Water or Slime? A platform for automating water treatment systems

Neil Klingensmith
University of Wisconsin
Computer Engineering
Madison, WI, USA
naklingensmi@wisc.edu

Anantharaghavan Sridhar
University of Wisconsin
Electrical Engineering
Madison, WI, USA
ananth.sridhar@wisc.edu

Zachary LaVallee
University of Wisconsin
Mechanical Engineering
Madison, WI, USA
zlavallee@wisc.edu

Suman Banerjee
University of Wisconsin
Computer Science
Madison, WI, USA
suman@cs.wisc.edu

ABSTRACT

We introduce a control system we call AWESOME, which is a platform for improving efficiency in water treatment systems. In this work, we study water softeners, which are a primary source of Chloride ion pollution in effluent sewage water [4, 1]. When water softeners regenerate, they dump large volumes of salt and water down the drain. Existing water softeners use open-loop controllers that do not include water quality sensors to trigger regeneration. As a result, they generally regenerate too frequently, wasting salt and water and polluting the environment. AWESOME aims to reduce the frequency of water softener regenerations by sensing water quality and applying domain knowledge to trigger regenerations at optimal times. The sensor readings, which may be noisy or inaccurate, are processed by backend algorithms to accurately predict when regeneration is required. In a pilot deployment on the UW campus, AWESOME decreased salt consumption by an average of 27%, saving $5240 in the first year after installation.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

General Terms

IoT

Keywords

Water Treatment; Building Automation; Sensor Platforms; Smart Buildings

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
BuildSys'15, November 4–5, 2015, Seoul, South Korea..
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3981-0/15/11 ...$15.00.
DOI: http://dx.doi.org/10.1145/2821650.2821652.

1. INTRODUCTION

Regional water politics are increasingly putting pressure on local governments to reduce their fresh water consumption and water pollution [13, 2]. For this reason, the availability of fresh water will set limits on population expansion in urban centers as well as the productivity of arable land, both in the United States and abroad.

In Madison and other metropolitan areas, water softeners are a primary source of Sodium and Chloride ion pollution [4, 1]. Pollution of fresh water sources is a major concern because it threatens the supply of potable water on which urban populations depend. Salt waste produced by softeners is discharged into waste water treatment facilities. Once dissolved, it is difficult and expensive to remove.

To combat the increasing concentration of pollutants being discharged by water treatment systems, we present AWESOME[1]. The goal of this work is to increase the amount of water that a softener can treat between regenerations, thereby reducing the amount of salt per gallon required to treat hard water. We accomplish this by introducing an adaptive feedback control system that can monitor the volume and quality of the treated water and automatically make real-time decisions about when to regenerate the filtration medium. AWESOME is unique because it is a *bolt-on* solution that can be added to any existing water softener system to save salt. In contrast to other water-saving control systems, it relies very little on human intervention.

The main challenge to building such a control system is that most commercially-available water quality sensors are delicate instruments whose output can drift over time. One key advantage of AWESOME is that it includes adaptive control algorithms that can tolerate sensor drift while detecting and responding correctly to true changes in water quality caused by depletion of the filtration medium.

What is a water softener? Water softeners stop lime buildup in pipes and equipment by removing dissolved minerals – Calcium and Magnesium ions – from tap water. This is typically accomplished by exchanging Calcium and Magnesium ions with Sodium ions. As Calcium and Magnesium-rich water passes through a filtration medium, Sodium ions,

[1]AWESOME stands for **A** **W**at**E**r **S**oftener **O**nline Opti**M**ization **E**ngine

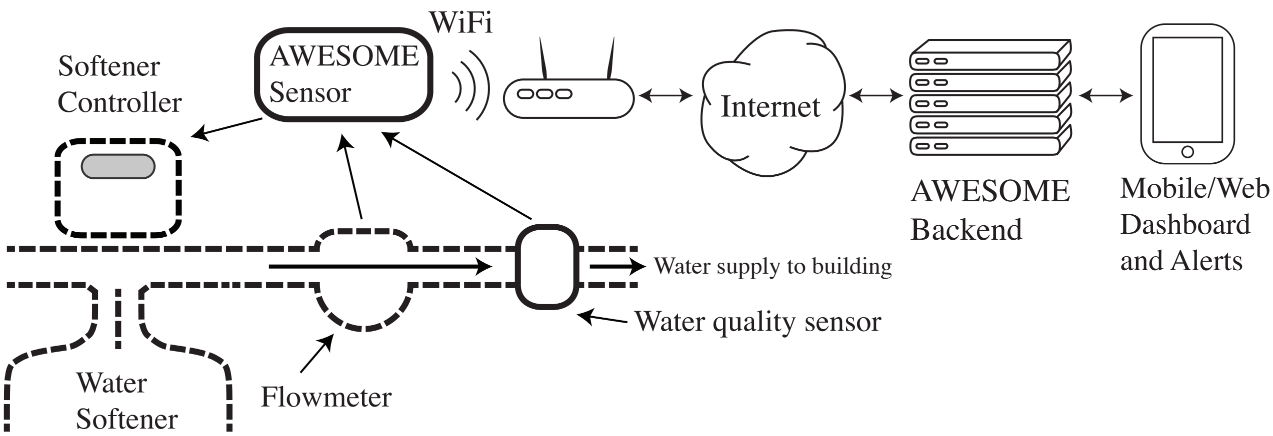

Figure 1: Dataflow diagram for the AWESOME system. Components with solid lines are part of AWESOME , and components with dashed lines are part of the existing softener system.

weakly bound to the medium, are released into solution. The Calcium and Magnesium ions replace the Sodium on the filtration medium. Eventually, the surface of the filtration medium becomes saturated with Calcium and Magnesium ions, and it can no longer treat water. It must be regenerated by flushing with concentrated salt brine, which replaces Calcium and Magnesium ions with Sodium, preparing the softener to treat more water.

However, large amounts of water pollution are not necessary to provide soft tap water in buildings. Anecdotally, we have observed that many softeners are configured incorrectly, leading to excessive backwashing and wasteful salt consumption. Furthermore, there is a lack of information about water softener salt consumption that leads to malfunctioning systems.

Other types of water treatment systems, such as reverse osmosis and deionizers, also discharge pollutants. We chose to work with water softeners in this study because we have access to several large softener systems on campus. In principle, our techniques could be ported to work with various other treatment systems.

In our pilot deployments in three UW residence halls (Chadbourne, Sellery, and Leopold), we used AWESOME to initiate water softener regeneration based on water quality measurements. Chadbourne Hall is an eight-story residence hall that houses 600 students and a dining facility. Sellery Hall a six-story residence hall housing 1,200 students that does not include on-site dining. Leopold Hall is a four-story residence hall housing 200 students. As a result, *we reduced the salt used by an average of 27%.*

1.1 The AWESOME System and How it Addresses Inefficiencies

Most water treatment systems use open-loop control, meaning that there is no sensor on the outgoing treated water to inform decisions about when to regenerate the filtration medium. Instead, the control systems usually use simple time-based (regenerate once a week on Tuesday) or flow-based (regenerate every 1000 gallons) schedules.

Inconsistencies in the incoming water quality, degradation of the filtration medium, or variation of the water usage pattern of the building can make the system unstable. In Madison, municipal water is drawn from a pool of 17 wells, each of which supplies water with different hardness. The variation ranges from 20 to 30 grains per gallon[2] [4].

Water softeners tend to be set to use more salt than strictly required for several reasons. Adding too much salt doesn't negatively impact the piping of the building and won't be noticed by the user. The cost of maintenance, especially in large buildings, for cleaning up lime is higher than the cost of additional bags of salt. This is primarily due to the cost of the labor involved. However, gross overusage of salt, as we have found to be common among the buildings we've studied, is also expensive in the long run.

Because water softener systems are typically located in remote service areas, problems can go undetected for months or years. Furthermore, existing softener controllers do not typically display enough information for maintenance staff to detect misconfigurations.

Misprogrammed controllers can cause inefficiencies in the softener system by allowing either too much or too little water to flow through the resin bed before it is depleted. Regenerating the resin bed too early can result in excessive salt use and increased Sodium Chloride pollution. Regenerating too late can cause the resin bed to become totally depleted, resulting in hard water being distributed to the building. This can cause pipes to clog and eventually destroy equipment. Unfortunately, many existing softener controllers are difficult to program. Even experienced maintenance personnel can make mistakes in programming the softener controllers, resulting in over-softened or under-softened water. Reprogramming softeners is routinely required after a power failure or after a changeover from daylight savings time.

Low flow rates through a softener tank can result in nonuniform flow of water through the filtration medium, causing some regions to deplete more quickly than others. Under low-flow conditions, water tends to move through a column in the center of the softener tank, quickly depleting the medium in that region. For this reason, it may be necessary to regenerate the water softener more frequently during periods of low consumption.

[2] A *grain* is a unit of mass, approximately equal to 65 mg. Water hardness, the concentration of Calcium and Magnesium ions, is conventionally measured in grains per gallon.

Variability of incoming water quality creates obstacles when provisioning a treatment system because it is difficult to predict the amount of water a system can treat before it needs to be regenerated. Since the hardness of incoming water commonly varies by 10% or more over the course of days or weeks, softeners are often configured to deal with the worst case hardness. Even when incoming water has relatively low hardness, the softener system will still be regenerated on the same schedule as under maximum hardness conditions because stock water softeners do not have any way of sensing water quality.

Proposed AWESOME System and its Advantages

When a water softener regenerates, it flushes its filtration medium with salt and water, which must be sent down the drain. The more frequently it regenerates, the more water and salt the system uses. Existing water softener control systems perform poorly because they cannot adapt their regeneration schedules to changes in water quality and usage. Since stock controllers are typically configured to regenerate more frequently than necessary to maintain soft water, their consumption of salt and water is wasteful.

Our goal in developing AWESOME was to design a system which will regenerate the softener tank only when strictly necessary. In so doing, the softener can treat more water between regenerations, which conserves resources.

By sensing the water after it has been treated by the water softener, AWESOME determines when it is appropriate to regenerate. However, we found that it is difficult to distinguish hard water from soft water using data gathered from just one sensor because water quality sensors have a tendency to drift over time. Instead, our backend ties together noisy sensor data using domain knowledge to generate a reliable regeneration output.

1.2 Key Contributions

The work described in this paper covers both significant research and engineering problems, making the following contributions:

- **We develop an adaptive control algorithm to predict when the softeners need to be regenerated based on our sensor inputs.**
- **We design a closed-loop system to measure, record, and respond to water quality sensor inputs.**
- **We deploy AWESOME in three locations over the course of two years.**
- **We evaluate our algorithm using data collected from our deployments, demonstrating that our methods can reduce the salt used by 15-45%.**

2. BACKGROUND

The main function of a water softener is to protect water heaters and fixtures from lime buildup. When hard water is heated, dissolved Calcium becomes insoluble and forms scale deposits on pipes and fixtures, requiring expensive and inconvenient repairs. Particularly in large commercial and industrial operations where hot water or ultrapure water has a mission-critical function, water softeners are often an indispensable component of the water supply chain.

The main working component of a water softener is a large tank filled with tiny plastic beads. The beads are made of a specialized plastic resin that is chemically engineered to attract metal ions—in particular, Calcium (Ca^{2+}), Magnesium (Mg^{2+}), and Sodium (Na^{+}). The beads are initially coated with Na^{+} ions, as shown in Figure 2 (a). As hard, Calcium-containing water flows past the resin beads, the Calcium ions adhere to the surface of the beads, releasing Sodium ions into solution, as shown in Figure 2 (b) and (c). This Sodium-rich *soft water* supplies boilers, deionizers, and other sensitive devices in the building. Soft water is safe for equipment because Sodium ions have far less propensity to precipitate out of solution than Calcium and Magnesium ions, meaning that soft water will not leave behind scale on the equipment it contacts.

In normal operation, before the resin has been depleted, water softeners remove nearly all Calcium and Magnesium ions from the water they treat. Before depletion, water coming out of the softener should have a Calcium ion concentration of nearly zero, a feature that AWESOME will take advantage of during backend data processing.

After the water softener has treated a large volume of hard water, its resin beads are fully saturated with Calcium ions, as shown in Figure 2 (d). Once this happens, the resin must be *regenerated* before it can treat more water. To regenerate the softener, the resin medium is soaked in a concentrated solution of table salt (NaCl) and water. Sodium ions in the salt brine replace Calcium ions on the surface of the resin beads, returning the filtration medium to the diagram shown in Figure 2 (a).

A controller, present on all water softeners, is responsible for deciding when the filtration medium needs to be regenerated. However, even most modern controllers do not use water quality sensors to determine when the resin needs to be regenerated. Instead, they track the total water that has flowed through the filtration medium since the last regeneration, and trigger a regeneration after the total flow has reached a pre-defined level. We believe that this is because most water quality sensors are expensive and produce noisy data which is not easily interpretable by a simple controller. Instead, controllers guess at when the filtration medium needs to be regenerated. If the controller regenerates too late, it will allow the resin to deplete, sending hard water to the building, which could be damaging. If it regenerates too early, it will use more salt than necessary, but it will not destroy any building components.

To err on the side of caution, most controllers regenerate the filtration medium long before it is actually depleted in order to avoid supplying hard water to the building, which could potentially require costly repairs. Since the same amount of salt is needed to regenerate a softener that is 80% or 99% depleted, it is wasteful to regenerate the filtration medium early.

3. DESCRIPTION OF AWESOME

The AWESOME system consists of custom hardware and software components that measure water quality, process the recorded data, and respond.

3.1 Sensing Water Quality With AWESOME

AWESOME uses water flow and quality sensors to track the status of a water softener in real time. We constructed custom hardware (shown in Figure 3) to gather and record

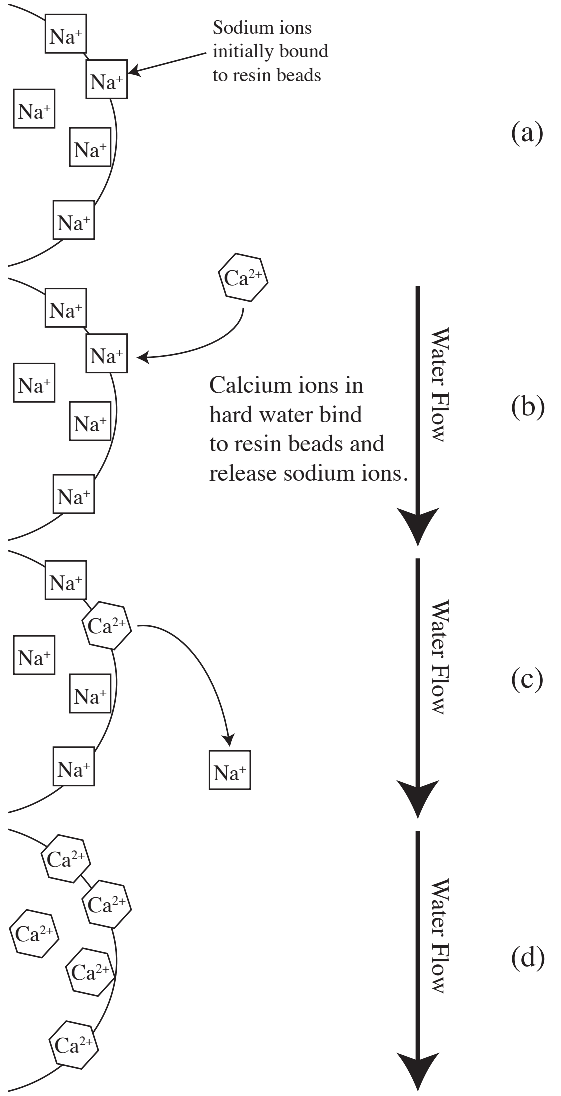

Figure 2: A diagrammatic overview of the way a water softener functions, explained in Section 2.

data from water softeners. The data it collects is transmitted to a remote database for storage and post-processing (described in Section 3.2). When backend algorithms running on the remote database server determine that the water softener needs to be regenerated, they transmit a signal to AWESOME, which forces the water softener to regenerate. A diagram of the dataflow used by AWESOME is shown in Figure 1.

The AWESOME board includes an embedded computer and several sensors to track the health of a water softener system:

The AWESOME embedded computer is a custom device featuring a Freescale ColdFire microcontroller operating at 60 MHz. The AWESOME board is a derivative of Emonix [12], and it includes many of the same software features—multithreaded tasking, POSIX-like programming environment, BSD sockets, etc. It also includes an xBee WiFi network controller which allows it to directly connect to the building's network. The board includes six programmable digital/analog sensor inputs through which sensors can be connected to collect data. An on-board 32 kbyte SRAM device is included to allow the board to cache data samples in the event of a network outage.

The AWESOME embedded device is responsible for sampling sensors, preprocessing and caching the data, and relaying it to the backend database through the building's WiFi network. When the backend determines that a regeneration is required, it sends a message to the the embedded device, which initiates a regeneration on the water softener.

A Calcium Ion Selective Electrode (ISE) produces a voltage signal that is proportional to the Calcium concentration of the water sample. These devices are extremely sensitive to temperature and other factors, and their outputs are known to drift over time. By itself, the Calcium ISE cannot tell us the concentration of Calcium in our water samples, so we need to use auxiliary sensors to augment our data. Manufacturers of Calcium ISEs warn that their products should not be held in test solutions for long periods of time. This concerned us because our approach requires that we use ISEs to continuously gather water quality samples, with the probes themselves constantly submerged in water samples. We expected to have to replace the sensors after a few months of data collection since we were not using the ISEs according to the manufacturer's recommendations. However, we have had several deployments of AWESOME running continuously for two years, and we have not had any failures in the ISEs.

We believe that the reason the ISE manufacturer publishes conservative guidelines about how to use its product is that the ISE is a scientific instrument intended for use in a laboratory setting. It would normally be required to make precision measurements of Calcium ion concentration in the measurement solution. However, since AWESOME only needs to make coarse-grained estimates of hardness—classifying the measurement solution as hard or not hard—an ISE output that is only accurate to within 10-20% is acceptable. Our backend processing can take noisy or unreliable sensor data and produce a reliable regeneration decision.

In Section 2, we mentioned that both Calcium and Magnesium ions contribute to total water hardness. Both are always present in hard water, though their ratios can vary depending on the source of the water. Our sensor system measures only Calcium ion concentration, since it is an indicator of total water hardness. Water that has a significant concentration of Calcium ions will also have a significant concentration of Magnesium ions.

A Water Temperature Sensor is used in combination with the Calcium ISE to estimate the Calcium ion concentration in the water sample. We use the Nernst Equation [9] to combine the Calcium ISE voltage and the water temperature to get an estimate of the Calcium ion concentration. Because the Calcium ISE's output voltage is prone to drift over time, we cannot use the raw output of the Nernst equation to make decisions about whether or not the water is hard.

A Flowmeter is already installed in most commercial water softeners. Flow rate data is used by the stock water softener controller to decide when the softener should regenerate—in fact, this is the only sensor included by default on most commercial softener systems. Fortunately, all water softener flow meters use the same interface to com-

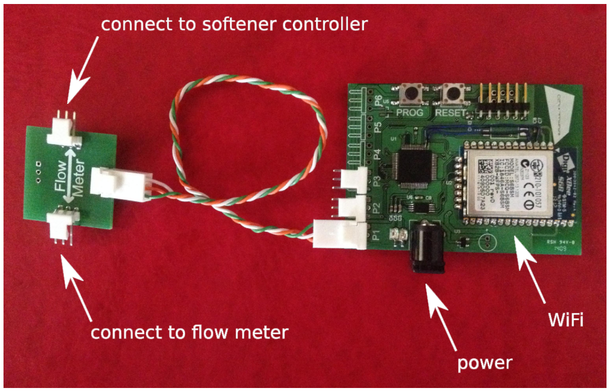

Figure 3: Photograph of an AWESOME board.

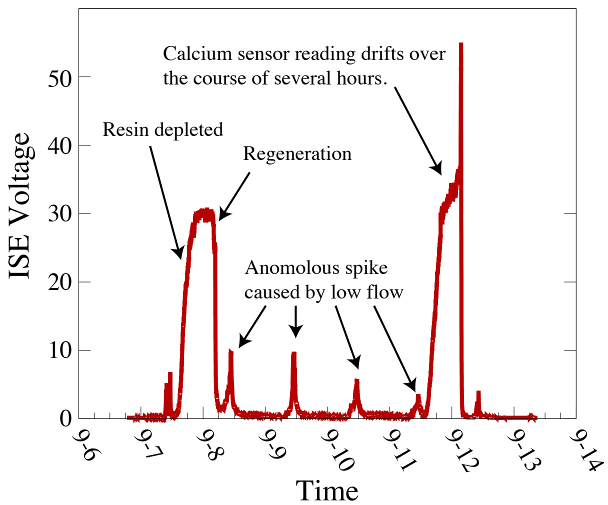

Figure 4: Example trace of some raw hardness data gathered from the Calcium ISE. At night, when the flow through the water softener drops, creating anomalous spikes that could be confused for hard water output. On 9-12, the sensor output drifts over the course of several hours, even though the water hardness remains constant.

municate with their controllers. AWESOME can intercept and record the flow rate signal.

3.2 Backend Data Processing in AWESOME

We now discuss several approaches to processing the data gathered by our sensors. The goal of our data processing algorithms is to correctly and quickly identify filtration medium depletion. To do so, the AWESOME embedded computer gathers data from the sensors at intervals of five minutes and passes that data to the backend for analysis. The backend clusters the each set of sensor readings into a feature vector and feeds it into our learning algorithm to decide whether the filtration medium needs to be regenerated.

As we mentioned in Section 2, the main challenge we face is that our water quality sensors produce noisy data. In particular, the output of our sensors may drift over time, and variability in flow may result in errant spikes in the sensor

Feature	Formula
ISE Output	$[Ca^{2+}] = e^{k(V-a)}$
Flow Rate	$f[gal/minute]$
Temperature	$T[^\circ C]$
Derivative of ISE Output	$\frac{d[Ca^{2+}]}{dt}$
Normalized ISE Output	$[Ca^{2+}]\left(1 - e^{-kf/T}\right)$

Table 1: Features used by our adaptive control algorithm to determine whether the softener system needs to be regenerated.

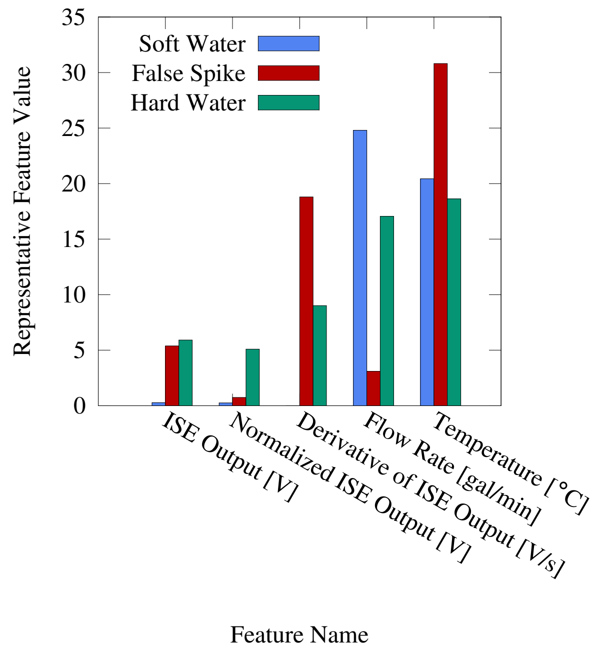

Figure 5: A demonstration of how the features used by the SVM differ for soft water, truly hard water, and false spikes (noise).

output, even though the softener is not actually producing hard water.

Figure 4 shows an example output trace of the data gathered by our hardness sensor before it has been processed by our backend algorithms. As we mentioned in 2, the Calcium ion concentration of the softened water should be nearly zero during normal operation of the softener. As the resin depletes, the Calcium ion concentration of the treated water transitions from zero to 20-30 grains per gallon (whatever the hardness of the municipal water is), usually over the course of a few hours. During resin depletion, the ISE voltage will increase synchronously with the Calcium ion concentration. It is these gradual increases in Calcium ion concentration that AWESOME seeks to identify. However, the true Calcium ion concentration may be coupled with noise, generated by the ISE:

- **Low flow rates**, particularly during the night, confuse the sensor because stagnant water in the pipes heats up, disrupting our readings.
- **Drifting voltage output of our Calcium ion selective electrode** caused by degradation of the ion-

selective membrane and the reference electrode can generate inaccurate readings.

The goal of our data processing algorithm is to identify true hard water samples correctly and reject anomalous readings. We fed the data gathered from our three pilot deployments into two learning algorithms (SVM and thresholding, described below), and evaluated the quality of predictions made by each algorithm.

Ground truth was gathered manually in each building using a chemical hardness test kit [8]. The test kit we used provides a reliable measurement of water hardness, but it is a time-consuming manual process which is not practical to automate.

The algorithms in this section were trained and tested by disabling regenerations on the water softeners in the pilot buildings and allowing their filtration media to deplete. Once the filtration media depleted, the outgoing water became hard, as shown in Figure 4. We allowed our sensors to collect data during the process of depletion, and we independently verified that the softener produced hard water by using a manual water hardness test kit [8]. Because we did not want to damage the buildings' hot water heaters, we only gathered data for a few hardness events in each building. For this reason, we have at most seven hard water events on which to train and test our algorithms.

Following the data collection, we disassembled one building's hot water heater to find out if the hard water events caused any damage. The building maintenance staff said there did not appear to be excessive lime buildup on the heating coils, leading us to believe that our experiments did not have a detrimental effect on the building's water system.

Why Simple Thresholding Does Not Work Well

In the simplest data processing approach, we used a threshold on the ISE voltage to determine whether the output water was hard. Above some user-defined output of the Calcium ISE, the water was considered hard, and a regeneration was initiated.

An advantage of this technique is that its performance is intuitive and predictable—an regeneration will always be initiated when the ISE voltage rises above the threshold, which can be made as low as desired. However, thresholding is not good at rejecting noise and drift in the ISE voltage output. Under low flow conditions, we frequently see spikes (usually at night, depicted in Figure 4) which the thresholding technique may interpret as medium depletion and erroneously regenerate the softener, wasting salt. Alternatively, we could increase the threshold above the maximum height that we expect for one of these erroneous spikes. However, this approach is also suboptimal because it does not regenerate the softener until the water has become very hard, which could damage pipes and equipment. We need a better approach which can accurately detect filtration medium depletion early, before the softener starts producing extremely hard water.

SVM-Based Technique

Using an SVM, we can integrate readings from multiple sensors to detect hard water early. SVM takes as input a set of labeled training vectors (sensor readings labeled with ground truth). The underlying SVM algorithm [5] finds a linear function separating feature vectors in the two classes—hard water sample vectors lie on one side of the function, and soft water vectors lie on the other. Unknown feature vectors are classified based on which side of the linear function they fall.

For each new feature vector the backend evaluates, a classification decision is generated by the SVM. If it classifies a new feature vector as being from a hard water sample, it immediately sends a message to the AWESOME embedded computer to initiate a regeneration. To reduce noise further, we considered adding an extra level of processing here—for example, requiring two or more consecutive hard water samples in a row before initiating regeneration. We found that post-processing of the SVM data did not improve noise immunity because the SVM classification decisions did not produce spurious false positives. In fact, SVM post-processing made our backend processing perform worse because it increased the delay between true medium depletion and regeneration.

The features we used as inputs to our learning algorithms are enumerated in Table 1. The first three features—ISE output[3], flow rate, and water temperature—are measured directly from sensors. The last two, computed as functions of the first three, were added to reject noise and drift of the Calcium ISE. Figure 5 shows how the features we chose for the SVM vary depending on whether the sample water is soft, hard, or a false spike.

The ISE output is a function of the water's Calcium ion concentration, and it should generally be low for soft water and high for hard water. Since we want to detect hard water before the Calcium ion concentration (ISE voltage) gets too high, we need to add more feature vectors to distinguish between false spikes and hard water.

To distinguish between hard water and false spikes, we also measure flow rate. False spikes generally occur because of low flow rate through the water softener. In addition, increased water temperature has a tendency to cause erroneous high ISE readings.

Perhaps the most important characteristic we studied is the shape of the ISE curve as it increases. In Figure 4, it is visually evident that the shape of a false spike is different from the shape of a true hardness event. We can take advantage of the difference in order to detect water hardness events early because we can identify the shape of the curve before we can identify its maximum amplitude. Also, measuring the shape of the ISE output gives us information about *how it changes with time*, which SVMs are generally not good at discerning.

We used the derivative of the Calcium ISE voltage output trace to measure shape. In Figure 5, it is clear that hard water events have a smaller derivative than false spike events. This is true almost universally in the dataset we collected from our pilot deployments. The reason for this is that at the end of a filtration cycle when the medium is almost depleted, the medium goes through an intermediate phase in which it slowly becomes less efficient at removing Calcium ions. That intermediate phase generally lasts for several hours. Erroneous spikes, on the other hand, are caused by changes in water use patterns which generally happen over much shorter time intervals.

---

[3] The *estimated* Calcium ion concentration $[Ca^{2+}]$ is computed as a function of the voltage output of the Calcium ISE (V). We augment this input with other sensor data to provide a more reliable estimate of whether or not the filtration medium needs to be regenerated.

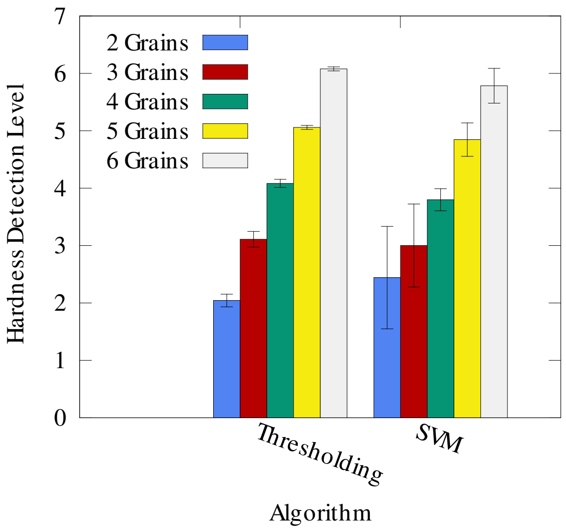

Figure 6: The hardness at which each algorithm detected filtration medium depletion. We show the detection level for each of five different training levels. Our goal is to detect medium depletion at the lowest possible hardness level while minimizing the number of false positive identifications.

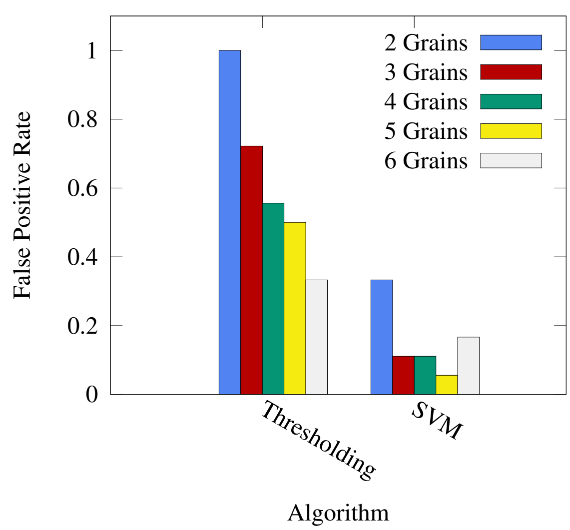

Figure 7: False positive rates for the learning algorithms evaluated. Each algorithm was tested with five different training thresholds. In general, lower training thresholds result in more false positives.

By tracking multiple sensor inputs as well as the shape of the ISE output, we can much more accurately distinguish between truly hard water and false sensor spikes. In the next section, we evaluate our SVM-based technique, comparing it to simple thresholding. We demonstrate that we can detect and respond to resin depletion before water gets too hard to threaten building systems like hot water heaters and faucets.

4. EVALUATION

In this section, we test our learning algorithms for robustness, showing that they can correctly identify filtration medium depletion in a water softener. We go on to discuss the results of pilot deployments of AWESOME, showing that it can reduce salt consumption of water softeners by 15-45% compared to stock water softener controllers.

4.1 Learning Algorithm Evaluation

To evaluate our learning algorithm, we feed data gathered from our pilot deployments into each of several candidate algorithms and test their performance. Our goal is to find out

1. How quickly can each algorithm identify depletion of the filtration medium?
2. How resistant is each algorithm to anomalous sensor inputs (noise)?

The answer to the first question will give us insight that is similar to a false-negative rate. In our application, false negatives are not very meaningful because AWESOME repetitively samples water from the same softener. If we get a false negative on one sample, we are almost certain to get a true positive on a subsequent sample. For AWESOME, a more meaningful figure of merit is the amount of time it takes the system to detect resin depletion. The longer it takes AWESOME to detect that the resin has been depleted, the harder the water that is produced by the softener, potentially endangering building systems.

The answer to the second question will tell us the false positive rate. When AWESOME produces a false positive, it regenerates the softener unnecessarily, wasting salt.

To identify filtration medium depletion, we will necessarily have to allow the medium to deplete and start producing hard water. Our goal is to minimize the amount of hard water that needs to be produced before our algorithm can identify that the medium is depleted. If we supply too much hard water to a building before initiating a regeneration, we could damage the water heater and other building systems.

SVM is a supervised learning algorithms, meaning that it requires a training phase in which input data labeled with ground truth is fed to the algorithm. When constructing the training sets for SVM, we need to label each feature vector with ground truth. This labeling process is subjective because the water hardness—concentration of Calcium and Magnesium ions—is a real number which must be converted to a binary value (0 or 1). So to label our training data, we must apply a threshold, which we call a **training threshold**, below which water is considered to be soft (0), and above which it is considered to be hard (1).

The threshold we apply will affect our algorithms' ability to distinguish between sets. If the threshold is too low, the algorithm may classify all unknown inputs as hard. If it is too high, we may miss all hardness events. Our goal in AWESOME is to produce training sets that will identify filtration medium depletion early (before the outgoing water has become too hard) while rejecting noise. In the parlance of hypothesis testing, we wish to optimize the ratio of true positives to false positives.

The lower we set the training threshold, the earlier our algorithms will in general be able to detect filtration medium depletion. Figure 6 shows a plot of the average outgoing

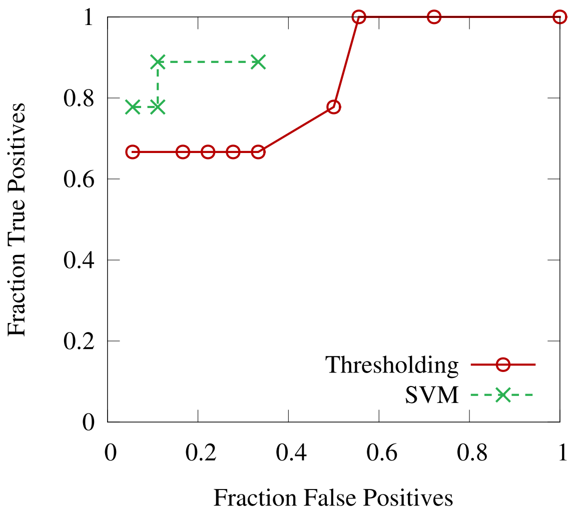

Figure 8: Receiver operating curve for the learning algorithms we analyzed. Algorithms that perform well will have ROCs that are close to the upper left corner of the graph. In that region, the algorithm generates a low fraction of false positives and a large fraction of true positives.

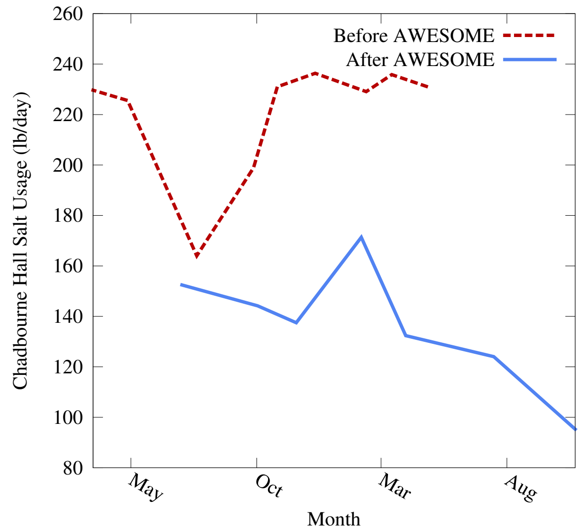

Figure 9: Salt usage in Chadbourne Hall over a two-year timespan. Salt consumption declined dramatically following the installation of AWESOME (solid blue line). This data is based on salt deliveries made by Kreger Salt Sales, the supplier of softener salt for Chadbourne Hall.

water hardness level at which each algorithm can detect medium depletion for several different training thresholds. The x-axis lists the algorithms we studied, and the y-axis gives the average water hardness required for each algorithm to identify a filtration medium depletion. To be practical, the detection level should be below five grains per gallon because allowing the hardness to get any higher would damage building systems.

In our evaluation, we trained and tested each algorithm to five different training thresholds (2-6 grains hardness). As one would expect, the higher the training threshold, the higher the output hardness needs to be for an algorithm to identify filtration medium depletion.

However, at low training thresholds, the learning algorithms tend to produce more false positives. Figure 7 shows the false positive rates on the y-axis with the algorithms we tested on the x-axis.

To strike a balance between minimizing false positives and minimizing the hardness detection level, we find that it is best to train the learning algorithms to identify depletion at three grains hardness.

The receiver operating curve (ROC) shown in Figure 8 shows the tradeoff between false positives and true positives for SVM and thresholding. On the y-axis, true positives are plotted as a function of false positives. These are parametric curves in which different hardness thresholds are used to train each control algorithm. Ideally, we would like our algorithms to operate as close to the upper left corner of the graph as possible: we want to increase the percentage of true positives as much as possible while keeping the fraction of false positives as low as possible. In the graph, SVM performs much better than thresholding because we only need to increase the proportion of false positives to 30% in order to achieve 90% true positive detection

4.2 Cost Savings

AWESOME was installed in Chadbourne Hall on the UW campus in July 2013, replacing the stock softener controller. Using our adaptive control algorithms to control the water softener, we were able to save $3369 in salt purchases over a twelve-month period ending July 2014. Table 2 shows the salt savings achieved after installing AWESOME.

For the pilot deployments discussed in this section, we used a thresholding algorithm to detect and respond to filtration medium depletion. To reduce the number of unnecessary regenerations caused by false positives, we used a threshold of ten grains per gallon. A downside of this approach is that it does not allow us to detect resin depletion early. We evaluated other learning algorithms (kNN, SVM, and neural network) in order to detect resin depletion early while rejecting noise.

Reduced Salt Usage

Figure 9 shows the average salt used by the softener system in Chadbourne over a three-year timespan starting in 2012. The dashed line represents the salt consumption of the building's water softener in pounds per day before installing AWESOME. The solid line shows salt consumption for the following year after AWESOME was installed. Following the installation of AWESOME in July 2013, the building's per-day salt consumption fell because the number of gallons between regenerations increased by more than 30%. Salt consumptions in the other buildings in which we piloted AWESOME also dropped. Figure 10 shows a comparison of the salt consumed by the water softeners in our three pilot buildings on the UW campus before and after installing AWESOME. Chadbourne Hall showed the most improvements because that building's water softener sys-

[5]These computations are based on orders of pallets containing 49 × 50-lb bags of salt, costing $7.64 per bag.

Building	Year 1	Year 2	Savings
Chadbourne	68,600 lb ($10,482)	46,550 lb ($7,113)	22,050 lb ($3369)
Sellery	41,445 lb ($6,333)	31,500 lb ($4813)	9,945 lb ($1520)
Leopold	17,000 lb ($2,597)	14,700 lb ($2,246)	2,300 lb ($351)
	127,045 lb ($19412)	92,750 lb ($14172)	**34,295 lb ($5,240) 27% saved**

Table 2: Comparison of salt consumed in the three pilot buildings before and after installation of AWESOME[5].

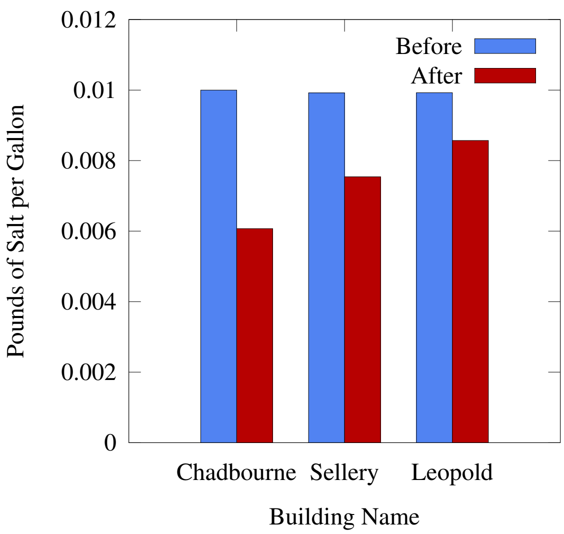

Figure 10: Amount of salt used per gallon of water treated in each of the three buildings where we deployed AWESOME.

tem was configured to regenerate *nearly twice as frequently* as it should have, by comparison to factory parameters. The other buildings, though configured correctly according to factory parameters, were still using too much salt because their stock controller settings were far too conservative.

We think that the salt consumption in Chadbourne hall improved more relative to the other two pilot buildings because the hardness of the municipal water supply to that building is the lowest.

Other Benefits of AWESOME

With the reduced volume of salt used to flush the water softener system in Chadbourne, there was a proportionate decrease in Chloride and Sodium pollution.

AWESOME makes it possible for maintenance staff to monitor the status of the water softener system from any computer. Water softeners are often installed in less accessed parts of the building as matter of aesthetic principle, however, this often makes access inconvenient and unwieldy. For example, before installing AWESOME , the only way to be sure the system was working properly was to climb a 6-foot ladder in a mechanical room in the basement of Chadbourne. Even then, it was not possible to get continuous historical data about the water flow through the softener.

5. RELATED WORK

Most existing work in the area of smart water systems focuses on trying to reduce water consumption in the context of human behavior. As such, much of the existing work deals with points of interaction between building occupants and the water distribution system, such as faucets, showers, and toilets. Hydrosense [6], NAWMS [11], and Driblet [3] are systems for identifying the sources of water consumption in a home using a single point of monitoring with the goal of classifying human behavior based on water usage. Winkler et. al. [17] devise a system to optimize the flow of water for irrigating lawns.

Another theme is to optimize the distribution of hot water throughout a residence. Circulo [7] and Hot Water DJ [14] aim to reduce the energy and water wasted while waiting for hot water to arrive at a sink or shower in a home. The Energy-Water Nexus [10] studies the relationship between energy and water consumption in a small collection of commercial buildings. Ranjan et. al. [15] leverage the energy-water nexus concept in their system by tracking localized energy and water consumption patterns to improve the accuracy of assignment of water consumption to building occupants.

Our work differs from the existing body of knowledge in that it addresses issues related to water treatment systems rather than water distribution systems.

6. CONCLUSION

In this work, we demonstrated that using closed-loop feedback control makes it possible to significantly reduce the amount of salt consumed by water softeners. We discussed the problems with simply applying a threshold to the output of a Calcium ion selective electrode, and we proposed an alternative algorithm that uses an SVM with feature vectors that include the outputs of several sensors. We validated our techniques by deploying AWESOME in three pilot buildings over the course of two years, measuring the amount of salt used in each deployment. In the pilot deployments, we found that it is possible to reduce salt consumption of softeners by at least 15%, saving thousands of dollars per year on salt purchases alone.

As the availability of water continues to diminish, the efficiency of water treatment and distribution systems will become increasingly important. California [13], New Mexico [16], and Puerto Rico [2] are all experiencing severe droughts, and they do not seem to have a clear path to conservation. With AWESOME, we hope to provide building managers and municipalities with tools to understand and control sources of inefficiency in water treatment and distribution systems.

7. REFERENCES

[1] Central arizona salinity study. Technical report, US Bureau of Reclamation, June 2015. http://www.usbr.gov/lc/phoenix/programs/cass/cass.html.

[2] L. Alvarez. In drought, puerto rico rations water, setting off a collection frenzy, 2015. New York Times.

[3] Driblet Labs. Driblet, 2015. http://driblet.co.

[4] R. Erickson and K. Lake. Third annual chloride progress report. Technical report, Madison Metropolitan Sewerage District, June 2013.

[5] D. C. et. al. scikit-learn: Machine learning in python, 2015. http://scikit-learn.org.

[6] J. E. Froehlich, E. Larson, T. Campbell, C. Haggerty, J. Fogarty, and S. N. Patel. Hydrosense: Infrastructure-mediated single-point sensing of whole-home water activity. In *Proceedings of the 11th International Conference on Ubiquitous Computing*, Ubicomp '09, pages 235–244, New York, NY, USA, 2009. ACM.

[7] A. Frye, M. Goraczko, J. Liu, A. Prodhan, and K. Whitehouse. Circulo: Saving energy with just-in-time hot water recirculation. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 16:1–16:8, New York, NY, USA, 2013. ACM.

[8] Hach. Total hardness test kit, model 5-b, 2015. http://www.hach.com/total-hardness-test-kit-model-5-b/product?id=7640219508.

[9] D. C. Harris. *Exploring Chemical Analysis*. W. H. Freeman, 2012.

[10] J. Kadengal, S. Thirunavukkarasu, A. Vasan, V. Sarangan, and A. Sivasubramaniam. The energy-water nexus in campuses. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 15:1–15:8, New York, NY, USA, 2013. ACM.

[11] Y. Kim, T. Schmid, Z. M. Charbiwala, J. Friedman, and M. B. Srivastava. Nawms: Nonintrusive autonomous water monitoring system. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, SenSys '08, pages 309–322, New York, NY, USA, 2008. ACM.

[12] N. Klingensmith, D. Willis, and S. Banerjee. A distributed energy monitoring and analytics platform and its use cases. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, BuildSys'13, pages 36:1–36:2, New York, NY, USA, 2013. ACM.

[13] I. Lovett. Californians keep up with Joneses' water use, 2014. New York Times.

[14] M. A. Prodhan and K. Whitehouse. Hot water dj: Saving energy by pre-mixing hot water. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '12, pages 91–98, New York, NY, USA, 2012. ACM.

[15] J. Ranjan, E. Griffiths, and K. Whitehouse. Discerning electrical and water usage by individuals in homes. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, BuildSys '14, pages 20–29, New York, NY, USA, 2014. ACM.

[16] M. Wines. Mighty rio grande now a trickle under siege, 2015. New York Times.

[17] D. Winkler and A. E. Cerpa. Distributed independent actuation for irrigation control. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, BuildSys '14, pages 148–151, New York, NY, USA, 2014. ACM.