

DoubleDip: Leveraging Thermoelectric Harvesting for Low Power Monitoring of Sporadic Water Use

Paul Martin
University of California,
Los Angeles
pdmartin@ucla.edu

Zainul Charbiwala
University of California,
Los Angeles
zainulcharbiwala@ucla.edu

Mani Srivastava
University of California,
Los Angeles
mbs@ucla.edu

Abstract

We present DoubleDip, a low power monitoring system for enabling non-intrusive water flow detection. DoubleDip taps into minute thermal gradients in pipes for both replenishing energy reserves and performing low power wakeup. One of the remaining issues with wireless water monitoring in residences and offices is that current solutions require installing sensor nodes with access to electrical wiring or replacing batteries frequently. DoubleDip (DD) significantly extends the lifetime of vibration-based non-intrusive water flow sensors by harvesting thermal energy from hot pipes wherever accessible. DoubleDip requires less than an inch of exposed metal pipe to attach a coupler for gathering sufficient energy to power itself, in some cases, into perpetuity. We observe that water use in homes and offices is incredibly sporadic, making continuous monitoring both impractical and wasteful. Instead, DD puts a thermoelectric harvester into double duty. It uses thermal gradients not only for gathering energy but also for extremely low power ($< 1\mu\text{A}$) wakeup. In this paper, we describe the DoubleDip design and demonstrate that thermoelectric wakeup is essential for longevity and accuracy. Since DD wakes up from its low power state only when there is a water flow event, it replenishes the energy it uses in sensing and transmitting data by the energy it harvests from the corresponding heat gradient. While DD nodes installed on cold water pipes harvest far less than those installed on hot water pipes, our pilot deployment over four weeks and five locations suggests that thermoelectric wake up is only slightly worse in latency for cold water monitoring and there is sufficient energy harvested from the hot water that it can be shared to extend the lifetime of nearby cold water nodes too.

Categories and Subject Descriptors

C.2.4 [Computer Systems Organization]: Computer Communication Networks—*Distributed Systems*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SenSys'12, November 6–9, 2012, Toronto, ON, Canada.
Copyright © 2012 ACM 978-1-4503-1169-4 ...\$10.00

General Terms

Design, Measurement, Experimentation, Performance

Keywords

Sustainability, energy harvesting, thermoelectric, water monitoring, low power

1 Introduction

Water monitoring in homes and office buildings promises to help alleviate some of the water reserve burdens we face today [27]. Reliable measurements of consumption are essential for evaluating the impacts of conservation projects on urban water demand [20]. Providing consumers with a fixture-level understanding of their usage through appropriate interfaces provides opportunities to make informed decisions about investments in high efficiency fixtures and appliances [11] and serves to encourage them toward more effective conservation. Sadly, a significant portion ($\sim 30\%$) of water use is actually wasted through leaks and oversight [20] which might have been detected early with continuous water monitoring coupled with intelligent analytics.

Gathering continuous usage data with fixture-level granularity and accuracy, however, should not be cumbersome. Presently, water flow data is only expected to be available in aggregate form for the entire household or building through smart water meters [25]. Previous work has shown that disaggregating this data to individual fixtures is non-trivial due to the similarity of fixtures in homes and offices [31]. Disambiguating across fixtures is especially hard in office buildings due to extensive plumbing infrastructure that may feed off just one meter. Conversely, installing in-line water flow sensors on each fixture or even arterial branches in the plumbing network requires much retrofitting, electrical wiring to power the sensors, or maintenance to change batteries. Consumers ideally desire to conceal the wireless monitoring system along with existing plumbing, but nodes that primarily rely on a fixed energy supply in particular have to remain accessible.

Improvements in energy efficiency of sensor nodes and power harvesting electronics have balanced power requirements to a degree wherein it is now feasible to tap into energy reserves that would otherwise have yielded only marginal improvements to system lifetime. For example, low power wireless motes such as Helimote [19] have had remarkable success using ambient solar radiation and adaptive scheduling mechanisms to power outdoor sensor networks. Yerva et

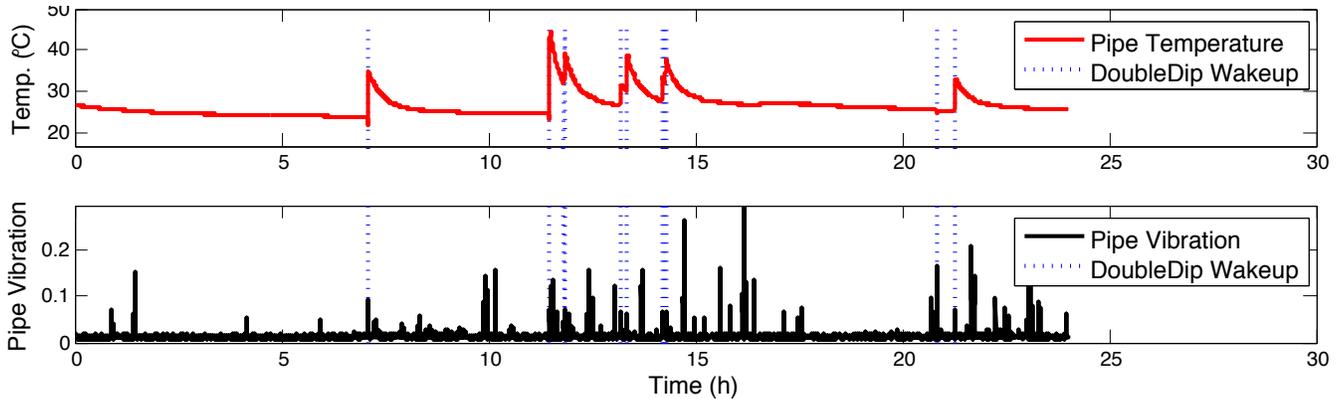


Figure 1: Sample data of pipe temperature, wakeup triggers from DoubleDip, and pipe vibration. Spurious vibrations would cause frequent false wakeups with vibration-based LPL even though water use is sporadic.

al. [33] were successful in using indoor photovoltaic (IPV) harvesting to achieve very low duty cycle sensing from ambient indoor lighting while maintaining mesh network functionality. Additionally, as sensor nodes decrease in size to allow for sensing of remote and physically constrained phenomena and their numbers increase for higher coverage, energy harvesting techniques move from luxury to necessity [33].

Hot water accounts for over 30% of water usage [20] and over 20% of household energy consumption for heating [1]. Of this, over 80% is used in faucets, showers, and baths [20]. Short uninsulated pipe segments close to these fixtures are usually accessible and thermal energy in the form of temperature difference between the surface of the pipe segment and the ambient environment can be transformed into electrical energy.

We have developed DoubleDip, a source-aware architecture and platform designed to enable non-intrusive water flow monitoring. DoubleDip senses vibrations during a water “event” much in the way that Kim et al. outline in NAWMS [14], except that DoubleDip leverages small thermal gradients that accompany water flow for both energy and wake up. The flow of water and thus the heat gradients from which DoubleDip harvests energy is incredibly sporadic, due to the nature of human interaction with both hot and cold water in both residential and office environments. The hot water input feeders to a kitchen or bathroom sink, for example, might only see a dozen uses over the course of a day or experience blackouts for a month at a time during user absences. The top half of Figure 1 shows the sparse temperature profile over a sample day for the surface of a hot water pipe leading to a faucet in a residential apartment. The profile follows a signature pattern of rapid transient when the faucet is turned on and slow settling after shut off as the water near the faucet loses its heat to the ambient environment. The total time that the faucet is on is almost imperceptible in this plot – 8 min over the 24 hour period.

1.1 Low Power Wakeup

Harvesting from unpredictable, bursty, and sporadic energy sources presents a host of challenges. Since the phe-

nomenon being sensed is also the one being harvested from, keeping the sensors active continuously to monitor a rare, ephemeral event is both impractical and wasteful [5]. One alternative to waking the system up from sleep for a water flow event is to use a form of low power listening (LPL) [26] with the vibration sensors already in place. A hypothetical vibration-based LPL mechanism can be imagined that activates an accelerometer ($\sim 300\mu\text{A}$ current) for a short period of time to estimate whether a water flow event has started. If not, the sensor is deactivated and the node goes back to sleep for the duty cycle period. A concern with this approach is that an event can only be detected with some latency, which could be the duty cycle period in the worst case. Assuming that our target average current consumption is $1\mu\text{A}$ and that it takes 250ms to detect the event reliably, the duty cycle period would have to be 75s to meet the $1\mu\text{A}$ target. Another issue is that of false wakeup. The bottom half of Figure 1 shows the vibration (standard deviation of accelerometer) readings over the same time duration as the top half, which shows many high amplitude spurious vibrations that would trigger even a finely tuned vibration-based LPL system. These vibrations originate from adjoining pipes in the plumbing network, the HVAC system, and general activity around the residence. While the NAWMS [14] system cancels out these “crosstalks,” the false wakeup is still a problem for power management.

To resolve these issues, DoubleDip employs a thermoelectric generator (TEG) in double duty. It uses thermal gradients not only for gathering energy but also for extremely low power ($< 1\mu\text{A}$) but sensitive wakeup. When a water flow event begins, there is a small but detectable voltage change across the TEG due to minute thermal gradients in the water stored in the pipe. Although this voltage is too small to harvest from, it is large and reliable enough to wake from. Figure 1 shows the triggers generated by DD’s wakeup circuitry compared to both the temperature profile and the vibration readings. While vibration has a number of excursions above a reasonable activity threshold even when temperature does not change, the DD triggers correctly identify water flow events (as evidenced by the temperature transient subsequent

to each trigger and a simultaneous spike in vibration). Additionally, DoubleDip wakes up almost instantaneously even while readings reported by temperature sensors are still below their noise floor, for example near the 21 hr mark in Figure 1. See Figure 13 and Section 3.6 for a detailed handling of wakeup performance and why temperature sensors cannot be used for wakeup directly.

In all, the key idea of the DoubleDip design is that it relies on an artifact of the phenomenon it is trying to capture to wake the system out of deep sleep through thermoelectric harvesting, removing the need for duty cycling, saving state, and preventing cold-boot delays [33].

1.2 Transient, Anemic Leaves

Interestingly, we find that there is sufficient temperature gradient in cold water pipes (coming from concealed plumbing) to wake a DD node attached to it. Occasionally, one may be fortunate enough to harvest some energy from cold water pipes when cooled water is fed to a heated ambient space or vice versa. In the worst case, when no additional energy is harvested, DD operates as a battery-powered sensor whose lifetime is dependent on the number of events sensed. Our experiments reveal that DD nodes on hot water pipes often have a net excess of harvested energy and, since hot and cold lines are often proximate, we implemented a *buddy charging* option by which more anemic nodes can request and receive energy from nodes with more plentiful energy reserves.

Finally, DoubleDip espouses a transmit-only radio topology as advocated by Schmid et al. in [28]. This is not just a design choice but an application necessity given the extreme scarcity of the energy reserve from which power is drawn. DoubleDip makes use of the state-of-the-art in harvesting, boosting, and radio technologies, but demands of the sensing application coupled with the scarcity of the energy reserve dictate that synchronization and wireless mesh routing be forfeited – without having to wake at a synchronized interval to listen for inbound packets we save on multiple fronts. We argue that this sacrifice is an inevitable one as such systems invade domains with increasingly harsh energy supplies and power requirements. These transmit-only *leaf* nodes need not be cut off from communication entirely; as described in [28], a simple scheme of queuing messages at *branch* nodes and sending them to a leaf destination will suffice, yielding a fully-networked *star* routing topology. Furthermore, the low rate and brevity of transmissions from DD or similar sensors means that packet collisions will be very rare.

1.3 Contributions of DoubleDip

The DoubleDip architecture is an attempt to ameliorate the effects of bursty sources on energy harvesting platforms. Towards this, we describe in detail the following contributions:

- We introduce the idea of thermoelectric wakeup for non-intrusive vibration-based water flow monitoring in residences and office buildings. We observe that water use even in busy areas like shared university restrooms is incredibly sporadic ($< 1\%$ total use), making continuous sensing both impractical and wasteful. Using minute changes in the thermal profile at pipe segments close to a faucet, DoubleDip is able to wake itself up

rapidly with high accuracy. The wakeup circuit employed in DD is simple and consumes less than $1\mu\text{A}$ of current. If no water flow event ever occurs, DD would last over 6 yrs in its deep sleep state.

- We design and implement DoubleDip. DD uses thermoelectric harvesting to recharge its low-leakage battery whenever an opportunity arises. DD further demonstrates an architecture for regulating and distributing harvested power, saving energy by voluntary and systematic power-gating of harvesting circuitry.
- We have evaluated DoubleDip in five locations over four weeks: two university restrooms and three residences. The university deployment spanned a break and the residences included one single occupant apartment and two multi-resident home. We find that hot water sensing frequently provides net excess energy, which can be shared with cold water DD nodes within proximity. Lifetime of DD nodes is sometimes limited only by the charge-discharge cycle lifetime of the Manganese Lithium batteries they use.

2 Related Work

The past decade has witnessed a sharp rise in the popularity of energy harvesting solutions for wireless sensor networks and power management in general. For small scale systems, harvesting solutions in the form of thermal, kinetic, solar, and piezoelectric transducers have found applications in new and exciting arenas. Improvements in low power digital and analog ICs such as the trail of low power motes marking the past decades [29, 21, 22, 6] as well as standardization of popular energy-aware embedded operating systems [18] have helped elucidate new low power computing paradigms and increase the lifetime of sensor networks, in some cases, into years.

The combination of today's energy harvesting techniques and the state-of-the-art in low power processing and wireless communication is indeed exciting, and there has been considerable research on this front in recent years. EnOcean has developed a range of commercially available products spanning solar and thermal harvesting and supporting various wireless protocols [7], and both Heliomote [19] and Yerva et al. [33] were successful in routing wireless packets despite frequently having to cold boot [33]. In contrast, the extreme scarcity of energy in water pipes dictates that DD nodes forgo routing altogether.

2.1 Thermoelectric Harvesting

Within the realm of thermoelectric energy generators (TEGs), much of the prior research can be divided into three categories – characterization and viability of TEGs for use in powering wireless sensor nodes, development of wearable body warmth energy harvesters for Body Area Networks (BANs), and application-specific systems for gleaning power from various waste sources.

Towards the characterization of TEGs for use in autonomous sensors, Ferrari et al. have measured the open circuit voltage and output power characterization for a range of devices, concluding by powering a low power mote with wireless transmitter when supplying the TEG with a temper-

ature gradient [8]. The careful measurement and design of the system presented in [8] lends credence to the efficacy of TEGs (and Peltier junctions) in powering sensor boards in addition to providing theoretical groundwork to compare expected output power from a given harvesting architecture and by which we are able to calculate several efficiency metrics herein.

There have been numerous attempts to create ‘wearable’ or ‘comfortable’ TEGs for body area networks (BAN), such as those described in [16, 12, 17] though in most cases the inclusion of a heatsink creates an unfortunate trade-off in power harvested and comfort. It must be noted that TEGs harvest energy from a temperature differential across two surfaces, but without an effective heatsink the cold surface slowly attains the temperature of the warm side by thermal conduction through the Peltier junctions reducing the power harvested despite a continuous heat source. Finding an appropriate heatsink is a challenge in DD as well—perhaps the most promising in this regard is a new flexible Peltier junction created by spraying doped silicon onto a flexible substrate [24] and air-cooled over a large area.

Many existing systems have the luxury of constant or pseudo-stationary energy sources – time constants may be in terms of hours or days. Zhang et al. demonstrated a “steam-powered” thermoelectric sensor network [34] using a Mica-2 [21] node and a Peltier device much like the one used in DD. In contrast, [34] has a fairly constant temperature gradient reaching as much as $80^{\circ}K$ and is able to generate over 0.5 W from an industrial steam pipe.

Despite the dissimilarities between DD and other platforms, DoubleDip faces many of the same obstacles encountered in other architectures and applications. In order to provide cold-boot power and effectively bootstrap their system, TwinStar uses a secondary solar panel [35]. Similarly, DD uses a low leakage lithium battery to avoid bootstrapping and race-condition issues. DoubleDip could be considered to be battery powered when opportunities for harvesting do not occur but recharges its battery when excess energy is available. Whereas TwinStar accepts leakage and designs control algorithms to counter it, DD reduces leakage by using a battery in lieu of a supercapacitor—a luxury afforded it given the low discharge depths of DD nodes on average and thus high cycle lifetimes of the rechargeable batteries used.

Recharging batteries from thermal harvesting has also been explored by Sodano et al. [30] for the application of Structural Health Monitoring (SHM). By directing solar irradiation to heat up one side of a TEG and using a large mechanical structure such as a bridge to cool the other side, [30] was able to charge a 300 mAh to capacity in a mere 3.5 min. An exhaustive summary of energy harvesting techniques for SHM including TEGs is presented by Park et al. in [23]. Therein, the authors cite interesting uses of thermal harvesting from such non-traditional sources as soil–air temperature gradients, vehicle exhaust pipe heat, and radioactive decay of isotopes in space (NASA) [23].

2.2 Water Flow Monitoring

Water flow monitoring and leakage detection continue to be active areas of research. Proposed techniques range from disaggregation of load monitors to distributed sensing archi-

tectures and even pipe topology mapping [32], each having inherent advantages and disadvantages. Non-intrusive disaggregation approaches such as Hydrosense [10, 2] use signature analysis approaches to determine the temporal profile of individual device usage as well as water flow rates. This approach offers high accuracy for detecting the on/off signatures for a number of devices after a manual training phase. While such systems have demonstrated success for controlled experiments, Froehlich et al. note that it is unclear how well these techniques scale when many devices overlap in time and in the presence of variable flow loads such as sinks [10]. On the other end of the spectrum, distributed approaches seek to give fixture-level accuracy without requiring a burdensome installation. TriopusNet [15], Fogarty et al. [9], and NAWMS [14] fall in this latter camp. In TriopusNet, sensors are deposited inside the feeder to a pipe network (residence, office, etc.) and then routed to unique pipe segments through a systematic turning on and off of faucets in order to direct nodes to appropriate branches for initial deployment and any maintenance thereafter. While this could potentially be automated, the infrastructure for doing so does not yet exist. Fogarty et al. attempt to map auditory signatures to water flow using distributed microphones, though they face problems with ambient noise and cross-talk.

NAWMS makes inferences about fixture-level water flow from individual pipe vibration data, self-calibrated using water flow information from a single feeder meter. The resulting system is non-intrusive and highly accurate, providing fine-grained flow estimates in environments where in-line sensors, manual calibration, and load disaggregation are prohibitively difficult.

3 System Architecture

The architectural design of DoubleDip is necessarily interdisciplinary – care must be taken not only to minimize electrical transients and steady state current but also to design thermal conduction paths to maximize temperature differentials across the TEG surfaces.

3.1 Thermoelectric Generators and the Seebeck Effect

As is the case with other thermoelectric generators, DoubleDip relies on the Seebeck effect – the corollary to the Peltier effect responsible for converting electricity into temperature differentials. DoubleDip uses a Peltier junction, an array of n-p junctions, to convert the difference in temperature across the array into an electrical signal, as shown in Figure 2 and given by:

$$V_{TEG} \approx \int_{T_{hot}}^{T_{cold}} (S_{cold} - S_{hot}) dT = (S_{cold} - S_{hot}) \Delta T \quad (1)$$

where S_{hot} and S_{cold} are the Seebeck coefficients of the materials used in the Peltier junction. This is a slight approximation, because in reality the Seebeck coefficients are non-linear functions of the temperature on an absolute scale.

V_{TEG} is therefore a function of the temperature gradient across the Peltier device, and so the use of a proper heatsink is of utmost importance whether producing temperature via the Peltier effect or harvesting it via the Seebeck effect. This

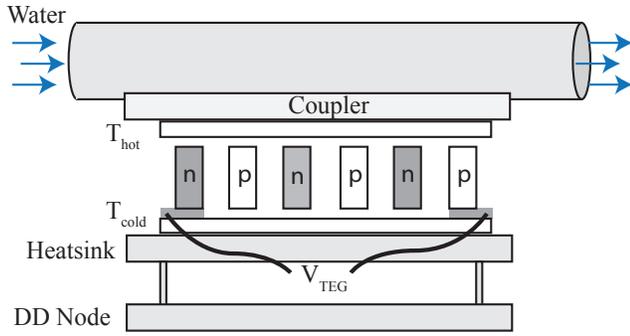


Figure 2: Diagram of DD node attached to water pipe, showing coupler, Peltier, heatsink, and the node itself.

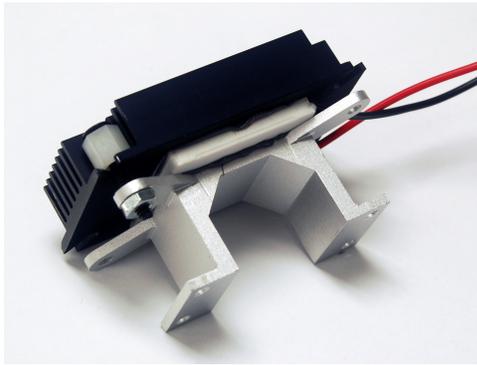


Figure 3: Peltier module including aluminum thermal coupler and heatsink. Four arms extend to fasten a heatsink to the Peltier module and coupler. The 0.75”-wide octagonal portion is covered in thermal compound prior to attachment to a right angle valve.

further complicates matters when trying to analyze the lifetime of such a system, as transients in temperature in the pipes will produce a power dependent on the temperature of the heatsink at any given time.

3.2 Mechanical Design

As shown in Figure 2, the TEG is attached to a pipe using a thermal coupler. In order to maximize heat transfer from the pipe to the aluminum coupler and Peltier module, the coupler-pipe contact must be of maximum surface area. This implies that the coupler itself must be designed to tightly fit a particular pipe. Any variation in this fit will result in variable coupler efficiencies as noted in Section 4. Fortunately, of the pipes surveyed in this study, each had a fairly standard right angle valve connecting the feeder pipe to a flexible hose pipe leading to a faucet. Coupling to this right angle valve requires a custom milled aluminum piece, and deployments on other locations such as shower heads would require a different design. Each DD node includes an aluminum thermal coupler designed to fit around these valves, and the resulting module with coupler, Peltier device, and heatsink is shown in Figure 3.

Of equal importance to the performance of DD is the heatsink used to bring the external side of the Peltier junction as close to ambient temperature as possible, and thus

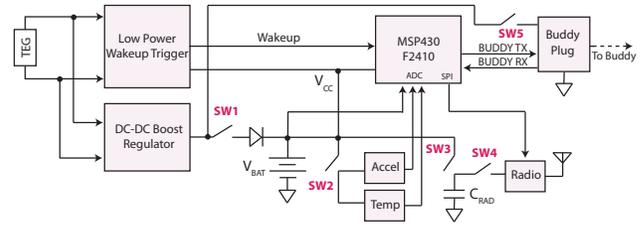


Figure 4: DoubleDip circuitry, including low power wakeup, RF, accelerometer, buddy plug for energy transfer, temperature sensors, and analog switches for power-gating and duty cycling.

increase ΔT . DD uses commercially-available air-cooled fin-style heatsinks secured to the external side of the Peltier device with thermal compound.

3.3 Electrical Design

Each DD node makes use of the state-of-the-art in both processor design and RF communication; the processor used is the MSP430F2410, boasting 300 nA low power mode with clock (LPM3) and 100 nA low power mode with RAM retention but requiring an external interrupt to wakeup. We have designed the node around the Nordic nRF24L01+ radio, the barebone radio used in recent Nordic ANT modules. The nRF24L01+ has a 900 nA deep power down mode and a full power (0 dBm) transmit current of 11.3 mA. In addition to processor and RF, the DD design consists of a power management section, low power wakeup, and a *buddy plug* for ‘social harvesting.’ The interaction of each of these subunits can be seen in Figure 4, and is described in detail below.

3.4 Power Management

Both the scarcity and amplitude of the energy events from which DD harvests dictate that the power regulation circuitry operate from very low start-up voltages and with very little quiescent current. Linear Technology offers a line of step-up DC converters and power management ICs for energy harvesting applications. Of these, the LTC3109 offers the ability to harvest from both voltage polarities (and consequently both temperature polarities, $\pm\Delta T$), and it has an input impedance similar to that found in most Peltier devices ($\sim 2 - 10\Omega$). This IC has one main output for powering a processor or other device as well as a storage output for placing excess energy harvested in a capacitor or battery. Though this is a convenient architecture, the complexity of the design has implications in terms of quiescent current draw. When not harvesting, the LTC3109 will draw 7 μA from the storage if available and from the output otherwise. To combat this, we use the IC on an as-needed basis, isolating it from the processor and sensing circuit completely when there is no energy to harvest. Figure 4 shows how this is implemented—switch *SW1* opens when the measured ΔT is below a certain threshold, indicating there is no energy to harvest. Additionally, this reduces the LTC3109 from an application-specific energy harvesting device to a load-matched step-up regulator. That is, the DD architecture is generalizable to any such step-up regulator or power management IC.

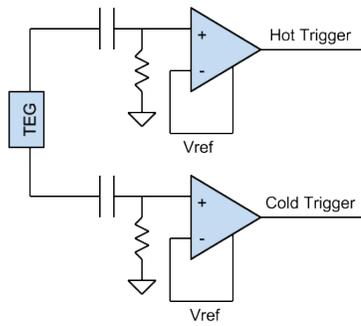


Figure 5: Comparator-based low power wakeup trigger. *Hot Trigger* or *Cold Trigger* will activate when there is any small deviation in ΔT , indicating water flow.

3.5 Battery vs. Capacitor

DoubleDip uses a battery-only design when it comes to energy storage. Despite technological advances in capacitors as storage devices, the leakage and energy density of supercaps cannot yet compete with current battery technologies. For example, the battery used in DD is an ML2020—a rechargeable 45 mAh manganese lithium battery with a volume of 63 mm³. If we were to draw, for example, 1 μA on average from this battery, the system lifetime would be 5.13 years from full charge. The equivalent capacitor to yield this lifetime would be $C = \frac{I_{\text{avg}} T}{\Delta V} \approx 60\text{F}$ with a volume of $\sim 1.2\text{ cm}^3$ [3], assuming an operating voltage of 2.5 V (the linear regime for the ML2020). A capacitor of this size would have much worse self leakage when compared to the equivalent battery—the same 60 F capacitor has a nominal leakage current of 47 μA [3], compared to the virtually negligible (1% per year, $< 51\text{nA}$) self-discharge of the lithium battery. Because of the uncertainty in the energy events and the possibility for extended periods of energy ‘droughts,’ we require an energy reserve capable of lasting for great lengths of time. The choice of using a rechargeable battery does not come for free; lithium batteries like the ML2020 have very low output current and suffer from a limited number of charge/discharge cycles. Because of the former, we must charge a capacitor to handle any transient loads like those demanded by radio transmissions. This is shown in Figures 4 and 6, where SW3 closes to charge C_{RAD} slowly from the battery and SW4 closes to cold-boot the radio. Because of the low quiescent current and aggressive duty cycling of DD nodes, the effective number of charge cycles of the battery (a decreasing function of the depth of discharge per cycle) is close to the maximum lifetime for that battery.

3.6 Low-Power Wakeup

An ideal water monitoring circuit is one that detects water flow immediately, wakes up to sense that flow, and then returns to sleep in an energy proportional manner. Detecting that a water flow event has occurred is a non-trivial challenge for external sensors, and the solution we have chosen is non-obvious, requiring due explanation.

Detection and trigger circuits can be divided into two classes—those that are actively polled and those that drive asynchronous interrupts. Those belonging to the former include such mechanisms as low power listen (LPL) accelerometers and temperature sensors, requiring a clock to periodically wake up and poll the sensors. In order to have a low latency trigger, we must perform this as frequently as possible. As mentioned in Section 1, duty cycling the accelerometer to reduce power will introduce an intolerable latency in the system, and accelerations caused by nearby pipes and other physical activity around the pipe will cause frequent false positives.

If we use a low power analog temperature sensor such as the MCP9700A, we can potentially sample at 1 Hz with a quiescent of 6 μA for only 1–2 ms. This gives a quiescent of $6\mu\text{A} \times 2\text{ms} + 0.3\mu\text{A} = 0.312\mu\text{A}$. At first glance, this seems like a very desirable wakeup method, but such low power temperature sensors come at a price in terms of resolution and accuracy, and thus false negatives will undoubtedly abound. We will see in subsequent sections that even more power hungry temperature sensors (with resolutions of around 0.125°K vs. 0.5°K for the MCP9700A) often times do not detect events that the wakeup trigger solution settled on in this work does until many seconds after the fact, potentially missing the event altogether. Those wakeup triggers belonging to the second camp—those that actively drive an interrupt pin—include battery monitoring circuits, temperature alert ICs, comparators, and the like. Of these, the battery monitoring circuits and temperature monitoring circuits are prohibitively power-hungry. Instead, DoubleDip employs low power comparators with finely tuned thresholds to meet its lifetime and accuracy goals.

The DD wakeup method, shown in Figure 5, combines a passive sensor and an active IC (ultra-low power comparators) to create an interrupt driving temperature monitor that consumes a mere 600 nA quiescent current. We leverage the fact that the very device from which we are harvesting energy creates small excursions in voltage (less than 10 mV in many cases) in the presence of a differential temperature. By passing this voltage through a low leakage (highly resistive) high pass filter, we can make use of the low power comparator ICs to send a wakeup trigger for both a positive and negative excursion from the baseline temperature. The cutoff frequency here (experimentally chosen to be $f_s = 10\text{s}$) is chosen to allow the baseline to readjust between events, allowing for detection of multiple subsequent events. Because this solution drives an active interrupt pin, we are able to wait in the MSP430’s lowest power setting, LPM4, without requiring power for a clock. This reduces the quiescent of the MSP430 from 300nA for LPM3 to 100nA for LPM4, giving a nominal wakeup current of 700 nA. Measured values indicate that this is closer to around 800nA. This means that, without any energy events to wake up and measure and with no additional harvested energy, DD will diligently wait for an event for 6.4 years given a 3V, 45 mAh battery.

3.7 Social Energy Harvesting

From our pilot deployment across 5 different locations, it is apparent that, given variations in time of day, weather, geographic location, plumbing, etc., some nodes will har-

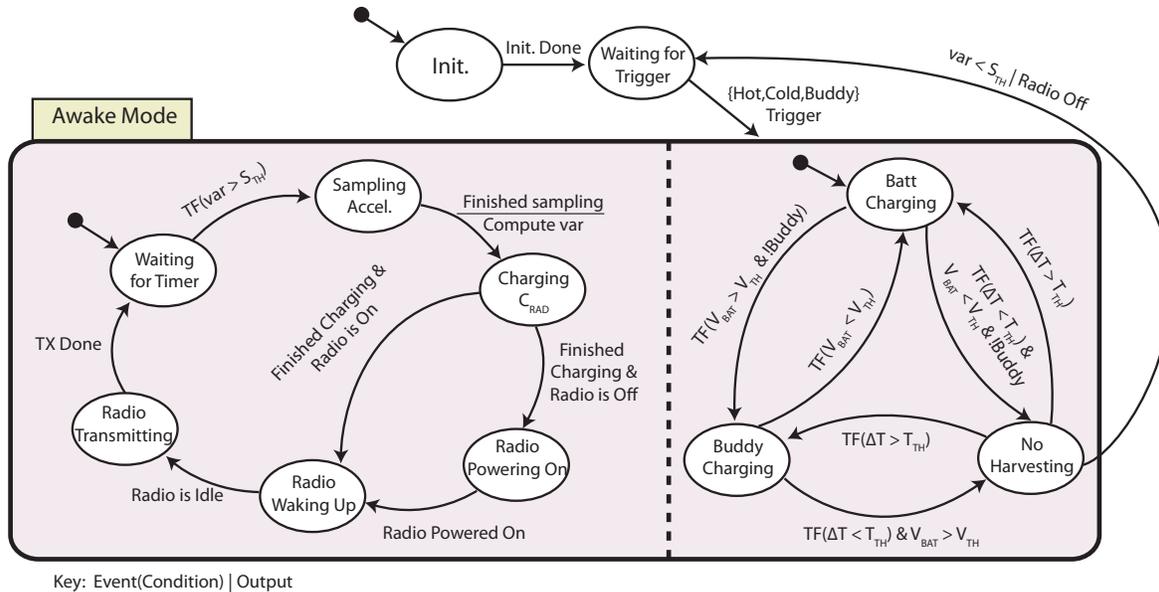


Figure 6: DoubleDip state chart, showing conditions for entering and exiting harvesting modes, buddy harvesting, sensing and relaying acceleration variance, and returning to low power wakeup. Events specified as TF are *Timer Fired* events. The dashed line within awake mode indicates concurrency.

vest more than others. Specifically, it is the case in those pipes tested in this pilot study that the hot water pipes experience a much larger ΔT than the cold water pipes due to the marginal difference between cold water and ambient temperatures. Obviously, cold water pipes in such an environment will harvest less than hot water pipes. In some situations the reduction in lifetime of these nodes is unavoidable and no nodes with surplus reserves are within reach, but in other cases it is possible and advantageous to transfer power harvested from DD nodes with excess energy to those suffering from more anemic sources. To accomplish this, we have added a four wire *buddy plug* for connecting one node to another. This four-wire interface (Power, Ground, TX, & RX) allows a node in surplus to signal to other nodes that it is about to send energy, after which it closes switch SW5 (Figure 4) to send energy over the *Power* line.

3.8 DD State Chart

Figure 6 illustrates the state transitions involved in the DD system flow in the form of a state chart. The diagram starts at the initialization state of the microcontroller, where GPIO pins, clocks, and communication registers are set. Upon completion of the initialization routines, DD immediately goes into the low power 800 nA state, waiting for an event to occur. Once an event has occurred (either a hot water trigger, a cold water trigger, or a signal from a nearby buddy node with an energy surplus), the system enters an awake state (LPM3) with a 1 Hz timer. The awake state has two independent logic loops, indicated by the dashed line.

In the left-most state loop, the MSP430 repeatedly samples the accelerometer and temperature (omitted from the state chart for simplicity) and calculates the accelerometer variance. It then proceeds to charge the transient load capacitor for the radio and transmit all relevant values over wire-

less to the base station. A given transmission period may contain several retransmits to mitigate packet loss due to the asynchrony of the network.

The right-most state loop determines the switch state of the power sector: the harvester will actively charge the battery as long as ΔT is above a certain threshold and as long as the battery itself is below a certain threshold. If we have enough energy, dictated by our specified battery threshold V_{TH} and $\Delta T > T_{TH}$ then we signal any neighbors available and pass energy to them. Similarly, if at any point we are being signalled by a buddy (BUDDY = 1 in Figure 6), we never transition from the *Batt Charging* state. If at any point we can gain energy from neither our own harvester ($\Delta T < T_{TH}$) nor a buddy harvester or our own battery reserves are below V_{TH} , we transition into the *No Harvesting* state. If we are in the no harvesting state and the variance of the accelerometer is adequately low ($var < S_{TH}$), indicating that there is no longer water flow to sense, we transition back into LPM4 and the low power wakeup mode.

4 Evaluation

The following section discusses results from hardware implementations of the proposed architecture as well as energy, accuracy, and lifetime statistics from the described deployment.

4.1 Energy Harvesting Characterization

In order to better understand the relationship between differences in temperature and energy harvested, we first characterized both the Peltier device and the step-up converter used. To characterize the Peltier device, a known temperature differential was applied across its two ceramic plates and the resulting open-circuit output voltage (V_{OC}) was measured. As described by Ferrari et al. in [8], V_{OC} is linearly

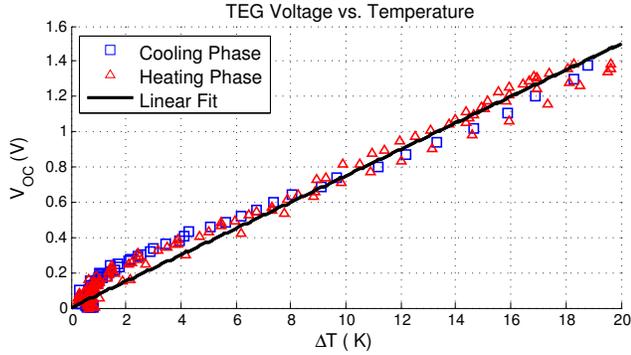


Figure 7: Open circuit output voltage for a CUI60333 Peltier module vs. differential temperature across the ceramic plates. The relationship is approximately linear, as per the physics governing the effect. For the heating and cooling phase, $r^2 = 98.9$ and 95.8 , respectively.

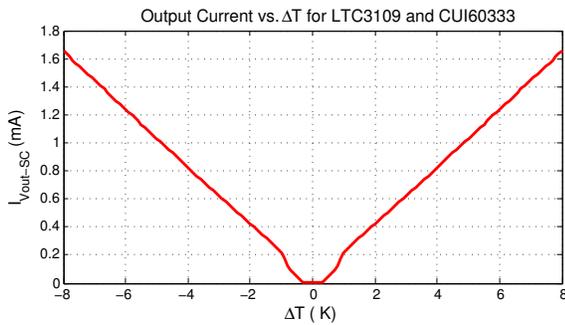


Figure 8: Short circuit output current of boost regulator vs. input temperature differential, ΔT .

dependent on ΔT , the difference in Seebeck coefficients of the materials involved, and the thermal conductivity across the ceramic plates. This linear relationship for a 3×3 cm Peltier device, the CUI60333, is shown in Figure 7.

Similarly, if we apply a known input voltage to the step-up converter such that $V_{in} = V_{OC}$, we can determine the short circuit current output of the power management sector and thus a complete mapping from ΔT to I_{SC} . Figure 8 shows this mapping, where temperature differences as little as $\pm 1^\circ K$ effect a modest I_{SC} . From this, given a known load resistance, we can calculate the expected output power based on ΔT . In the following section, we use this model and the actual energy collected to compute the efficiency of the thermal couplers used.

4.2 Fine-Grained Test Data Collection

In order to verify the DD architecture, we designed a separate testing board for fine-grained data collection with identical harvesting and triggering circuitry but with a different processor (for convenience) and with the possibility of running on both battery and USB power. This board, shown in Figure 9, samples pipe temperature, coupler temperature, heatsink temperature, accelerometer variance, output voltage, and system status (showing wakeup triggers, etc.) with one second granularity. The output of the harvesting circuitry on the testing board feeds into a large (15 mF) ca-



Figure 9: Testing PCB for 24/7 fine-grained monitoring of temperatures, accelerations, and energy harvested from both hot and cold water pipes.



Figure 10: DoubleDip node—ultra low power energy harvesting node with low latency wakeup, pipe vibration sensing, temperature sensors, and RF. The board size is roughly 3×3 in.

pacitor from which we can infer both the amount of energy harvested and the amount that would have been lost due to quiescent current draw of the boost regulator, had we not switched it off. If the capacitor on this board goes above a threshold voltage (2.5 V in these experiments), the capacitor is drained and the system status indicates that there was a discharge event. This keeps the capacitor from saturating to prevent it from rejecting any available harvested energy. We deployed this board on 6 distinct pipes feeding into 5 different sinks in 4 buildings (3 residential and 1 at a university). Of these pipes, 4 were hot water intake pipes to sinks, and 2 were cold water intakes. The statistics of this deployment are summarized in Table 1. Data was collected over a total of 103.5 node-days over which a total of 1,356 water flow events and 12.6 hours of induced vibrations were recorded.

Figure 11 shows data from a test board deployed for an example weekday during the month of March. The top plot shows temperature curves for the water pipe as well as the difference in temperature between the pipe and heatsink while the second and third plots show output voltage and cumulative energy harvested, respectively. Energy harvested during this day totaled around 1000 mJ, while energy expended is equal to energy spent waiting in low power wakeup mode plus the energy spent sensing and transmitting water

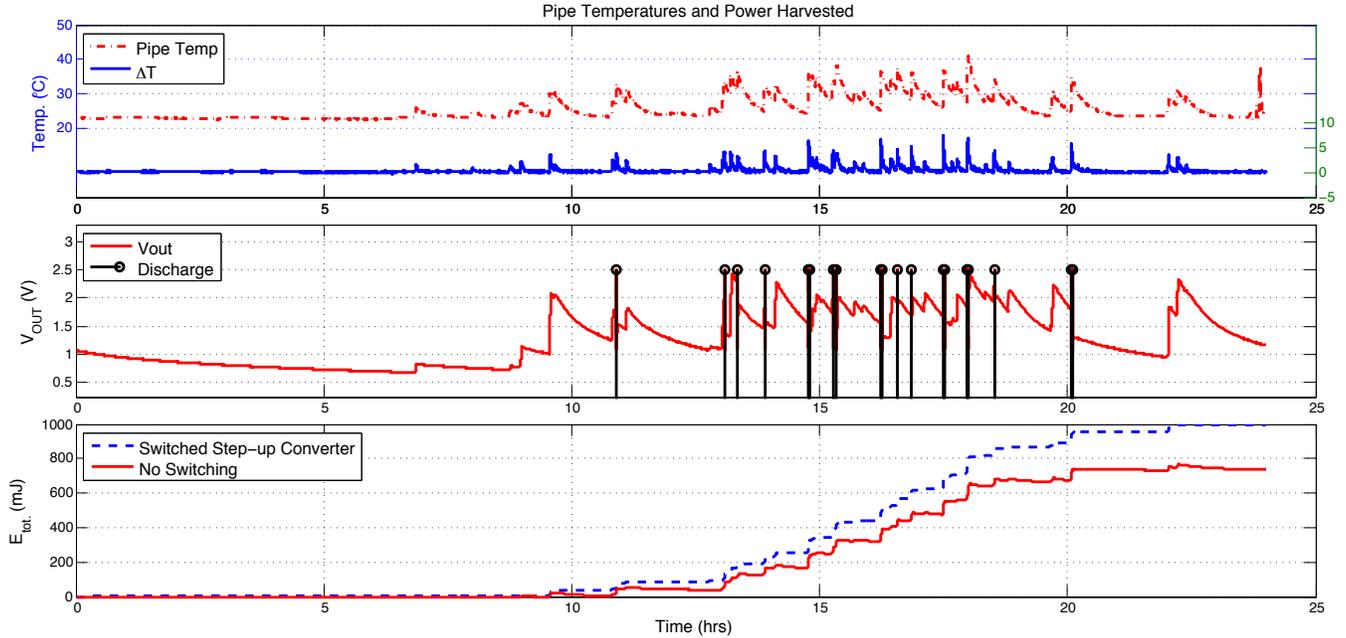


Figure 11: Power harvested for an example weekday during March (24 hours starting at 12:00am). From top to bottom, plots show (1) Pipe temperature and $\Delta T = T_{coupler} - T_{heatsink}$; (2) output voltage across a 15 mF capacitor with discharge signals when the capacitor gets close to saturation; and (3) cumulative energy harvested over the entire day, in mJ. Cumulative energy is shown when the step-up DC-DC converter is always on and when it is switched on an as-needed basis.

	Node Time (hrs)	On Time (hrs)	Events
Univ. 1 (Hot)	295.6	2.4	389
Univ. 1 (Cold)	289.2	1.9	215
Univ. 2 (Hot)	239.2	2.3	249
Res. 1 (Hot)	99.2	0.5	29
Res. 2 (Hot)	1127.1	2.7	291
Res. 3 (Cold)	435	2.8	183

Table 1: Total deployment statistics for 6 nodes across 4 buildings, showing the total amount of time each node was installed (Node Time) and active during an event (On Time) as well as the total number of events sensed.

event data. This number will be explained in detail in the following sections.

4.3 Wakeup Power, Latency, and Accuracy

The low power wakeup mode draws quiescent current from two comparators operating at a nominal 300 nA and the MSP430F2410 operating at a nominal 100 nA, making the nominal current 700 nA. As Figure 12 shows, this current is closer to around 800 nA, with an average of 820 nA. Figure 12 also shows that if the pins are configured to communicate with the radio, the average current goes up to around 900 nA, and if we were to use the MSP430’s LPM3 (with low power oscillator), the average current would be closer to 1.2 μ A. The comparators used (the LTC1540) have a high offset voltage, meaning that for best performance a comparator’s reference voltage (also generated by the LTC1540) should be tuned to just above the triggering threshold. Even with

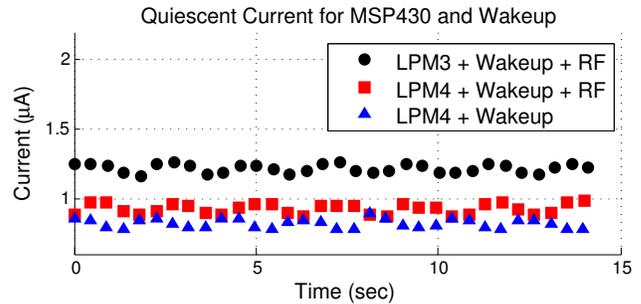


Figure 12: Current consumption for low power wakeup circuitry and deep sleep MSP430 (LPM4). Current is shown with and without SPI pins configured for RF and compared to a sleep mode with low power clock (LPM3).

all comparators tuned, DD nodes in different physical locations and buildings are likely to see a range of wakeup times, due in large part to the pipe’s distance from the main water heater, the ambient temperature of the building, the fitting of the coupler, and even the frequency with which water flows through that pipe—these factors and others contribute to the rate of heat transfer from pipe to TEG. Figure 13 explains the procedure for estimating the wakeup latency from the vibration data and the triggers from the wakeup cycle. Figure 14 shows a CDF of the wakeup latencies across all nodes, and Table 2 shows the number of false positives and the average amount of water flow time missed per day. There occur many more false negatives than false positives (none of which we observed), but by relating the false negatives to the

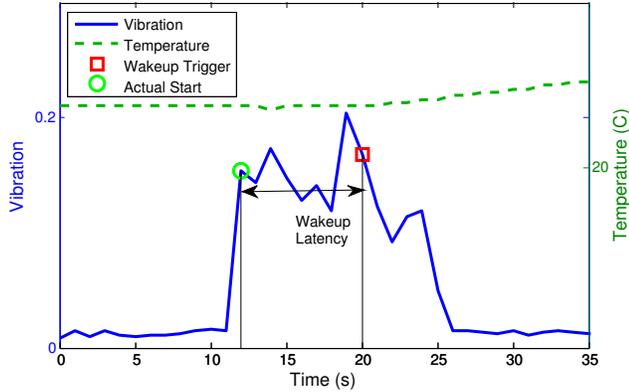


Figure 13: Example calculation of wakeup latency. The green circle indicates the start of the water event, as determined by the vibration variance (y-axis). The red square indicates the time the wakeup trigger fired. The difference between these two is the latency.

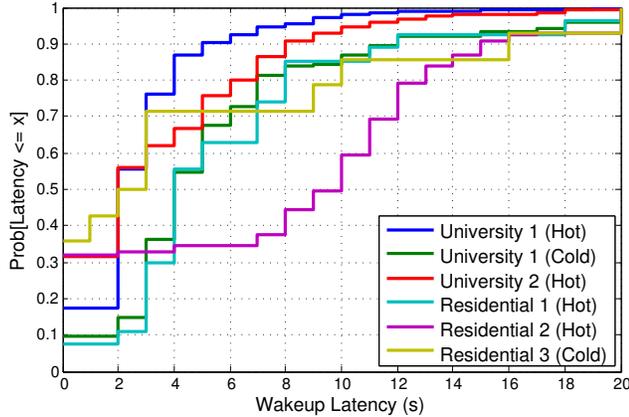


Figure 14: Cumulative distribution function (CDF) for wakeup latencies across all nodes.

corresponding length of the water event missed, we see that the actual percentage of water flow left unobserved is fairly small ($< 2\%$) in a university setting and moderate in residential settings. Where the pipes are farther from the main boiler or the water pressure is lower (such as residential housing) the latency observed and % water flow time unobserved is considerably higher. However, this can be ameliorated to some extent through careful tuning. For example, the excellent latency for all University deployments is in part due to better tuning given their ease of access. Residential deployments, on the other hand, received no further tuning after the initial deployment. Figure 15 shows how drastically wakeup latency can be improved by fine-tuning the comparator threshold. Furthermore, this tuning can be completely automated. We have created a prototype autotuning system using a digital potentiometer, yielding much finer resolution tuning and thus lower latency wakeup. This comes at only a marginal cost in energy—about 50 nA additional—and will thus be incorporated into future versions.

	%Extra Events	%Missed Events	%Missed Time
Univ. 1 (Hot)	0.00	3.62	1.54
Univ. 1 (Cold)	0.00	2.40	0.73
Univ. 2 (Hot)	0.00	1.61	1.79
Res. 1 (Hot)	0.00	13.79	20.31
Res. 2 (Hot)	0.00	34.60	19.32
Res. 3 (Cold)	0.00	15.38	17.59

Table 2: Wakeup latencies and accuracies per day, per node.

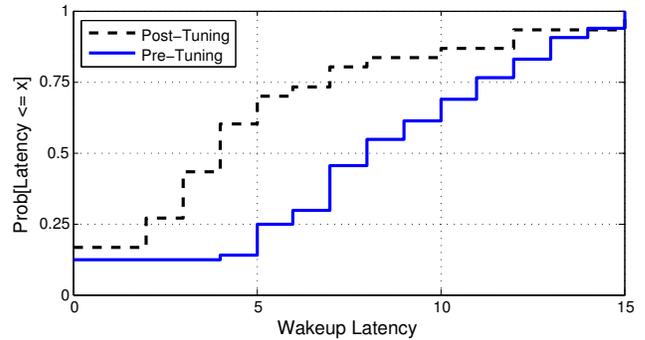


Figure 15: CDF of wakeup latency for a single node before and after tuning the wakeup threshold.

4.4 Subsystem Power Characterization

Since the quiescent current of DD is relatively minuscule and the voltage versus capacity curve of the battery is fairly flat, we see little or no change in battery voltage as energy is consumed or returned to the battery. To get a good estimate of the energy harvested, therefore, the test board shown in Figure 9 is used. In order to determine the overall system lifetime and the effect that harvesting has on the lifetime, however, we must perform a careful analysis of the DD node itself, shown in Figure 10. Specifically, we characterized the operation of the Nordic nRF24L01+ 2.4 GHz radio while operating from a capacitor for transient loads and the power consumption of the MSP430F2410 when transitioning from sleep to active mode and sampling an accelerometer at 100 Hz. By doing so, we can calculate the amount of energy spent per day by knowing the amount of time spent sampling events and transmitting RF. The total amount of energy consumed per day is then given by:

$$E_{day} = I_{idle} * t_{idle} + I_{adxl} * t_{on} + I_{RF-boot} * N + I_{TX}(t_{on} - N) \quad (2)$$

Where N is the number of events detected in that day, t_{on} is the total amount of system on-time during that day, $I_{RF-boot}$ is the average cold boot wakeup current of the nRF24L01+ radio after power-gating the radio, and I_{TX} refers to the average transmission cost of the radio when one packet is sent every second.

Figure 16 shows the instantaneous and average currents consumed by both the accelerometer and the MSP430 while sampling the accelerometer for 32 samples. If we naively sample the accelerometer, an ADXL335, by power-gating only at the very beginning and end of the sampling session, the system consumes 296 μ A on average. If however

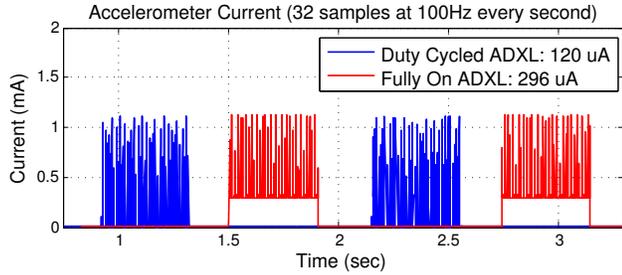


Figure 16: Current consumption of an ADXL335 3-axis analog accelerometer, sampling 32 bytes every second at 100 Hz. Current for a constantly on ADXL335 vs. one that is power-gated ('Duty Cycled') are shown, offset in phase for clarity. Power-gating reduces average current consumption from 296 μA to 120 μA .

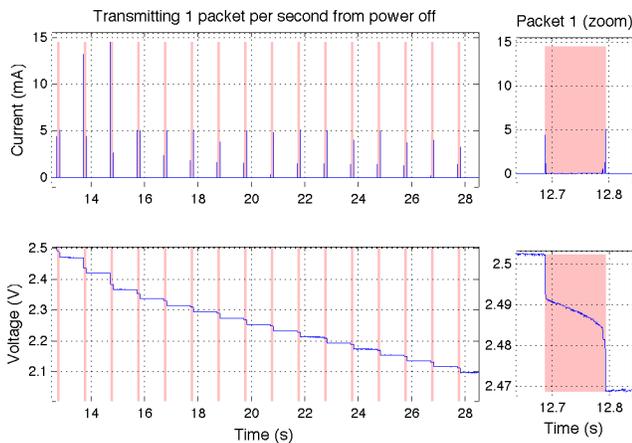


Figure 17: Current consumption of a Nordic nRF24L01+ radio operating from a capacitor. The radio sends one 32 Byte packet every second and then is turned off (power-gated). The right-most plot shows a zoomed in version of what occurs to the voltage across the transient capacitor during a single transmission.

we power-gate the accelerometer in between samples (something that is only possible if the decoupling capacitance is kept below 0.1 μF), we obtain an average power of 120 μA .

Figure 17 shows the operation of the radio from a transient load capacitor. A modest reduction in capacitor voltage occurs during each transmission, and the capacitor must be recharged in a periodic fashion to send multiple packets or retransmissions. The current consumed by the radio as measured by an ammeter and verified from Figure 17 is 4.96 μA with $\sigma = 0.99\mu\text{A}$ for power-gating and 3.43 μA with $\sigma = 0.37\mu\text{A}$ for power-down (RAM retention) for one 32 Byte packet per second. By moving away from LPL radio reception and heavily duty cycling stackless radios, we can greatly reduce the consumption of what was once and in some systems remains the power bottleneck [4, 13].

4.5 Net Energy Profile

Combining these numbers as per Equation 2 and comparing with the total amount of energy harvested per day given the test boards yields the net harvested per day, per node.

Figure 18 shows these numbers plotted per day-of-the-week (left) and the net energy per day taken by subtracting energy consumed from energy harvested (right) along with the average net energy per day. As expected, the most energy gained per day is realized by the two hot water intake pipes in the university restroom. Additionally, these two numbers are heavily influenced by both the day of the week (weekend vs. weekday) and the season – the weekdays leading up to the second week are during a university break. Residential 1 (a single residence), for which only 6 days of data were collected, is also marginally in positive numbers, despite a low number of total events. This is due to the energy proportionality of DD, scaling down to just 820 nA when no activities are present. University 1 Cold and Residential 2 and 3, however, are in negative numbers, indicating that their lifetime is limited by more than just the shelf life and cycle count of the battery chosen. The DD system lifetime can then be estimated by:

$$t_{life} = \begin{cases} t_{bat-shelf-life} & \text{if } E_{avg} \geq 0 \\ \frac{A \cdot h(bat) \times V_{cc} \times 3600}{|E_{avg}| \times 365.25} & \text{if } E_{avg} < 0 \end{cases} \quad (3)$$

where the capacity of the rechargeable lithium battery used is 45mAh, and $V_{cc} = 3\text{V}$ during normal operation. The energy burned per day for a node where no events occur and no energy is harvested is 228mJ and the total lifetime is 6.4 years. On the other hand, the worst case E_{avg} per day from Figure 18 is -128 mJ / day, belonging to Residential 3 cold. This has a total lifetime of $45\text{mAh} \times 3 \times 3600 / 128\text{mJ} \times 365.25 = 10.4$ years, almost 10 years shy of the shelf life of the ML2020 battery used. This energy can be made up for if we request energy from a nearby node in order to meet the full shelf life specification of the battery, though in some cases nodes with energy deficits may be isolated.

Finally, it is also particularly interesting to note unexpected trends (or even lack thereof) in terms of power harvested by week-day. There are many confounding variables responsible for some of the change in energy harvested during the first week (vacation) and over the weekends, not the least of which is the status of HVAC in the building. For example, if air conditioning is turned off to save power, the ambient temperature of the cold water close to the valve might rise while the water in walled-in pipes remains cool, causing an increased ΔT when water begins to flow.

4.6 Efficiency

There are many conversions from the path of temperature to usable energy, and it is worth noting the various inefficiencies at each junction and the effect each has on the system as a whole. Perhaps the most important of these is the pipe→coupler+heatsink junction. The inefficiencies here are very observable; an ideal heatsink would remain at ambient temperature while an ideal coupler would track the pipe temperature with zero error. Instead, the heatsink converges to within some ΔT of the coupler given time and the coupler resembles a low-pass filtered version of the pipe temperature. Figure 19 shows an example of the temperature curves observed during two bursts of hot water on University 1 Hot.

By combining the temperature measurements with the model derived in Figure 8, we can calculate the

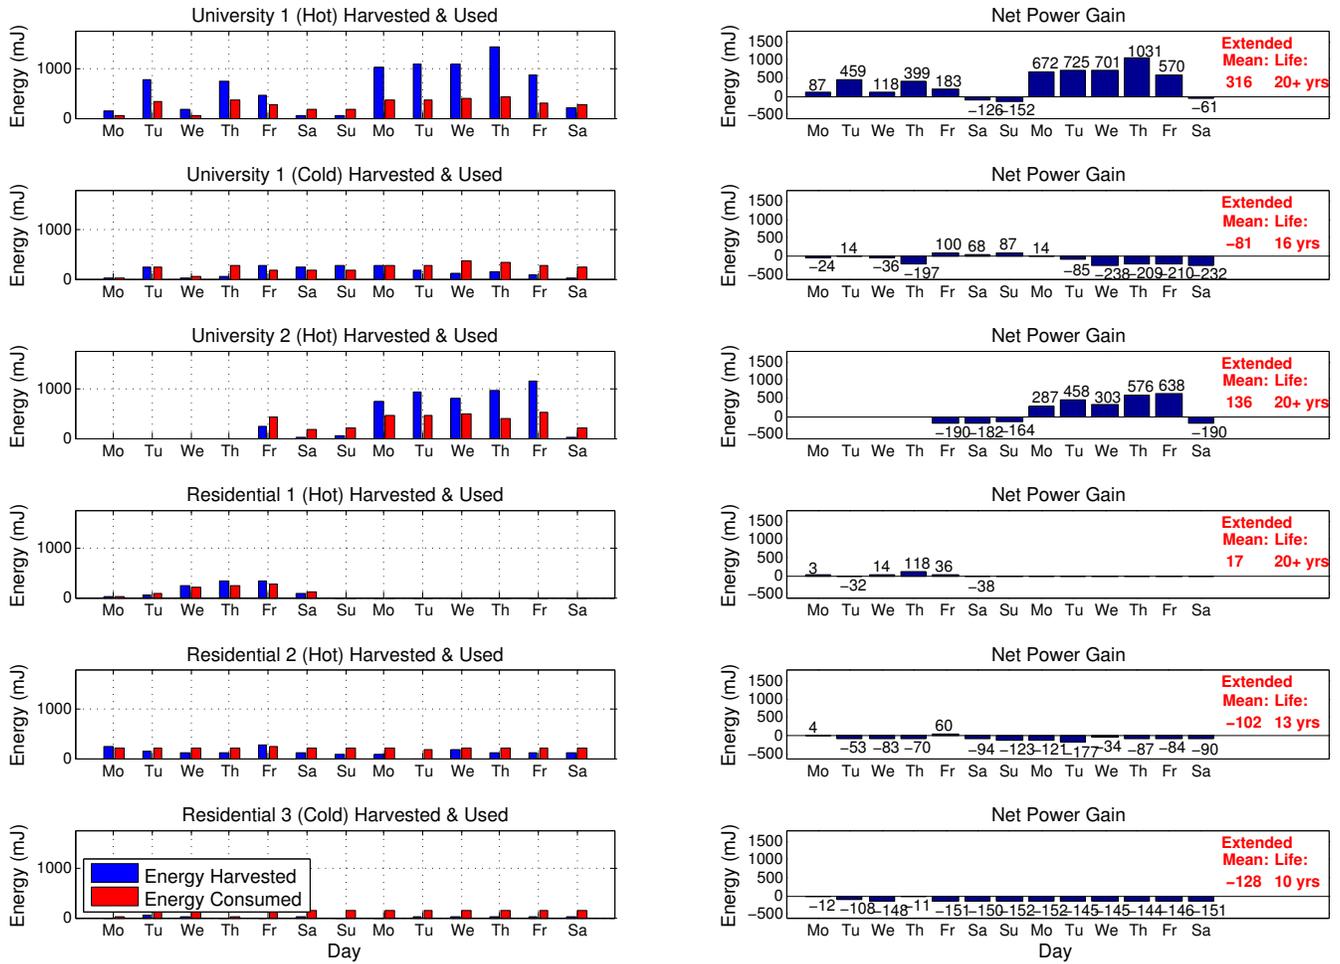


Figure 18: Energy harvested and consumed per day for 2 example weeks. By subtracting energy consumed from energy harvested (left) we get the net energy per day (right). Mean daily energy and lifetime estimates are calculated using the entire data set for each node (not shown). University 1 (Hot), University 2 (Hot), and Residential 1(Hot) show a net positive in energy gained, indicating that the node lifetime is limited only by the battery shelf life. The others (while suffering a net loss in energy) are still considerably better than they would have been had the same circuitry been used without harvesting.

Location	Coupler Efficiency (%)
University 1 (Hot)	65.82
University 1 (Cold)	19.55
University 2 (Hot)	41.59
Residential 1 (Hot)	27.09
Residential 2 (Hot)	17.96
Residential 3 (Cold)	21.65

Table 3: Coupler efficiencies

pipe→coupler efficiency as the ratio between power delivered to the TEG and ideal power given the pipe and ambient temperatures. These values are given in Table 3, ranging from 65% all the way down to 18%. The variation can be explained in part by the installations—each installation requires fastening the coupler to the pipe, and some installations were more successful than others due to manufacture variations in the right angle valves and couplers as well as variations in how tightly each coupler was fastened to the

right angle valve. Improving this efficiency can be accomplished by increasing the surface area in contact with the coupler, decreasing the mass and radiant surface area of the coupler, and improving the fastening process. The calculation of the efficiency of converting the thermal energy to electrical energy, on the other hand, is quite complex and contains many opposing forces. In order to capture more heat from the pipe, more surface area must be in contact with the coupler as stated above. Yet in increasing the surface area in contact one necessarily increases the mass of the coupler itself, increasing the rise time of the coupler and decreasing the responsiveness of the wakeup triggers. Measures can be taken to decrease the thermal mass of the coupler while retaining a larger surface area in contact with the pipe, but in doing so care should be taken to not increase the radiant surface area of the coupler, thereby creating a virtual heatsink on the side on you wish to retain heat. Figure 20 shows how the current thermal coupler responds to changes in contact

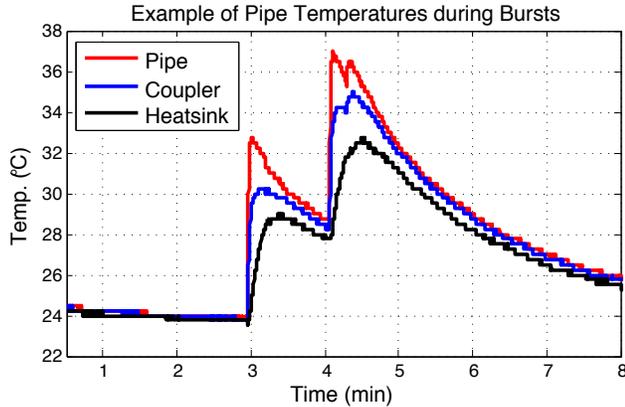


Figure 19: Example temperature curves for hot water bursts, showing the difference between T_{Pipe} and $T_{Coupler}$ as well as the non-ideality of the heatsink.

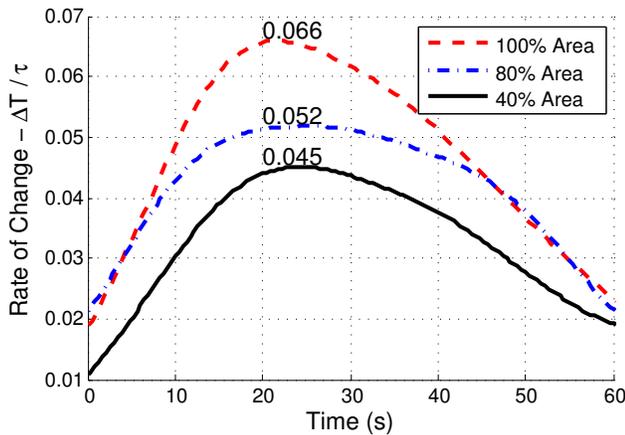


Figure 20: Temperature change rate for varying contact areas.

area. Decreasing contact area clearly decreases the rate at which heat is transferred to the TEG, which will decrease energy harvested and increase wakeup latency.

The conversion of the small electrical signal (V_{OC}) produced by the Peltier (and indeed the Peltier junction itself) to a usable voltage is another significant source of inefficiency. The specified efficiency of the step-up converter used (and as validated by tests described earlier) is at most 25% and perhaps better estimated at around 15% given the range of input voltages seen. This efficiency increases dramatically if we can afford to increase the operational start-up voltage. That is, if we can increase V_{OC} by means described above, we can not only harvest more but lose less in the boost conversion process. Finally, there is a loss in the diode used to drop the output voltage of the step-up converter from 3.3 V to 3.1 V—the nominal charging voltage of the lithium battery—and in the charging / discharging routine of the battery itself. The efficiency of the diode can be calculated as $100\% \times \frac{3.1V}{3.3V} \approx 94\%$, and despite lengthy tests of charge & discharge cycles at varying currents, no appreciable inefficiency in the battery could be detected. Given these inefficiencies, the potential to harvest from energy embed-

ded in water pipes is clearly much greater than we are able to achieve with our current prototype, and improvements in technologies and methods mentioned throughout this paper will likely allow for a host of similar applications.

5 Future Work

Despite the success of DD in meeting the shelf life limits of the battery used for those nodes with a net positive energy per day and those in close proximity and able to share surplus energy reserves, there remain interesting challenges to both increasing the efficiency and quality of DD and applying these lessons elsewhere. Potential future work currently under consideration includes:

1. A finite element analysis of heat transfer and dissipation via the thermal coupler.
2. A more advanced two-way communication protocol between proximal harvesting nodes, allowing nearby nodes access to battery state of all neighboring nodes and ways to more elegantly meet lifetime specs.
3. A continued deployment across more residencies to better understand the variations of wakeup latencies, temperatures, and power numbers from location to location.
4. Developing a hybrid capacitor system to amortize the effect of battery shelf-life on overall system lifetime.

6 Conclusion

We have demonstrated an architecture—DoubleDip—capable of sensing pipe vibrations and enabling non-intrusive water flow detection for, in some cases, the entire lifetime of the lithium battery used (20+ years). We accomplished this by leveraging the energy embedded in the differential temperature between the pipe (hot or cold) and the ambient room temperature for both compensating energy expenditures and triggering the system to awaken from a deep sleep mode. DD espouses an extremely low power architecture, wherein the sacrifice of mesh networking, time synchronization, and other luxuries is a necessity given the anemia and irregularity of the energy (water) events.

Our pilot deployment of 6 nodes across 5 sinks and 4 buildings showed that our low power wakeup trigger had a reasonable latency (around 4 seconds after calibration) and a high accuracy of detection, with virtually no false positives and a small fraction of water flow time missed with respect to the total time of water flow per day. We observed that some pipes experience much greater changes in temperature and thus were able to harvest much more than others, even given physical proximity. To mitigate this, we developed a method by which a node with excess energy reserves can wake up a nearby node and transfer a fraction of its harvested energy. Net energy gain from day to day varied from an average of +316 mJ/day on a university hot water pipe feeding a sink faucet to -128 mJ/day on one residential cold water pipe, meaning the lifetime of some nodes was limited only by the battery shelf life, and even those nodes with net losses in energy saw an improvement through harvesting and the low power DD architecture.

Finally, all hardware schematics have been released as open source material and can be found at

<http://nesl.ee.ucla.edu/fw/dd>.

7 Acknowledgments

This material is based upon work supported by the NSF under awards # CNS-1143667 and OIA-0963183. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

8 References

- [1] C. Aguilar, D. White, and D. Ryan. Domestic water heating and water heater energy consumption in Canada. *Canadian Building Energy End-Use Data and Analysis Centre (CBEEDAC '05)*, April 2005.
- [2] T. Campbell, E. Larson, G. Cohn, J. Froehlich, R. Alcaide, and S. N. Patel. Watt: a method for self-powered wireless sensing of water activity in the home. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*, pages 169–172, New York, NY, USA, 2010.
- [3] Cooper industries supercaps. http://www.cooperindustries.com/content/dam/public/bussmann/electronics/resources/data_sheets/bus_elx_ds_4375_hb_series.pdf.
- [4] N. Dang, E. Bozorgzadeh, and N. Venkatasubramanian. Quares: Quality-aware data collection in energy harvesting sensor networks. In *The Second International Green Computing Conference (IGCC '11)*, pages 1–9, July 2011.
- [5] P. Dutta, M. Grimmer, A. Arora, S. Bibyk, and D. Culler. Design of a wireless sensor network platform for detecting rare, random, and ephemeral events. In *IEEE Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN '05)*, Piscataway, NJ, USA, 2005.
- [6] P. Dutta, J. Taneja, J. Jeong, X. Jiang, and D. Culler. A building block approach to sensor network systems. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys '08)*, pages 267–280, New York, NY, USA, 2008.
- [7] EnOcean. Thermo energy harvesting for energy harvesting wireless sensors. <https://www.enocean.com>. 2012.
- [8] M. Ferrari, V. Ferrari, M. Guizzetti, D. Marioli, and A. Taroni. Characterization of thermoelectric modules for powering autonomous sensors. In *IEEE Proceedings on Instrumentation and Measurement Technology (I2MTC '07)*, pages 1–6, May 2007.
- [9] J. Fogarty, C. Au, and S. E. Hudson. Sensing from the basement: a feasibility study of unobtrusive and low-cost home activity recognition. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*, pages 91–100, New York, NY, USA, 2006.
- [10] 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 ACM International Conference on Ubiquitous Computing (UbiComp '09)*, pages 235–244, New York, NY, USA, 2009.
- [11] E. Geller, J. Erickson, and B. Buttram. Attempts to promote residential water conservation with educational, behavioral and engineering strategies. *Population & Environment*, 6(2):96–112, 1983.
- [12] D. Hoang, Y. Tan, H. Chng, and S. Panda. Thermal energy harvesting from human warmth for wireless body area network in medical healthcare system. In *International Conference on Power Electronics and Drive Systems (PEDS '09)*, pages 1277–1282, Nov. 2009.
- [13] L. Huang, V. Pop, R. de Francisco, R. Vullers, G. Dolmans, H. de Groot, and K. Imamura. Ultra low power wireless and energy harvesting technologies; an ideal combination. In *IEEE International Conference on Communication Systems (ICCS '10)*, pages 295–300, Nov. 2010.
- [14] 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.
- [15] T. T.-T. Lai, W.-J. Chen, K.-H. Li, P. Huang, and H.-H. Chu. Triopusnet: automating wireless sensor network deployment and replacement in pipeline monitoring. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks (IPSN '12)*, pages 61–72, New York, NY, USA, 2012.
- [16] V. Leonov, B. Gyselinckx, C. V. Hoof, T. Torfs, R. F. Yazicioglu, R. J. M. Vullers, and P. Fiorini. Wearable self-powered wireless devices with thermoelectric energy scavengers. *Integration Issues of Miniaturized Systems*, pages 1–8, April 2008.
- [17] V. Leonov, T. Torfs, P. Fiorini, and C. Van Hoof. Thermoelectric converters of human warmth for self-powered wireless sensor nodes. *IEEE Journal on Sensors*, 7(5):650–657, May 2007.
- [18] P. Levis, S. Madden, J. Polastre, R. Szewczyk, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler. Tinyos: An operating system for sensor networks. In *Ambient Intelligence (Aml '04)*. Springer Verlag, 2004.
- [19] K. Lin, J. Yu, J. Hsu, S. Zahedi, D. Lee, J. Friedman, A. Kansal, V. Raghunathan, and M. Srivastava. Heliomote: enabling long-lived sensor networks through solar energy harvesting. In *Proceedings of the 3rd ACM International Conference on Embedded Networked Sensor Systems (SenSys '05)*, pages 309–309, New York, NY, USA, 2005.
- [20] P. W. Mayer, W. B. DeOreo, E. Towler, and D. M. Lewis. Residential indoor water conservation study: Evaluation of high efficiency indoor plumbing fixture retrofits in single-family homes in the east bay municipal utility district service area. *Aquacraft, Inc.*, July 2003.
- [21] Crossbow mica2 mote. http://bullseye.xbow.com:81/support/support_pdf_files/mpr-mib_series_users_manual.pdf. June 2007.
- [22] Crossbow micaz oem module. http://bullseye.xbow.com:81/support/support_pdf_files/mpr-mib_series_users_manual.pdf. June 2007.
- [23] G. Park, T. Rosing, M. Todd, C. Farrar, and W. Hodgkiss. Energy harvesting for structural health monitoring sensor networks. In *Journal of Infrastructure Systems*, volume 14, pages 64–79, March 2008.
- [24] Perpetua flexible thermoelectric generators. <http://www.perpetuapower.com/technology.htm>. 2012.
- [25] Pike research. <http://www.pikeresearch.com/research/smart-water-meters>. 2012.
- [26] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd ACM International Conference on Embedded Networked Sensor Systems (SenSys '04)*, pages 95–107, New York, NY, USA, 2004.
- [27] P. Raskin, E. Hansen, and R. Margolis. Water and sustainability. In *Natural Resources Forum*, volume 20. Wiley Online Library, 1996.
- [28] T. Schmid, R. S. Shea, M. B. Srivastava, and P. Dutta. Disentangling wireless sensing from mesh networking. In *The Sixth Workshop on Hot Topics in Embedded Networked Sensors (HotEMNETS '10)*, June 2010.
- [29] Sentilla tmote sky. <http://www.sentilla.com/pdf/eoltmote-sky-datasheet.pdf>.
- [30] H. Sodano, G. Simmers, R. Dereux, and D. Inman. Recharging batteries using energy harvested from thermal gradients. In *Journal of Intelligent Material Systems and Structures (JIMSS '07)*, volume 18. SAGE Publications, 2007.
- [31] V. Srinivasan, J. Stankovic, and K. Whitehouse. Watersense: Water flow disaggregation using motion sensors. In *The 3rd ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys '11)*, in conjunction with ACM SenSys, 2011.
- [32] T. te Lai, Y. han Chen, P. Huang, and H. hua Chu. Pipeprobe: A mobile sensor droplet for mapping hidden pipeline. In *The 8th ACM Conference on Embedded Networked Sensor Systems (SenSys '10)*, pages 113–126, Jun 2010.
- [33] L. Yerva, B. Campbell, A. Bansal, T. Schmid, and P. Dutta. Grafting energy-harvesting leaves onto the sensor tree. In *Proceedings of the 11th ACM International Conference on Information Processing in Sensor Networks (IPSN '12)*, pages 197–208, New York, NY, USA, 2012.
- [34] C. Zhang, A. Syed, Y. Cho, and J. Heidemann. Steam-powered sensing. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems (SenSys '11)*, pages 204–217, New York, NY, USA, 2011.
- [35] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z.-L. Zhang. Leakage-aware energy synchronization for wireless sensor networks. In *Proceedings of the 7th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys '09)*, pages 319–332, New York, NY, USA, 2009.