics Toolbox Documentation; The MathWorks, Inc., Online. <http://www.mathworks.com/help/stats/index.html>