# Water Nonintrusive Load Monitoring

Christopher Schantz, John Donnal, Brian Sennett, Mark Gillman, Sean Muller, and Steven Leeb

***Abstract*— Resource conservation decisions require detailed consumption information. This paper presents sensors and signal processing techniques that use pipe vibration signatures to non-intrusively identify water consumption at the appliance level. The method requires as little as one easily installed vibration sensor. This method provides a no-fuss retrofit solution for detecting the operation of a building's water consuming appliances. In addition, flow rate is nonintrusively obtained from a conventional water meter via a new, high sensitivity strap-on magnetic sensor. Together, these two sensors track load operating schedule and water consumption in a building, demonstrated here at three different field test sites.**

***Index Terms*— Smart homes, water conservation, vibration measurement, fluid flow measurement, magnetoresistive devices.**

## I. Introduction

BETTER information leads to better resource consumption decisions [1], [2]. While highly granular resource tracking can find waste [3]–[5], utility bills often contain only very coarse consumption information. Many water utility companies are examining the adoption of smart water meters [6] to gather higher resolution data. However, current systems for tracking water consumption either fail to provide adequately detailed reporting, or require intrusive and expensive installation, or both. A resource constrained future will demand effective and actionable information at the lowest acquisition cost. A nonintrusive load monitor for water (WaterNILM), is a system of sensors and signal processing techniques that together act to transform a buildings metal pipe network into a sensor for detailed water consumption tracking. This paper presents circuits, signal processing techniques and field demonstrations of the WaterNILM approach.

WaterNILM can determine water fixture identity and operating schedule from vibration measurements, and perform detailed consumption tracking by fusing data from vibration and flow measurements. Pipe vibration signatures indicate the operation of water-consuming loads. Vibration sensors can be low cost and require no specialist labor for installation, making retrofit attractive. However, the vibration sensors do not quantitatively determine flow rate. The vibration sensors and identification accuracy can be enhanced by determining instantaneous flow rate, sensed by a pair of externally attached quadrature magnetic sensors on the home's water meter. The time and volume resolution our magnetic sensor technique is significantly higher than fixed volume-per-pulse counters typical of instrumented water meters. A system like WaterNILM, combined with similar electrical monitoring capabilities [7], [8], opens the door to fully correlated resource consumption tracking on the individual level. The non-intrusive nature encourages wide adoption without relying on the slow spread of smart appliances. The information produced will allow the enormous computation capabilities of the modern era to be applied to a wide range of practical problems, including revealing new resource conservation strategies.

Manuscript received October 10, 2014; accepted November 4, 2014. Date of publication November 20, 2014; date of current version February 5, 2015. This work was supported in part by the Grainger Foundation, in part by the Kuwait MIT Research Program, in part by the Kuwait Foundation for the Advancement of Science, in part by the Office of Naval Research Structural Acoustics Program, and in part by the MIT Energy Initiative. The associate editor coordinating the review of this paper and approving it for publication was Prof. Subhas C. Mukhopadhyay.

C. Schantz, J. Donnal, B. Sennett, M. Gillman, and S. Leeb are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: cschantz@mit.edu; jdonnal@mit.edu; bsennett@mit.edu; gillman@mit.edu; sbleeb@mit.edu).

S. Muller is with Merrimack High School, Merrimack, NH 03054 USA (e-mail: seanp.muller@merrimack.k12.nh.us).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSEN.2014.2372053

## II. Approaches to Load and Flow Rate Sensing

### A. Other Work

Various methods have been used to track load or fixture level water consumption. Direct sensing employs temperature sensors, inline flow meters, and electric power sensors on water appliances [9], [10]. Flow trace analysis uses flow rate signals from an inline flow meter to determine water use. A commercially popular flow trace analysis software, AuquaCraft Corporation's Trace Wizard, employs an iterative procedure requiring trained human experts for trace interpretation and statistical parameter tuning. In addition flow trace analysis does not distinguish which fixture is operating in a home, only fixture class [11], and assumes consistent operation of the target fixtures. It has significant challenges dealing with multiple simultaneous flows [12] and variability in human operated fixtures. Investigators have tried microphones to identify loads acoustically with mixed results [13]. A single sensor approach to both load disaggregation and flow rate estimation is the HydroSense project, [14], which uses a garden hose tap pressure sensor to record and classify transient pressure signals during valve turn on and turn off events. Like flow trace methods, consistent valve operation is beneficial, and like WaterNILM, the unique characteristics of the propagation path between the valve and sensor impart identifying information into the measured signal. The method has difficulty with simultaneous transients or even concurrent loads due to modification of pressure transients during flow. WaterNILM's vibration based method does not rely on transients, avoiding this problem. The clamp-on sensors also avoid a wet connection to the water supply [15]. The Nonintrusive Autonomous Water Monitoring System (NAWMS) project

of [16] uses many distributed wireless vibration sensors for flow rate tracking using the ideal relation of increasing turbulent and pipe vibration intensity with flow rate. Fixture identity is determined by a one-sensor-per-load approach, similar to direct sensing. However, only laboratory test bed results are given. The work presented here revealed that pipe vibration intensity was not monotonic and even decreased with increasing flow rate in some cases. The literature currently does not report a water consumption monitoring scheme that uses machine learning and classification techniques on vibration signatures and highly sensitive magnetic sensors to recover instantaneous flow rate information from in home water meters.

### *B. Pipe Network Vibration and Transmission*

Understanding 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. For a topical summary of the relevant characteristics see [17]. Intra-building water supply networks often consist of segments of metal pipes connected by bends and tee junctions. Turbulence is a broadband vibration source on the order of 100 Hz to 1 kHz range for flow scenarios in typical buildings [18]. Strong higher frequency vibration sources on the order of 10 kHz also occur due to cavitating bubble collapse. There are also external sources of vibration from operating appliances and building occupants. Vibration sensor bandwidth should overlap the turbulent dominated frequencies because most domestic pipe flows are turbulent. The occurrence of cavitation depends on flow valve setting and its spectral characteristics vary dramatically [19] making it an unreliable source of identifying information. We remove this content with low pass filtering.

The forces acting on a pipe are influenced by locally generated vibration and traveling pressure waves propagating throughout the pipe network. The strongest vibration sources are at flow disturbances like valves, elbows and tee-junctions [20]. Tee-junctions vibrate strongly during splitting flow, which has implications for the pipe vibration during multiple load operation. Sensor placement too close to a flow disturbance may cause the measured signal to be overwhelmed by local effects. The internal pressure waves are coupled to the motion of the pipe wall through geometric coupling, most strongly at bends in the pipe [21]. Pipe segments can be modeled as quasi-one dimensional beams [22], the vibration response of each segment being dependent on the supporting boundary conditions. Essentially, the flow path to each water load is excited by distributed vibration sources. Each path's vibratory characteristics are coupled to the pipe's internal pressure field, providing the basis for identifying loads with unique pipe vibration signatures. The pipe network behaves as is its own flow destination sensor.

### *C. Water Meters and Flow*

A feature of most building pipe networks is a positive displacement water meter to track volumetric consumption for utility billing purposes. The meter is an opportunity for high resolution flow rate sensing. The display unit of a meter can be mechanical or electronic, but the measurement assembly often has a rotating magnetic coupling for motion transfer through the water tight boundary of the meter [23]. The magnetic coupling is either tracked with follower magnets to increment a mechanical volume register or counted with a reed switch or similar pulse generating device in the meter's display. The magnetic field is detectable outside of the meter. Because the field weakens with distance from the coupling, its peak to peak amplitude is often less than the earth's magnetic field in many locations outside the meter. Fig. 1 shows the field measured in the periphery of a common nutating disk type meter.

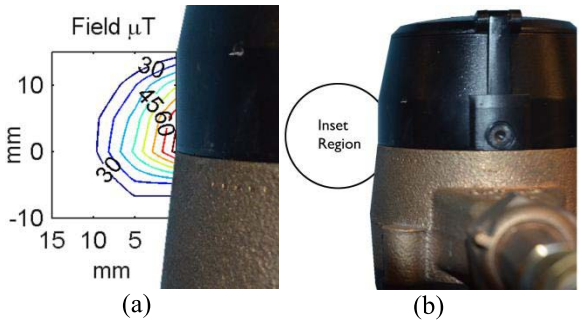


Fig. 1. External peak to peak field strength during flow.

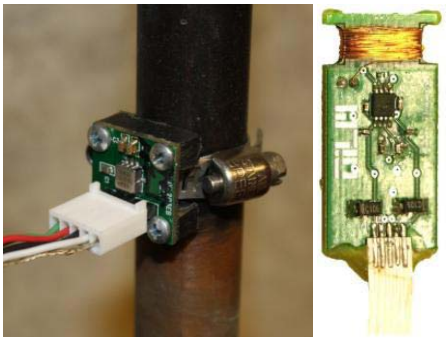


Fig. 2. (a) MEMS accelerometer attached to pipe. (b) Compensated TMR sensor board.

## III. Sensing Hardware

### *A. Pipe Vibration Sensors*

Vibration is sensed with commercially available micro-electro-mechanical systems (MEMS) accelerometers. An Analog Devices ADXL 203 was selected for its bandwidth, low noise, high sensitivity, and analog output. The ADXL 203 is rated to 2.5kHz bandwidth and has a sensing range of $\pm 1.7$g full scale. The sensor has an inherent resonance at 5.5kHz, which is followed by a sharp drop in sensitivity as frequency increased, negating the need for anti-aliasing filters given sufficiently high sample rates. The sensor was attached to the pipe with a pipe clamp and plastic mounting block, pictured in Fig. 2a.

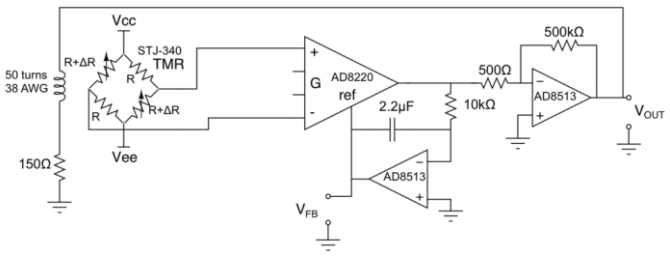


Fig. 3. Compensated TMR circuit.

### *B. Magnetic Sensor for Flow Rate*

Measuring instantaneous flow rate from the weak magnetic fields on the periphery of a water meter requires high quality waveform measurement at low cost. Ref [24] offers an overview of the options. Recent advances in materials and fabrication techniques have improved Tunneling Magnetoresistive (TMR) devices [25]. We have developed a compensated magnetic sensor circuit built around a TMR element for flow rate sensing [26] (Fig. 2b).

The circuit employs an STJ-340, a TMR Wheatstone bridge sensor produced by Mirco Magnetics [27]. Field changes imbalance the bridge, which can be measured by a differential amplifier. While the STJ-340 can detect very small fields (25mV/G as constructed), there are two significant challenges in using it as a magnetic waveform sensor. First, DC offset errors quickly saturate the sensor output. The offset errors from the environment and from imbalance in the bridge itself (which can be up to 10%) must be removed before applying any significant gain to the output. The sensor also has a nonlinear response to large changes in the applied field.

As seen in schematic in Fig. 3, the DC offset error is removed by an integrator connected to the REF pin of the instrumentation amplifier, resulting in an AC signal. The REF pin is sampled to recover DC information. The amplifier output is then fed through a gain stage that drives an air core solenoid wrapped around the STJ-340. The current through this solenoid builds a magnetic field that opposes the applied field, creating a feedback loop that zeros the operating point of the STJ-340. Exposing the sensor to small fields improves linearity and increases operating range. The current in the compensation coil is sensed as a voltage across a resistor, and amplified for output.

### *C. Vibration Signal Feature Generation*

To identify water loads we must extract consistent features from the pipe vibration signal. The vibration signal is sampled at 12kHz or 16 kHz, shown in Fig. 4a. Cavitation dominated frequencies are removed by digital re-sampling to 4 kHz, and the signal is divided into 0.75 second segments, Fig. 4b. Example segments are shown in Fig. 4c. The intensity of turbulent excitation depends on flow rate, but it acts on the relatively fixed vibration modes of the pipe structure. This makes frequency domain features attractive for identification. In Fig. 4d a power spectral density (PSD) is calculated for each segment. The log PSD reveals more detail, and is normalized

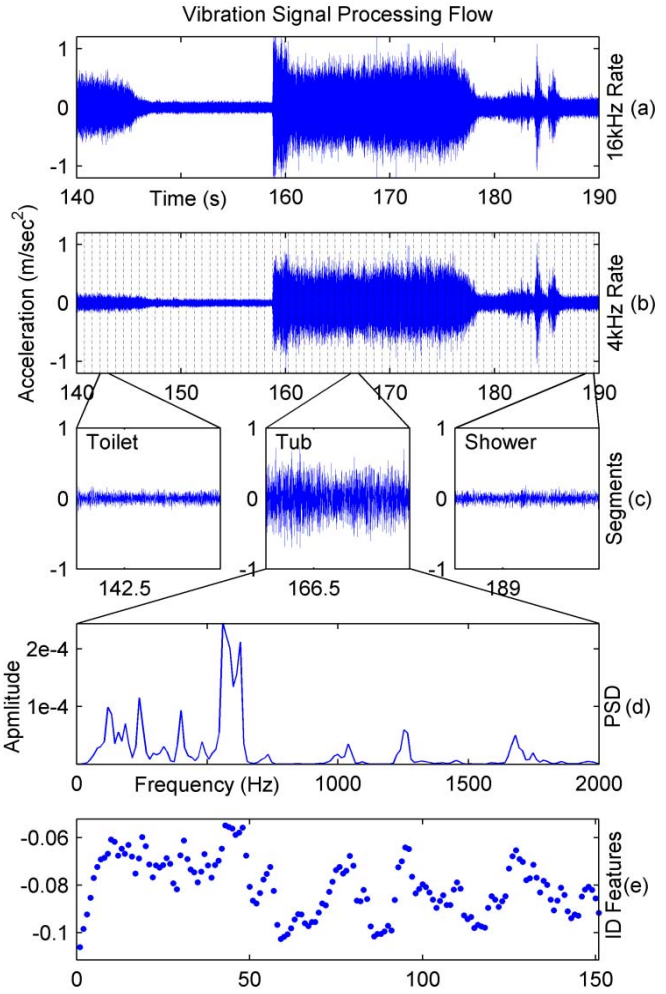


Fig. 4. Steps for feature vector generation. (a) Sampled vibration time series. (b) Low pass filtering, re-sampling, and segmentation. (c) Example segments. (d) Power spectral density. (e) Feature vector elements are the normalized frequency bins of the log PSD.

by the log PSD norm to reduce spread in the cluster of a given load's feature vectors caused by differences in flow rate and ensuing turbulent excitation level. A final feature vector is shown in Fig. 4e. The feature vectors have a large number (151) of dimensions. A parallel coordinate plot is a way to visualize high dimensional data. Fig. 5 shows feature vector differences between a few example loads.

## IV. Signal Processing for Load ID and Flow Rate

### *A. Classification and Clustering*

Feature vectors can be classified by many methods. Linear Discriminate Analysis (LDA) and Quadratic Discriminate Analysis (QDA) are cluster based classification methods. LDA and QDA with diagonal covariance matrix assumption were chosen because they handle situations where the dimension of the feature space is large compared to the number of points in the training clusters, [28]. Data to train the algorithms must be acquired during system commissioning, and minimizing the duration of training data required will save time and water.

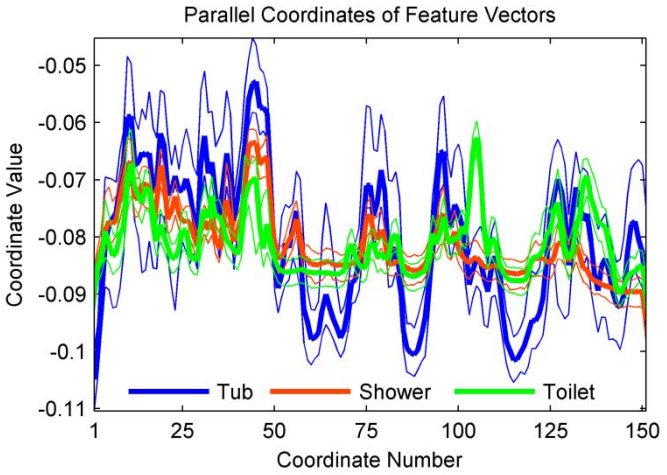


Fig. 5. Collections of load feature vectors with 5% to 95% spread shown in thinner line. Each coordinate number corresponds to a frequency bin of the PSD based feature vectors.

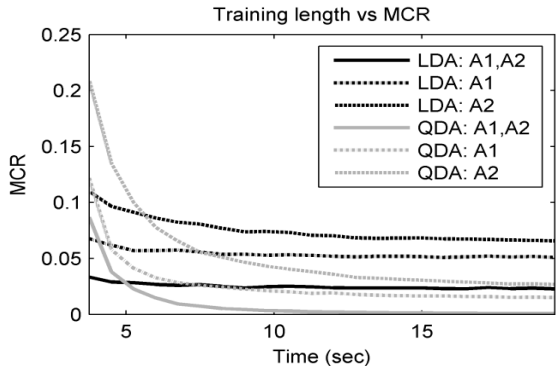


Fig. 6. Mean of 100 trials of segment MCR versus. training data duration. Computed with leave-one-out validation if less than ten segments in training data, and ten fold cross validation more than ten segments in training data.

We used data from fourteen load or load combinations from field site one to investigate the Mis-Classification Rate (MCR) vs. training data duration tradeoff. Details of collection are given in section V. Fig. 6 shows that 10 seconds of training for each load should suffice for 90% or better segment ID accuracy. Accuracy can be improved by employing simple usage rules for the loads being monitored. For example if flow rate is constant, majority voting over a window of three or more consecutive segment IDs will correct misidentified segments in many cases.

### *B. Analytic Signal Construction for Flow Rate Estimation*

Positive displacement meters have a fixed volume per rotation constant. This means the instantaneous frequency (IF) of the meter's magnetic field is proportional to the flow rate through the water meter. IF is defined for a monocomponent analytic (i.e. complex valued) signal. Real valued signals must be converted to analytic form prior to IF estimation. For intermittent flow, the signal will contain frequency content approaching DC, rendering the usual Hilbert transform method for constructing analytic signals inaccurate due to spectral leakage across DC [29]. Instead, we use a signal from a

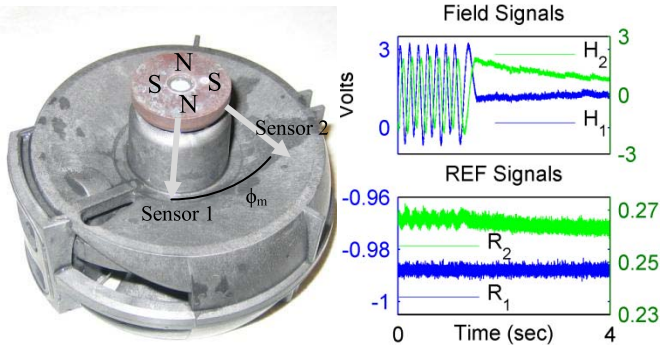


Fig. 7. Sensor position axes and example signals.

second sensor mounted at an angular difference $\phi_m$ from the first sensor. Two sensors also allow unambiguous flow direction sensing.

Each sensor gives two outputs, the AC coupled magnetic field measurement denoted $\mathrm{H_i}(t)$ and the low frequency or DC content of the reference pin voltage of the instrumentation amplifier denoted $\mathrm{R_i}(t)$. The position vectors of two sensors relative to the internal magnet disk are shown in Fig. 7, along with example signals.

Mounting constraints or an unknown magnetic pole count may preclude knowledge of the magnetic angle between sensors. A correction angle, $\phi_c$ in eqn (1), must be found to compensate. The magnetic angle is $P\phi_m$, where $P$ is the pole pair count of the magnetic drive being monitored.

$$z\,(t) = H_1\,(t) - e^{i\phi_c} H_2(t) \tag{1}$$

Rotation of $H_2(t)$ by the correction angle will minimize the negative valued frequency content of $z(t)$. This angle is a geometric parameter and need only be computed once during system installation. PSDs plots illustrate the procedure in Fig. 8a-c.

### *C. Pseudo DC Response Recovery*

The magnetic elements in the flow meter are stationary during no flow and the sensor is exposed to a constant DC field depending on the resting position of the magnets. The DC offset corrector will adjust the reference pin voltage, causing the output of the I-amp to trend to zero. This effect can be seen around t = 1.5 sec in Fig. 7. When flow resumes the integrator will again modify the reference voltage based on the average revolving field. This new mean is generally different from the stationary field preceding flow, causing unwanted startup transient distortion. These distortions can potentially cause difficulty with IF estimation. To remove these distortions the AC content in $\mathrm{H_i}\,(t)$ and the low frequency or DC content in $\mathrm{R_i}\,(t)$ must be combined.

Because stationary periods expose the sensor to a random dc field value, the offset correction can be modeled as piecewise time varying. Combining AC and DC content in the sensor's signals must also be done piecewise. The first step determines the boundaries between flow and no-flow regions. High pass filtering $z(t)$ constructed in (1), then applying a median

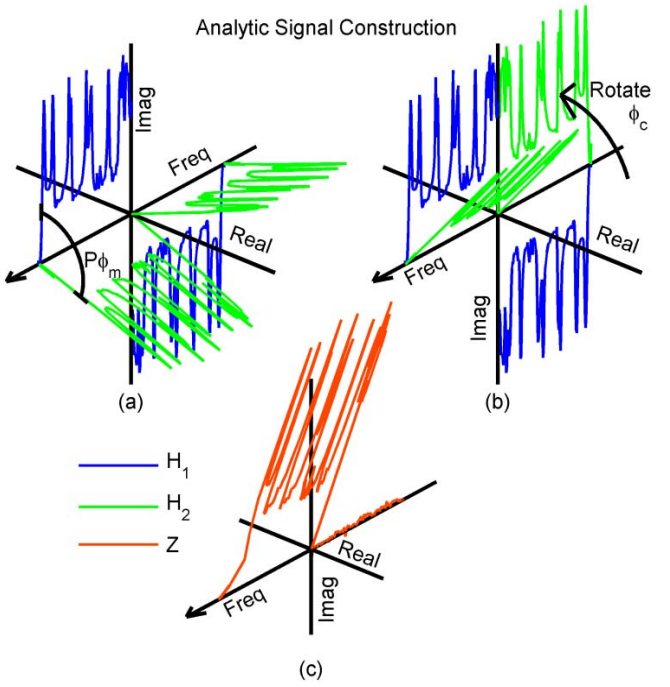


Fig. 8. (a) PSD of sensor signals showing relative phase of positive and negative frequency components. While the PSDs of $H_1$ and $H_2$ are real valued, the phase angle relationship between the positive and negative sequence components of each signal is represented graphically by rotation in the complex plane. (b) $H_2$ (PSD shown) rotated by correction angle in $\phi_c$ in eqn (1). (c) Sum of signals creates analytic z(t) (PSD Shown, rotated relative to original phase of $H_1$.

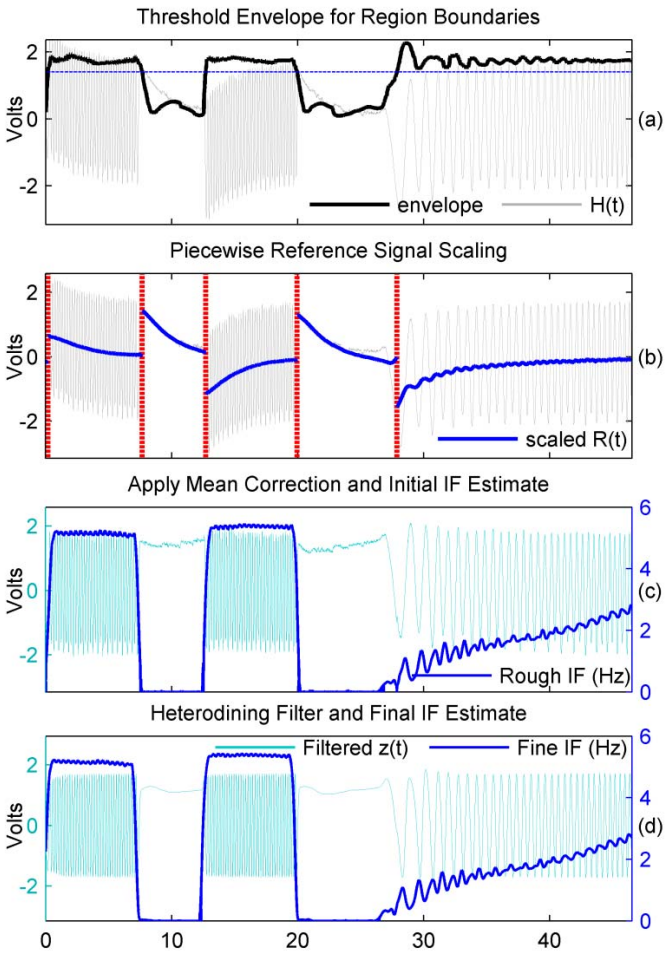


Fig. 9. Correction of field measurement and IF estimation: (a) flow versus no flow regions, (b) scaling of reference pin signal, (c) rough IF estimation from dc corrected signal, and (d) heterodyne based variable center frequency band pass filtering and final IF estimation.

filter and thresholding its magnitude envelope reveals these boundaries, visible in Fig. 9a and Fig. 9b. Doing this sets a minimum detectable flow rate. One may design the filter and DC offset corrector for a desired minimum flow rate.

After the boundaries between flow regions of $H_i(t)$ are found, the corresponding segment of a low pass filtered $R_i(t)$ is scaled and shifted via least squares to match, as in Fig. 9b, and then subtracted from the $H_i(t)$ segment to correct the distortion in the flow segments. The no flow segments are now zero mean, and are further shifted by the first value of the scaled $R_i(t)$ segment as an approximation to the true dc value of stationary field during no flow. The segments are reassembled into a corrected $H_c(t)$ for Fig. 9c. The correction is done to $H_1(t)$ and $H_2(t)$ and a new $z(t)$ is formed.

### *D. Filtering and Instantaneous Frequency Estimation*

IF is applicable only to mono-component signals, [28]. The magnetic signals from a water meter contain harmonics and noise, which are removed by variable center frequency band pass filtering. This filter is implemented via a heterodyning process. First, an initial rough IF estimate, shown in Fig. 9c, is used to create a mixing signal, which is multiplied by the analytic signal $z(t)$ to shift the frequency range of interest to dc for narrow low pass filtering. The results are then multiplied by the inverse of the mixing signal to return the filtered $z(t)$ to the original domain, allowing an improved IF estimate. Fig. 9d plots the real part of the filtered $z(t)$ and the fine IF estimate.

IF is computed from the derivative of the instantaneous phase. The inherent numerical derivative makes signal noise a primary concern. We use the Kay-Tretter algorithm described in [30] and implemented in [31]. It is based on an assumption of short time linear phase variation and operates on relative phase differences between data points to avoid wrapping problems. A discontinuity will only contribute one erroneous differential phase to the estimation window. Slope of the window of phase differences is computed in a least squares sense. The window length sets the frequency response vs. noise susceptibility tradeoff. The rough IF estimate for the mixing signal used to filter out harmonics is performed with a longer window than the final IF estimate.

## V. Field Tests and Results

Field tests were performed at three single family residences. The sites are denoted H1, H2, and H3. Each site had access to the point of water pipe ingress and metering, located in the structure's basement. Pipe vibrations during flow to different water loads in isolation or in combination were recorded at two locations. Accelerometer one (A1) was always installed

TABLE I
H1 LOAD DESCRIPTION AND MCR

| **Type**: | **Count** | **Description** | | |
|---|---|---|---|---|
| Binary Valve | 3 | Washing machine, 1st floor toilet, 2nd floor toilet | | |
| Throttling Valve | 6 | 1st floor, 2nd floor, and basement hot & cold faucets | | |
| Kitchen Faucet | 4 | Hot setting, cold setting, hot sprayer, cold sprayer | | |
| Mixed hot and cold | 2 | Bathtub tap, shower head | | |
| 2 concurrent loads | 4 | Shower and (1st floor toilet, 2nd floor toilet, 2nd floor cold faucet, cold kitchen tap) | | |
| 3 concurrent loads | 2 | Shower and 1st and 2nd toilets, Shower and 2nd floor cold faucet and 2nd floor toilet. | | |
| **Miss-Classification Rate Percentage** | | **A1 features** | **A2 features** | **A1 and A2 features** |
| LDA | | 9.1% | 8.6% | 4.6% |
| QDA | | 6.7% | 5.0% | 1.7% |

TABLE II
H2 LOAD DESCRIPTION AND MCR

| **Type**: | **Count** | **Description** | | |
|---|---|---|---|---|
| Single binary valve | 2 | Basement toilet, 1st floor toilet. | | |
| Single throttling valve | 6 | 1st floor and basement hot & cold faucets, front hose tap, rear hose tap | | |
| Kitchen Faucet | 4 | Hot setting, cold setting, hot sprayer, cold sprayer | | |
| Mixed hot and cold | 2 | Bathtub tap, shower head | | |
| Concurrent with shower head | 6 | 1st floor or basement toilet, 1st floor hot or cold faucet, hot or cold kitchen tap | | |
| Concurrent with 1st floor toilet | 5 | 1st floor hot or cold faucet, hot or cold kitchen tap, Bathtub tap | | |
| Concurrent with Basement cold faucet | 4 | 1st floor or basement toilet, bathtub tap, 1st floor hot faucet | | |
| 3 concurrent loads | 2 | Shower and 1st floor toilet and 1st floor cold or hot faucet | | |
| **Miss-Classification Rate Percentage** | | **A1 features** | **A2 features** | **A1 and A2 features** |
| LDA | | 9.5% | 8.9% | 7.1% |
| QDA | | 6.3% | 5.1% | 3.8% |

(a) (b) (c)

Fig. 10. (a) A1 installed on vertical pipe above water meter. (b) A2 installed on hot water heater output. (c) Prototype magnetic sensor mounting detail, also visible at bottom of (a).

downstream of the water meter prior to the first tee-junction. Accelerometer two (A2) was installed on the output pipe of the home's hot water heater. The installed sensors are shown in Fig. 10a and Fig. 10b. At H3, flow rate was monitored via the TMR magnetic sensor attached to the home's water meter, Fig. 10c. A laptop with USB data acquisition box recorded all signals.

The H1 testing objective was to gather proof of principle data. Ten or more seconds of steady state flow (discounting transients from valve actuation) were recorded for the nineteen flow configurations described in Table I. Notably, the single lever controlled kitchen faucet accounted for four flow configurations depending on setting and sprayer use. The bathtub tap was recorded at a range of flows, but the shower head was tested at full hot flow. Segment identification accuracy easily exceeds 90%.

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

The data from H2 allows an exploration of creating artificial vibration signatures for simultaneous load flows from signatures acquired during individual load training. If a method could be found to achieve this, it would reduce the effort required for training. Two single valve flow configurations were specially tested: cold water flow at a bathroom faucet and flow to refill the toilet reservoir in the same bathroom. The test first activated the faucet valve for approximately 10 seconds before actuating the flush lever of the toilet. The setting of the faucet handle remained constant throughout the test to remove operator variability. The toilet was also recorded in isolation, and it was selected as the partner load for repeatability.

We determined, however, that linear superposition of single load PSDs will not succeed in creating useful composite signatures. This outcome is consistent with the nonlinear nature of fluid mechanics, especially at tee-junctions. The silhouette plot in Fig. 11, displays a measure of cluster

TABLE III
H3 LOAD DESCRIPTION AND MCR

| Type | Count | Description | | | | | |
|---|---|---|---|---|---|---|---|
| Binary Valve | 6 | Four toilets and two bidets. | | | | | |
| Throttling Valve | 10 | Hot & Cold faucets in basement, and three bathrooms, front hose tap, garage hose tap. | | | | | |
| Four "one lever" faucets | 12 | Hot, mixed, and cold setting at each faucet | | | | | |
| Mixed hot and cold. | 6 | Three shower heads and three bathtub taps. | | | | | |
| Concurrent with front hose | 6 | Half bath toilet, full bath toilet, hot or cold kitchen faucet 1, hot or cold kitchen faucet 2 | | | | | |
| Concurrent with shower 1 | 2 | Toilet or bidet | | | | | |
| Concurrent with shower 2 | 4 | "one handled faucet" hot or cold or mixed setting, or toilet | | | | | |
| **Miss-Classification Rate Percentage** | **No sub-clustering** | | **2 group sub-clustering** | | **3 group sub-clustering** | |
| | **LDA** | **QDA** | **LDA** | **QDA** | **LDA** | **QDA** |
| A1 features | 36.1% (32.3%) | 35.3% (30.8%) | 22.6% (20.9%) | 22.3% (21.2%) | 15.1% (14.8%) | 16.7% (16.2%) |
| A2 features | 28.5% (24.2%) | 24.6% (20.4%) | 16.0% (12.2%) | 15.1% (12.1%) | 10.8% (7.4%) | 12.1% (8.8%) |
| A1 and A2 features | 22.8% (18.2%) | 17.7% (13.1%) | 11.3% (8.4%) | 8.2% (5.9%) | 5.1% (3.4%) | 4.3% (3.5%) |

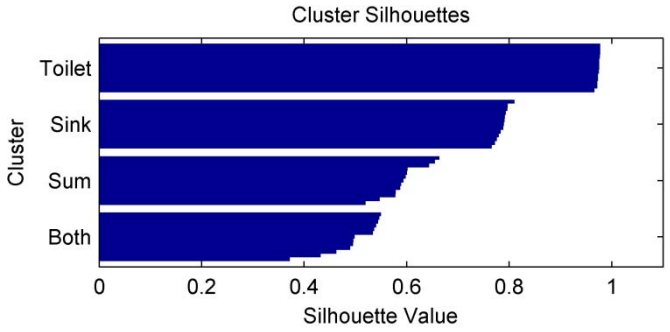


Fig. 11. Silhouette plot for composite signatures. Positive values indicate candidate point is closer to own cluster than any other, and show cluster distinctness.

cohesiveness. The silhouette value of a point is the difference between the average distance between the point and its own cluster members and the average distance between the point and the members of the nearest "other" cluster, normalized by the largest of the two average distances. Fig. 11 shows that the artificial signature cluster labeled "Sum" and the true signature cluster, recorded during operation of both loads, labeled "Both" have positive silhouette values e.g. are in distinct clusters. Concurrent load operation must be trained explicitly. Despite the increase in the number of clusters for the classifier to consider, classification accuracy remains above 90%. Table II 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. Priority can be given to training loads in combination with longer duration loads or site specific usage patterns to reduce training burden.

H3, described in Table III, represents a challenging case. Three full bathrooms (one containing two faucets), one half bathroom and two kitchen sinks gives H3 more water fixtures than H1 or H2. Flow rate was recorded at H3. One objective was to investigate the effect of varying flow rate on vibration signature consistency. Testing procedure included at least four flow rates settings spanning the full range of every variable flow valve. The mixed temperature tests of the single lever controlled faucets were also measured at four flow settings each. MCR is generally higher for the H3 data set compared to other sites due to the variation in vibration signatures caused by flow rate changes.

A technique called sub-clustering was applied to the H3 data set with good results. The LDA and QDA 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. Sub-clustering divides the diffuse clusters of variable flow valves into multiple sub-clusters using the k-means algorithm. This improves cluster compactness and generally results in better classification boundaries.

For the sub-clustered cross validation results in Table III, 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 sub-clustering is evident. Note: when applied to H3 data, sub-clustering may create a cluster with fewer than 10 seconds of associated data. These clusters were removed completely before calculation of the MCR. One tradeoff of sub-clustering is increased training time of variable flow valves.

Flow rate in the H3 data set allows a second evaluation of MCR and sub-clustering performance. The identifying information in pipe vibration signatures is largely turbulence driven. The MCR percentages enclosed in parenthesis in Table III were computed after the removal of segments with flow rate insufficient for turbulence in the pipes of H3. MCR improves, notably due to the removal of the electronically controlled bidet from the set of loads. The bidet load is hard to identify with vibration, but is easier to identify by its small flow rate and consistent duration. Fig. 12 shows the vibration, flow rate, and load identity tags from a test recorded at H3. The flow rate information has been used to identify "Transition" regions. This is necessary as the vibration-only classifier, being trained on steady state vibration signatures, assigns identities spuriously during load transitions. The "No-Flow" regions were mostly identified via vibration, but further improved by comparing to a minimum flow rate. Further improvements to ID accuracy are possible by applying site specific or load specific rules. For example, the bidets

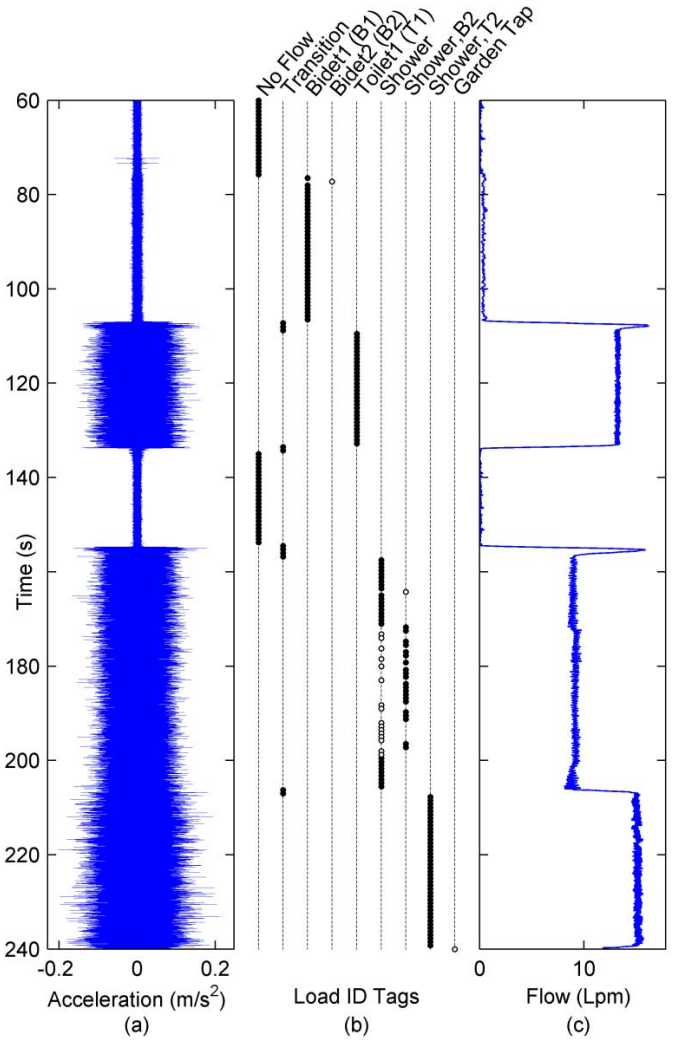


Fig. 12. Sample data record of real flows at H3. (a) pipe vibration at A1. (b) LDA classifier results using A1 and A2 features. Open circles are misclassified. (c) Flow rate.

at H3 have a consistent duration that could be exploited to correct many of the misclassified tags in Fig 12. Techniques from flow trace analysis can also be employed to glean further load information from the flow rate signal.

## VI. Conclusions

Two distinct non-intrusive sensors for tracking fluid consumption in a pipe distribution network have been presented in this paper. The first, a vibration based technique, uses MEMS vibration sensors that are inexpensive and easily installed to track water consumption by appliance using pipe vibration signatures. The second, a magnetoresistive sensor, estimates water flow rate by detecting the small magnetic fields associated with the operation of conventional mechanical flow meters. Together, these sensors provide a fully electronic and easily installed sensor system for tracking water flow. Results identifying the operation of individual water consuming appliances in homes have been demonstrated with better than 90% accuracy in field tests at three different homes. The sensors can also be used either individually or together as a basis for detecting pathological operating conditions and fluid leaks. We are working to apply these sensors for the detection of subtle leaks before they would otherwise become noticeable on conventional utility bills or aggregate consumption estimates.

## References

[1] S. Darby, "The effectiveness of feedback on energy consumption. A review for DEFRA of the literature on metering, billing, and direct displays," Environmental Change Inst., Univ. Oxford, Oxford, U.K., Tech. Rep., 2006.

[2] A. Faruqui, S. Sergici, and A. Sharif, "The impact of informational feedback on energy consumption—A survey of the experimental evidence," *Energy*, vol. 35, no. 4, pp. 1598–1608, Apr. 2010.

[3] P. Ester, *Consumer Behavior and Energy Conservation: A Policy-Oriented Experimental Field Study on the Effectiveness of Behavioral Interventions Promoting Residential Energy Conservation*. Boston, MA, USA: Martinus Nijhoff Pub., 1985.

[4] A. H. McMakin, E. L. Malone, and R. E. Lundgren, "Motivating residents to conserve energy without financial incentives," *Environ. Behavior J.*, vol. 34, no. 6, pp. 848–863, 2002.

[5] P. C. Stern, "Information, incentives, and proenvironmental consumer behavior," *J. Consum. Policy*, vol. 22, no. 4, pp. 461–478, 1999.

[6] Oracle Corp., Redwood Shores CA, USA. (Sep. 2009). *Smart Metering for Water Utilities*. [Online]. Available: http://www.oracle.com/us/industries/utilities/046596.pdf

[7] M. D. Gillman *et al.*, "Energy accountability using non-intrusive load monitoring," *IEEE Sensors J.*, vol. 14, no. 6, pp. 1923–1931, Jun. 2014.

[8] S. R. Shaw, S. B. Leeb, L. K. Norford, and R. W. Cox, "Nonintrusive load monitoring and diagnostics in power systems," *IEEE Trans. Instrum. Meas.*, vol. 57, no. 7, pp. 1445–1454, Jul. 2008.

[9] J. S. Weihl and W. Kempton, "Residential hot water energy analysis: Instruments and algorithms," *Energy Buildings*, vol. 8, no. 3, pp. 197–204, 1985.

[10] G. O. Ladd and J. L. Harrison, *Electric Water Heating for Single Family-Residences: Group Load Research and Analysis*. Palo Alto, CA, USA: Electric Power Research Institute, 1985.

[11] W. B. DeOreo, J. P. Heaney, and P. W. Mayer, "Flow trace analysis to assess water use," *J. Amer. Water Works Assoc.*, vol. 88, no. 1, pp. 79–90, 1996.

[12] A. K. Nguyen, H. Zhang, and R. A. Stewart, "Analysis of simultaneous water end use events using a hybrid combination of filtering and pattern recognition techniques," in *Proc. 6th Biannial Meeting Int. Environ. Modelling Softw. Soc.*, Leipzig, Germany, 2012, pp. 1–8.

[13] J. Fogerty, C. Au, and S. E. Hudson, "Sensing from the basement: A feasibility study of unobtrusive and low-cost home activity recognition," in *Proc. ACM Symp. User Inter. Softw. Technol. (UIST)*, 2006, pp. 91–100.

[14] E. Larson *et al.*, "Disaggregated water sensing from a single, pressure-based sensor: An extended analysis of HydroSense using staged experiments," *Pervasive Mobile Comput.*, vol. 8, no. 1, pp. 82–102, Aug. 2012.

[15] C. Schantz, B. Sennett, J. Donnal, M. Gillman, and S. B. Leeb, "Water non-intrusive load monitoring," presented at the Urban Water II, Algarve, Portugal, 2014.

[16] Y. Kim, T. Schmid, Z. M. Charbiwala, J. Friedman, and M. B. Srivastava, "NAWMS: Nonintrusive autonomous water monitoring system," in *Proc. 6th ACM Conf. Embedded Netw. Sensor Syst.*, 2008, pp. 309–322.

[17] C. Schantz, "Methods for non-intrusive sensing and system monitoring," Ph.D. dissertation, Dept. Mech. Eng., Massachusetts Inst. Technol., Cambridge, MA, USA, 2014.

[18] Y. F. Hwang, W. K. Bonness, and S. A. Hambric, "Comparison of semi-empirical models for turbulent boundary layer wall pressure spectra," *J. Sound Vibrat.*, vol. 319, no. 1, pp. 199–217, 2009.

[19] W. K. Blake, *Mechanics of Flow-Induced Sound and Vibration: Complex Flow-Structure Interactions*, vol. 1. New York, NY, USA: Academic, 1986.

[20] M. Norton and D. Karczub, *Fundamentals of Noise and Vibration Analysis for Engineers*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2003.

[21] D. C. Wiggert and A. S. Tijsseling, "Fluid transients and fluid-structure interaction in flexible liquid-filled piping," *Appl. Mech. Rev.*, vol. 54, no. 5, pp. 455–481, 2001.

[22] C. de Jong, "Analysis of pulsations and vibrations in fluid-filled pipe systems," Ph.D. dissertation, Eindhoven Univ. Technol., Eindhoven, The Netherlands, 1994, ISBN 9038600747.

[23] F. Arregui, C. Enrique, and C. Ricardo, *Integrated Water Meter Management*. London, U.K.: IWA Pub., 2006.
[24] J. Lenz and S. E. Alan, "Magnetic sensors and their applications," *IEEE Sensors J.*, vol. 6, no. 3, pp. 631–649, Jun. 2006.
[25] S. Ikeda *et al.*, "Tunnel magnetoresistance of 604% at 300 K by suppression of Ta diffusion in CoFe$B$/Mg$O$/CoFe$B$ pseudo-spin-valves annealed at high temperature," *Appl. Phys. Lett.*, vol. 93, no. 8, pp. 082508-1–082508-3, Aug. 2008.
[26] C. Schantz, J. Donnal, S. B. Leeb, P. N. Marimuthu, and S. Habib, "WaterWOLF: Water watch on load flow," presented at the Urban Water II, Algarve, Portugal, 2014.
[27] *STJ-340: Four Element Bridge Magnetic Sensor Data Sheet, Micro Magnetics*. [Online]. Available: http://www.micromagnetics.com/docs/STJ-340_datasheet.pdf, accessed May 1, 2014.
[28] T. Hastie, R. Tibshirani, and J. Friedman, "High-dimensional problems: $p \gg N$," in *The Elements of Statistical Learning*, 2nd ed. Berlin, Germany: Springer-Verlag, 2009, ch. 18, pp. 649–694.
[29] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals," *Proc. IEEE*, vol. 80, no. 4, pp. 520–538, Apr. 1992.
[30] S. Kay, "A fast and accurate single frequency estimator," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 12, pp. 1987–1990, Dec. 1989.
[31] F. Auger, P. Flandrin, P. Gonçalvès, and O. Lemoine, *Time-Frequency Toolbox*. France: CNRS/Rice Univ., 1995–1996. [Online]. Available: http://tftb.nongnu.org

**Christopher Schantz** received the B.S. degree from the California Institute of Technology, Pasadena, CA, USA, in 2008, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2011 and 2014, respectively, all in mechanical engineering. His research interest lies in signal processing for sensing and control systems.

**John Donnal** received the B.S. degree in electrical engineering from Princeton University, Princeton, NJ, USA, in 2007, and the M.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2013, where he is currently pursuing the Ph.D. degree. His research interests include nonintrusive load monitoring synthesis, energy harvesting, and communications systems.

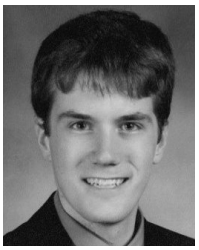

**Brian Sennett** received the B.S. degree in electrical science and engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2013, where he is currently pursuing the degree at the Laboratory for Electromagnetic and Electronics Systems. He is interested in audio signal processing.

**Mark Gillman** received the B.S. degree in electrical engineering from United States Military Academy, West Point, NY, USA, in 2004. He is an active duty Army Major and a Professional Engineer in the state of Missouri, and is currently pursuing the M.S. degree in electrical engineering at the Massachusetts Institute of Technology, Cambridge, MA, USA. His current research interests include the measurement and analysis of microgrid power systems in military applications.

**Sean Muller** is currently a Teacher with Merrimack High School, Merrimack, NH, USA. Since 2000, he has served as a Research Affiliate with the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, where he is involved in wide range of research projects in materials science and engineering. He has made refereed publications based on his collaborative work at MIT, and was recognized as a Cubist Pharmaceutical's New England Science Teacher of the Year in 2012.

**Steven Leeb** received the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1993, where he has been a Faculty Member with the Department of Electrical Engineering and Computer Science since 1993. He holds a joint appointment with the Department of Mechanical Engineering, MIT. He is concerned with the development of signal processing algorithms for energy and real-time control applications.