ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home

Sidhant Gupta

A dissertation
submitted in partial fulfillment of the
requirements for

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Shwetak N. Patel, Chair

Gaetano Borriello

Joshua R. Smith

Program Authorized to Offer Degree:
Computer Science & Engineering

University of Washington

Abstract

ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home

Sidhant Gupta

Chair of the Supervisory Committee:
Dr. Shwetak N. Patel
CSE

Imagine an energy feedback system that displays not only your total power consumption, but also continuously shows real-time usage while breaking it down categorically by electrical appliances. In addition, such a system provides personalized and cost-effective energy saving recommendations, for example, it could report, "Based on your usage patterns, you could save $215 per year by switching to a more efficient heating unit, which will pay for itself in 27 months". The challenge in this scenario is to sense end-uses of energy to provide feedback at the appliance level. Though an individual measurement device could be installed on each and every appliance to obtain such fine-grained energy usage patterns, such an approach is not only expensive but requires tedious maintenance and installation.

In this thesis I formalize and evaluate a new sensing techniques that requires minimal instrumentation in the form of a single plug-in sensor for measuring disaggregated electrical energy usage in the home. In particular, this sensor leverages Electromagnetic Interference (EMI) produced by most consumer electronics as signatures to detect and classify their use in real-time. In addition, I highlight the complementary and exclusive machine learning features can be extracted from this EMI which can profoundly strengthen prior algorithms used for non-intrusive energy disaggregation. I also show that EMI from appliances can also be used to estimate their operating states or mode of operation yielding implications for not only better energy disaggregation but also human activity sensing.

To my parents for their unwavering faith and encouragement
and to my wife Maneet, for her relentless support and love
which makes no task seem impossible.

TABLE OF CONTENTS

Page

List of Figures iv

List of Tables ix

Glossary xi

Chapter 1: Introduction and Motivation 1
1.1 Why Energy Monitoring and Disaggregated Feedback 2
1.2 The Value of Disaggregated Energy Feedback 6
1.3 Sensing for Energy Disaggregation 8
1.4 Thesis Contributions & Statement 9

Chapter 2: Background & Related Work 11
2.1 Distributed Sensing 12
2.2 Non-Intrusive Load Monitoring 13
2.2.1 Low Frequency Sampling Based Features 14
2.2.2 High Frequency Sampling Based Features 18
2.2.3 Voltage Domain High Frequency Sampling Features 20

Chapter 3: Sensing & Fingerprinting Electromagnetic Interference 22
3.1 ElectriSense: Electrical Event Detection & Classification 22
3.1.1 Theory of Operation 23
3.1.2 Prototype System Implementation 26
3.1.3 Proof of Concept Hardware 28
3.1.4 Signal Processing, Event Detection & Feature Extraction Algorithm . 29
3.1.5 Performance Evaluation 31
3.1.6 Discussion and Further Insights 39
3.2 Implications for Energy Disaggregation Algorithms 45
3.2.1 Complementary EMI Features 47
3.2.2 Exclusive EMI Features 48

3.2.3 Combining ElectriSense with Real-time Power Consumption Data . . 48
3.3 Discussion & Overview of Contributions 52

Chapter 4: Going Beyond Sensing Appliance On/Off Events: Using EMI for Operating State Estimation 55
4.1 Overview of Operating State Estimation 55
4.2 Case Studies of Common Appliances 57
4.2.1 EMI Variations of Compact Fluorescent Lamps 58
4.2.1.1 Variations in CFL's Line-Conducted EMI 59
4.2.1.2 Signal Processing Pipeline & EMI Tracking Algorithm 61
4.2.1.3 Observable Signals and Operating States 66
4.2.1.4 Experimental Validation of CFL's Operating States Estimation 70
4.2.1.5 Discussion 76
4.2.2 EMI Variaitons of Modern Televisions 77
4.2.2.1 Overview 77
4.2.2.2 System Description 80
4.2.2.3 Selection of TVs and Movies 80
4.2.2.4 Data Collection Procedure 82
4.2.2.5 Signal Conditioning and Automated Signal Identification . . 83
4.2.2.6 Analysis and Results 83
4.2.2.7 Discussion 88
4.2.3 EMI Variations from LCD Driver Circuitry 89
4.2.3.1 Theory of Operation 90
4.2.3.2 Experimental Procedure and Analysis 93
4.2.3.3 Results and Discussion 96
4.3 Extending to Other Appliances 96
4.3.1 Laptop Power Supply 97
4.3.2 Desktop PC 98
4.3.3 Vacuum Cleaners 99
4.4 Generalized Temporal Tracking Algorithm 100
4.5 Overview of Contributions & Implications for Energy Disaggregation Algorithms 103

Chapter 5: System for Collecting Longitudnal & Naturalistic Appliance Level Power 107
5.1 The Need for Labeled Appliance Level & Naturalistic EMI Dataset 107

5.2 FireFly: Ground Truth Datacollection System 109
5.2.1 Overview of the FireFly System 109
5.2.2 Block 1: AC Interface 112
5.2.3 Block 2: Firefly Node Motherboard 115
5.2.4 Block 3: WiFi Linux Daughterboard 124
5.2.5 Block 4: Local Aggregation Server 130
5.2.6 Block 5: Remote Aggregation Server 134
5.2.7 Block 6: GPS Disciplined Clock and NTP Server 135

Chapter 6: Conclusion and Future Directions 137
6.1 Implications of EMI Sensing for Different Communities 137
6.2 Future Directions 140

Bibliography 143

LIST OF FIGURES

Figure Number Page

1.1 A prototype HTML5 application running on a tablet showing not only the total energy consumed as a dollar value, but also the real-time power consumption while breaking it down by individual appliances. The green circle in the bottom half of the interface lights up when a particular appliance is being actively used. 4

1.2 Fine grained appliance level energy usage information can further be used to provide highly targeted personalized feedback. In this example smart bill, the message highlighted with red outline recommends that the homeowner change their dishwasher; and in doing so the expected savings will pay for the new dishwasher itself in a few years time. 5

2.1 (*left*) Power vs. Time showing step like changes when appliances are turned on or off. (*right*) Clusters of appliances in complex power space. Image courtsey of Hart *et al.* 15

3.1 (*left*) Circuit model of a SMPS with placement of the voltage probe. (*right*) Frequency domain analysis at the voltage from probe showing EMI at 10 KHz. 25

3.2 Continuous voltage noise signatures of various devices during various periods of operation. Colors indicate amplitude at each frequency with blue to red as low to high. Observe that the EMI is observed for the duration of an appliance's operation. 26

3.3 Prototype system consists of a single plug-in smodule, acquisition hardware and the supporting software 27

3.4 Block diagram of major components of our system. 28

3.5 Schematic of the plug-in power line interface module. 29

3.6 (*left*) Background noise observed on a particular power line. (*center*) A new device is turned on, producing EMI that introduces new signals to the power line. (*right*) After background subtraction the new signal features are extracted. The resulting Gaussian fit and its features amplitude (A), mean (μ) and variance (σ) are also shown. 30

3.7 Visual confusion matrix highlighting misclassification due to physical proximity of similar fixtures in H5 and H6. 35

3.8 Variation of features over 6 months for four devices shown in the feature space. Note that no cluster overlaps. 38

3.9 Small, but discernable variation in the mean of the EMI peaks for four of the same model and brand Compact Fluorescent Lamps. 40

3.10 Same CFL lamp plugged into different regions of the home producing EMI amplitude variations. 41

3.11 Band limited EMI generated by a dimmer shown at various dim levels. . . . 42

3.12 Startup burst of EMI signal generated by CFL lamps on ignition. 43

3.13 EMI sensing complements prior power based appliance usage detection. Shown here are two vacuum cleaners being operated briefly. Note the change in power as well as EMI that can be used together for a more robust event detection and feature extraction than either one independently. 46

3.14 A strong EMI signal can be observed despite the appliance consuming only 8.2 Watts. Reliably detecting such small power changes using power-based disaggregation algorithms proves to be a challenging task in practice. On the other hand, the strong EMI complements the prior approaches and makes electrical event detection robust. Also note that the device's power consumption changes over time (step downs and ups) reducing to 1 watt, but the EMI remains detectable. 47

3.15 Center frequency of the EMI from a CFL drifts by 0.5-1KHz over time as the lamp warms up to finally reach a steady state. Sensing such appliance specific behavior allows extracting machine learning features that can make disaggregation algorithms more robust. 49

3.16 (left) Conventional real-time power monitoring requires current transformers to be installed *inside* the breaker panel. Such a task requires professional installation. (right) In contrast, Contact less Current Sensor (CCS) I developed measures real-time power consumption using a sensor that is installed on the *outside* of the panel and is end-user deployable. 50

3.17 ElectriSense combined with Contactless Power Consumption sensor shows how much and where energy is being consumed in real-time. 51

4.1 Simplified block diagram of a CFL. L1-Lamp-C1 forms a series resonant circuit. Human proximity causes the effective capacitance to change and detune the oscillator. 59

4.2 (*left*) A user bringing their hand closer and then touching the lampshade. The blue line shows the energy in the spectrum change as a result. (*right*) Spectrogram showing change in observed EMI as a result of human proximity to the CFL bulb, which causes a large capacitive change resulting in increased energy. 61

4.3 Stages of the signal processing chain that condition the input signal, apply differentiators, detects and segments events. 62
4.4 Summed energy in band over time shows increase in energy as a result of a hover (left), filtered summed energy (middle) and first derivative waveform along with start and stop of the hover event as detected (right). 64
4.5 Summed energy in band (blue line) tracked over time overlaid with the spectrogram showing EMI from CFL. Notice the bump in blue line and increase in harmonic energy at 182 KHz in the spectrogram as a user performs a *hover* gesture or induces a hover operating state for 2.5 seconds. 67
4.6 Notice the difference in amplitude of energy between a one finger touch on bulb and a three finger touch. Blue line is filtered sum of these energies over time. 68
4.7 Amplitude of the observed energy depends on the proximity to the lamp. Amplitude gradually rises when a hand is brought closer over time. 69
4.8 Shift in frequency (top) follows similar trend as ambient temperature (bottom). 70
4.9 Lamps that we used for our experiments. Two lamps (top left, top middle) were chosen to be similar for comparison. All lamps were on during the experiment. 71
4.10 ROC showing tradeoff between the true positives rate and false positives rate for different threshold values. 74
4.11 Cross Correlation results at various query lengths. 87
4.12 5 gestures or induced operating states that can be inferred from the time varying EMI a LCD's internal driver circuitry produces: (1) full-hand touch, (2) five-finger touch, (3) hover, (4) push, and (5) pull. 90
4.13 Simplified schematic of a small section of the LCD panel and the row and column drivers. 91
4.14 EMI seen at the row rate of 67.5 KHz is observed well above the noise level during the user's touch of the LCD. 92
4.15 The summed amplitude (top, blue) of the EMI during a touch event, indicated by the gray highlight. The filtered amplitude (red) and derivative (bottom) are shown along with the threshold used to determine the start and end of the event. 94
4.16 Summed energy curves representing full-hand and five-finger touches (top) as well as the push and pull gestures (bottom). The gray shaded area indicates the touch duration. 95
4.17 (*left*) EMI from a Macbook Pro laptop charger. (*center*) Change in energy density extracted from the EMI. *(right)* Laptop's actual CPU load. Notice how the extracted signal from EMI closely correlates to the CPU load. 97

4.18 EMI of frequency-shift type observed from a desktop PC. Dotted lines indicate manual labels of when the computer was turned ON/OFF and when the OS was deemed to be fully booted. OS shutdown label indicates when the user clicked the Windows OS shutdown button. 98

4.19 Two types of EMI variations and the extracted signals are shown. Appliance operating state changes manifest as (left) change in the center frequency of the EMI causing it to shift; and (right) as change in the energy density in a band-limited frequency region around the center frequency. 101

5.1 Block diagram showing the various components of the FireFLy data acquisition system. 111

5.2 Schematic of the plug-in AC line interface for galvanically isolated and stepped down current and voltage waveforms suitable for input into FireFly DAQ ADC. An AC line interface acts as a pass-through device that is connected to each appliance being monitored. 113

5.3 The black box on the left consists of the the AC interface circuitry that acts as a pass-through circuit for the electric Kettle plugged into it. The voltage and current waveforms for the kettle are captured by the FireFly unit connected to the AC interface via a 4-conductor ribbon cable. 115

5.4 Schematic for the FireFly Node motherboard. 117

5.5 Firefly DAQ node with the daughterboard plugged into the base motherboard. 118

5.6 Strict timing and hard real-time deadlines are maintained on a FireFly DAQ by doubling the amount of time the CPU has to process each buffer using a ping-pong buffer approach. To further minimize CPU load, DMA chaining is employed with a dedicated DMA channel for each buffer bank that hands over to the alternate DMA channel when done requiring no reconfiguration of DMA controllers while operational. 122

5.7 High level flowchart depicting the various functions the microcontroller firmware performs. 126

5.8 The daughterboard on a FireFly DAQ node is a micro Linux computer running the popular OpenWRT distribution. Note the various interfaces such as Serial Peripheral Interface (SPI), Ethernet, USB, TTL232, GPIOs and WiFi that make it a versatile platform. This board communicates with the XMega motherboard using USB and TTL232. *Image courtesy of 8devices.com.* 127

5.9 High level overview of the various tasks the local aggregation server performs. All DAQ nodes deployed at a site communicate with one such server that archives the data, computes application specific metrics and also acts as a central point for the researcher to command and control the various aspects of the node hardware. Not shown is the configuration functionality that allows DAQ nodes in field to update their configuration by requesting from the same server. 133

5.10 A custom GPS disciplined Stratum-1 NTP server built using a GPS module and a Raspberry Pi running Linux. On a local area network, timing coupling or offset of an average of 300 nanoseconds with a jitter less than 3 microseconds can be achieved. 135

LIST OF TABLES

Table Number Page

3.1 A summary of the homes showing the style, size, number of appliances we tested and the number of events. *Longterm 6 month deployment. 31

3.2 The performance using 10-fold cross validation and minimal training for classification for each home. The results of combining the same devices that are spatially near each other as a single event are also shown (*). 34

3.3 The performance of four devices across different homes using a 10-fold cross validation classification to evaluate stability of EMI signatures independent of home. 36

4.1 Average hit rate per user (P1-P10) for user-induced hover state's 1st run. . . 73

4.2 Average hit rate for touch gestures across different participants. 73

4.3 Average energy change (dB/dt) for each lamp and gesture. Notice the difference between the energy values for a hover vs. bulb touches vs. shade touches. Shade touch energy for L2 is low due to the grounding effect of the metal surface. 75

4.4 Hit rate from the two in-home experiments. 76

4.5 A summary of the tested TVs showing make, technology, date of manufacture, and price for each device. 78

4.6 List of twenty movies that we included in our dataset selected such that they span various genres. 81

4.7 A summary of the homes showing the style, size and number of floors. (*) 800 sq. ft refers to only one apartment in a converted multi-family home. . . 83

4.8 Frequency regions and EMI signal types for our 8 TVs. 84

4.9 Analysis of EMI signal similarity given identical and different video content within a TV - averaged over the 15 minute intro segment of all movies. . . . 86

4.10 Analysis of EMI signal similarity given identical and different video content between TV pairs (averaged over all movies). For this we used randomly selected 15 minute streams from each 60 minute movie 86

4.11 Specifications of the six monitors (M1-M6) and two laptops (L1, L2) used in the experiment. 93

5.1 Choosing optimal burden resistor for the CT. Maximum current at 2.5V FS and resolution at 12-bit are shown for each burden resistor. 114

5.2 Timing delay between successive USB transfer function calls and each DMA controller completion. All timings are presented in seconds. Configuration is formatted as LS_X for low-speed USB and HS_X for high-speed USB, where X is the size of buffer in bytes. Notice that HS_322 performed optimally with low average USB write time and least jitter (see spread and standard deviation). 125

GLOSSARY

AC: Alternating Current

ADC: Analog to Digital Converter

AP: Access Point

CCS: Contact-less Current Sensor

CFL: Compact Fluoroscent Lamp

CT: Current Transformer

CV: Cross Validation

DAQ: Data Acquisition

DC: Direct Current

DMA: Direct Memory Access

EMI: Electromagentic interference

FCC: Federal Communications Commission

FFT: Fast Fourier Transform

FS: Full Scale

FSM: Finite State Machine

KSPS: Kilo Sample Per Second

LCD: Liquid Crystal Display

NIALM: Non Intrusive Appliance Load Monitoring

NN: Neural Network

NTP: Network Time Protocol

PGA: Programmable Gain Amplifier

PLI: Power Line Interface

RPM: Revolutions Per Minute

SAR: Successive Approximation Register

SOC: System on a Chip

SMPS: Switched Mode Power Supplies

SVM: Support Vector Machine

ACKNOWLEDGMENTS

I would like to sincerely thank all the amazing people that have helped me on the way to reach this momentous milestone in my life. I could not have done this without their constant support, encouragement and help.

First I would like to thank all my family and friends who never gave up believing in me and were always on my side to give me the much needed encouragement that I can accomplish whatever I want. In particular, I would like to thank my mom, who would would always convince me to keep trying and move forward despite heart-breaking failures. She has taught me how to remain optimistic and seek the light even when it may seem that there is only darkness. I would also like to thank my dad for instilling in me the confidence to accomplish things that I never thought I was capable of and imbibing in me high integrity principals and good will. I can only hope that one day i can be the amazing and kind person that you are. Words cannot do justice to how thankful I am to my wife Maneet for her unwavering love, dedication and selflessness. I could not have done this without you and I am amazed that you have calmly put up with me. I have learned so much from you and you will always be an inspiration and my fuel that keeps me reaching for the sky.

I would like to thank my advisor Shwetak from the bottom of my heart. He has not only been a phenomenally amazing advisor, but also a great friend. I still remember the phone call 5 years back when you asked me consider doing a Ph.D. as my next step and I declined thinking it would be such a bad idea. Well, here we are and I could not be happier that I took the decision to pursue a Ph.D, and in particular under your guidance. The last 5 years have been a profound humbling and character building experience. I am glad I listened to you and will continue to

do so. I would also like to thank you for all the wisdom you have imparted. It still amazes me that I would setup a meeting with you to discuss how I have failed at a particular task and you would show me how I can turn that failure into a success. I cannot remember a time where I have left your office not feeling optimistic and confident that I could accomplish the impossible! Thank you so much.

I would also like to thank Gregory Abowd, my grand-advisor and the reason I am in research! Had you not made the effort and gone out of your way to give me a chance at research, I would have never recognized my strengths. I will forever be in your debt. Thank you for being one of the most kind, humble and amazing person I have ever known.

I would also like to thank all my other advisors that have helped me through graduate school and made me a better person and a better researcher. Matt Reynolds, you have had a profound impact on how I approach my work now. You have taught me to be meticulous and systematic. Your depth of knowledge and intuition amazes me and I can only hope that someday I can accomplish a fraction of what you can. Anthony LaMarca, you have been an absolutely amazing advisor to me. I do not think I would have been convinced to join graduate school had it not been for your persuasion. It amazes me how you so strongly believe in me — thank you so much for that. It really means a lot. I will always seek your help as I have done for the last 7 years. Jeffrey Hightower, thank you so much for being my mentor and introducing me to the world of research. I really appreciate you hiring a first year Masters student with no track-record as your intern and then giving me the intellectual freedom to grow and prove myself. Thank you! Joshua Smith, it has always been a pleasure working with you and it has been wonderful that I can just run up to you and have my questions answered. Gaetano Boriello, I would to thank you for mentoring me the last 5 years and giving me invaluable advise that has shaped both my graduate work as well as professional goals. I truly appreciate all the efforts you made to see me succeed.

I have been so fortunate to have been part of an amazing research group called the Ubicomp Lab! I could have never imagined that "working" could mean so much fun while changing the world for good. All of you guys are amazing and are destined for greatness. I would like to thank Jon Froehlich, Eric Larson, Gabe Cohn, Tim Campbell, Keyu Chen, Tien and Mayank Goel. Jon, my research career was basically jump-started with you teaching me how to write and present it well. I cannot imagine where I am without you taking charge and showing me how its done! Eric, thank you for being a mentor. Without your machine learning and signal processing expertise, I would still be taking my qualifier exams right now. Gabe, we started school together and it is amazing to see how we have grown professionally together to have such a complimentary set of skills. Your attention to detail amazes me and I can only hope I can be half as good as you are. I am really looking forward to accomplishing great thing at MSR alongside you. Tim, you have been an absolute delight to have around. I could have never published my first paper had it not been for your relentlessness help with machining and many sleepless nights of writing. It amazes me how selflessly you chose to help me! Thank you for that. Keyu, thank you for being one of the nicest person I have known. I truly enjoy working with you and your integrity and dedication towards everything you do is phenomenal. I look forward to collaborating with you in future. Tien, it has always been a pleasure working with you. I am really grateful to you for spending countless sleepless nights working and hacking algorithms to support my research. I could not have done many of the projects I did without your support. Mayank, you have been more than just a colleague and a great friend. You have been like family to me and it makes me proud that you are such an awesome researcher. I truly hope we can work together again.

I would like to thank the CUE research group at Microsoft Research for all the research opportunities and trusting me with solving some of the hardest problems. Desney, words cannot express how I highly think of you. You have been an incredible advisor to me and have me helped shape my career, so far and to come. I am very

excited to be working alongside you as my next step! I could not have asked for a more stellar 'boss'.

Finally, I would like to thank all the students at University of Washington who found me worthy of letting me mentor them. I am not sure if I was able to live up to your expectations, but each one of you sure helped me become a better researcher and mentor.

Chapter 1

INTRODUCTION AND MOTIVATION

A key challenge in Ubiquitous Computing is to infer human activity and context to enable and support seamless interaction with the increasingly digitized world around us. Low-cost and easy-to-install technologies to sense and model human activity in the home have thus long been a central focus of ubicomp research. Sensing appliance-level energy usage (i.e., detecting when and which electrical device a user operates and how much energy it uses) has emerged as one particularly promising area for activity-inference research because of how often it reveals the residents current activity (e.g., stove usage implies cooking). In addition to ramifications for activity sensing, sensing energy usage has a huge impact on sustainability applications and eco-feedback, as it enables homeowners insight into not just how much energy they consume but also where it is spent. This appliance-level energy monitoring, also called energy disaggregation has spurred an entire industry that aims to empower users with utility bills that are akin to grocery or credit card bills *itemizing* energy usage allowing users to monitor and act on the spending.

My thesis focuses on development of *sensing techniques that are easy to install, low cost and enable continuous measurement of electrical appliance usage in the home.* In particular, I have applied this sensing for performing accurate and robust disaggregation of energy usage for eco-feedback applications. I accomplish this by developing an important approach that re-purposes signals from already ubiquitous technologies. For example, re-purposing electromagnetic noise produced by electrical devices as a unique "signatures" or signal to infer which devices are being currently used and consequently to disaggregate a home's energy usage.

My approach of leveraging signals from existing technologies requires minimal installation of new hardware making retrofit installations particularly easy while lowering the barrier to widespread adoption. In addition, my sensing systems require fewer sensors mak-

ing them lower cost and easier to maintain, while providing comparable quality of data to traditional multi-sensor systems. In developing these systems, I have built on my belief that the fundamental working principles of already ubiquitous technologies can be leveraged to offer new capabilities that paves the way for next generation of technologies.

The fundamental sensing approach I have invented that I discuss in this dissertation, ElectriSense, is general purpose enough to have a profound impact on multiple high-impact application domains. For instance, not only can it be used to detect when a particular appliance is used but also to help build better predictive models of higher level human activity. In this dissertation I have chosen to focus on one particular application of my sensing technique: Disaggregated Energy Monitoring.

1.1 Why Energy Monitoring and Disaggregated Feedback

Despite accounting for less than 5% of the world's population, the United States consumes a quarter of the world's energy resources. Contrary to popular belief that this is a result of industrial consumption, studies have shown that everyday human behavior such as home energy use and transportation accounts for as much as 28% of the US energy consumption [8]. When commercial building such as office spaces are also taken into consideration, these account for as much as 40% of the U.S energy and green house gases emission. Similar to the US, in the European Union, residential energy usage alone accounts for as much as 26.7% [15] . With exponential increase in energy demand, relying on efficiency and conservation measures to lower the overall energy footprint becomes imperative. To date, the primary methods applied to reducing consumption have been technological and economic. For example, the hybrid vehicle technology has been emphasized as a key solution to green house gases reduction and oil dependence. While many research and policy programs exist that aim to both develop and increase adoption of more efficient home appliances, it should be clear that this can be a slow process and suffer from what is commonly known as the *retrofit* problem or the problem of upgrading existing technology. That is, it is easy to have new homes fitted with efficient appliances and sophisticated sensing, however majority of preexisting homes and buildings are still using older and perhaps, inefficient technology, overhauling which is a very slow and long process. Thus, solutions that promise similar

efficiency gains with minimal changes to existing homes can have a substantial impact.

One such class of solution is human centered behavior change. That is, a set of technologies, policies and programs that inform and motivate environmentally sustainable behaviors. Despite growing evidence that suggests human behavior can significantly affect resource consumption and that intervention strategies to promote conservation should be aggressively pursued, experts believe that behavioral change, or in other words, unwillingness of homeowners to adopt new and efficient technologies is a significant barrier that has held back energy reductions [36]. This is understandable as it is non-trivial to convince people to upgrade their otherwise perfectly working water heater for instance. Substantial research in persuasive and feedback technology has shows that continuous and personalized feedback can motivate homeowners to upgrade in lieu for long term monetary and conservation benefits. However, I believe that the lack of easily deployable and low-cost sensing technologies that enable such feedback applications has held back the proliferation of persuasive eco-feedback systems and policies that they drive.

Around fifty scientific studies have studied the effects of providing homeowners with feedback on their energy consumption and providing specific feedback or more precisely appliance-level usage has been shown to produce the greatest reductions [10] . The case for disaggregated appliance-level feedback is further bolstered by evidence from a recent study of 60 homes across several continents that has shown that significant reductions of up to 12% in energy usage could be achieved through *continuous* and *disaggregated* energy usage feedback to consumers [19]. It further articulates that maximum energy savings were observed using direct feedback such as real-time appliance level usage information as opposed to indirect feedback like monthly bills or weekly advise.

Disaggregated energy feedback promises an in-home feedback display that tells a homeowner not just their total power consumption, but also continuously shows real-time usage while breaking it down categorically by electrical appliances. In addition, such a system could provide personalized and cost-effective energy saving recommendations. For example, it could report, "Based on your usage patterns, you could save $215 per year by switching to a more efficient heating unit, which will pay for itself in 27 months". Figure 1.1 shows a prototype application running on a tablet, which not only reflects total energy consump-

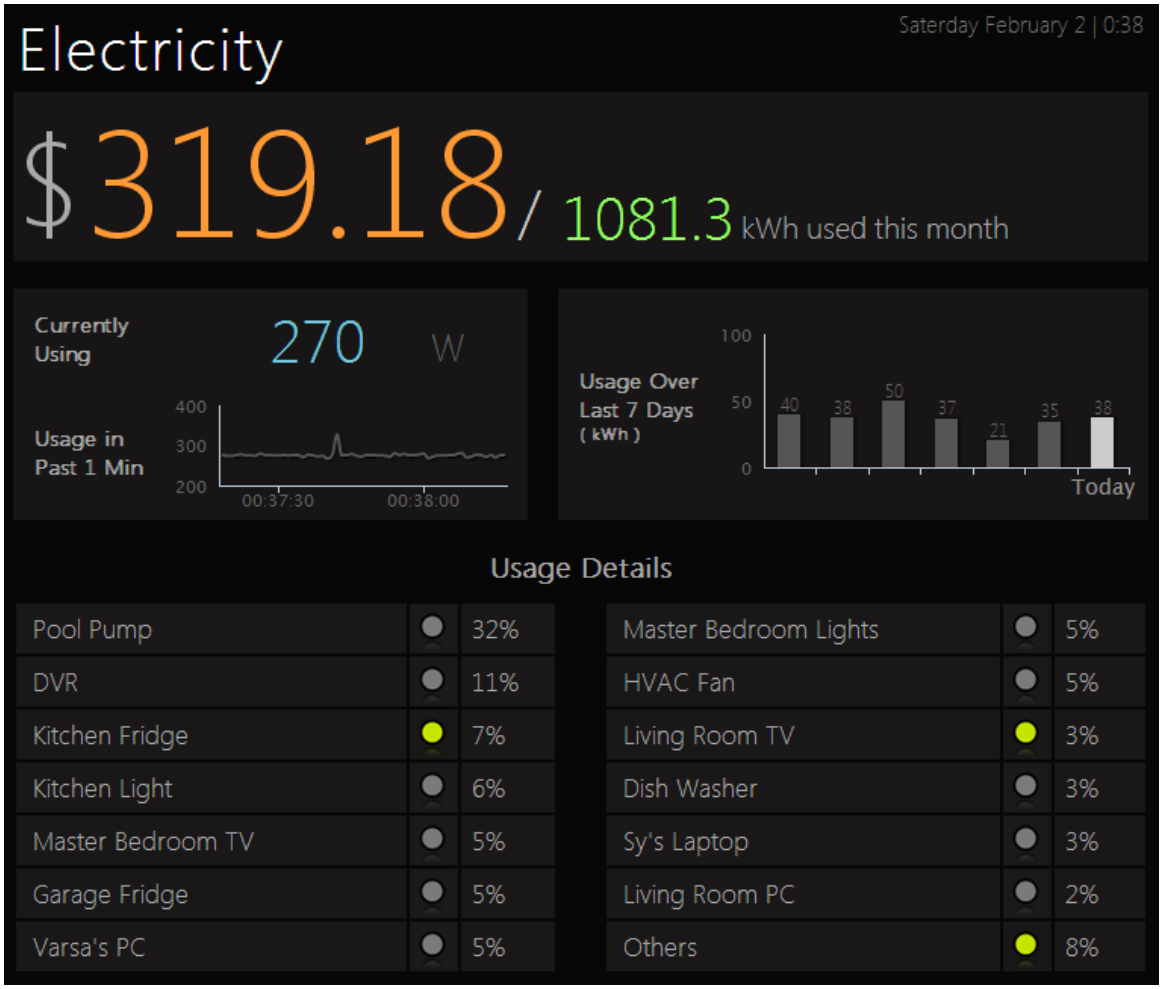

Figure 1.1: A prototype HTML5 application running on a tablet showing not only the total energy consumed as a dollar value, but also the real-time power consumption while breaking it down by individual appliances. The green circle in the bottom half of the interface lights up when a particular appliance is being actively used.

tion in dollars so far and the real time energy usage, but also disaggregates it into various appliances in the home. Appliances being currently used are indicated with a 'green light' on the interface, along with what percentage it contributed the home's net energy usage.

In addition, such fine grained appliance-level energy usage information can be used to provide highly targeted and personalized feedback. For instance, Figure 1.2 shows an elec-

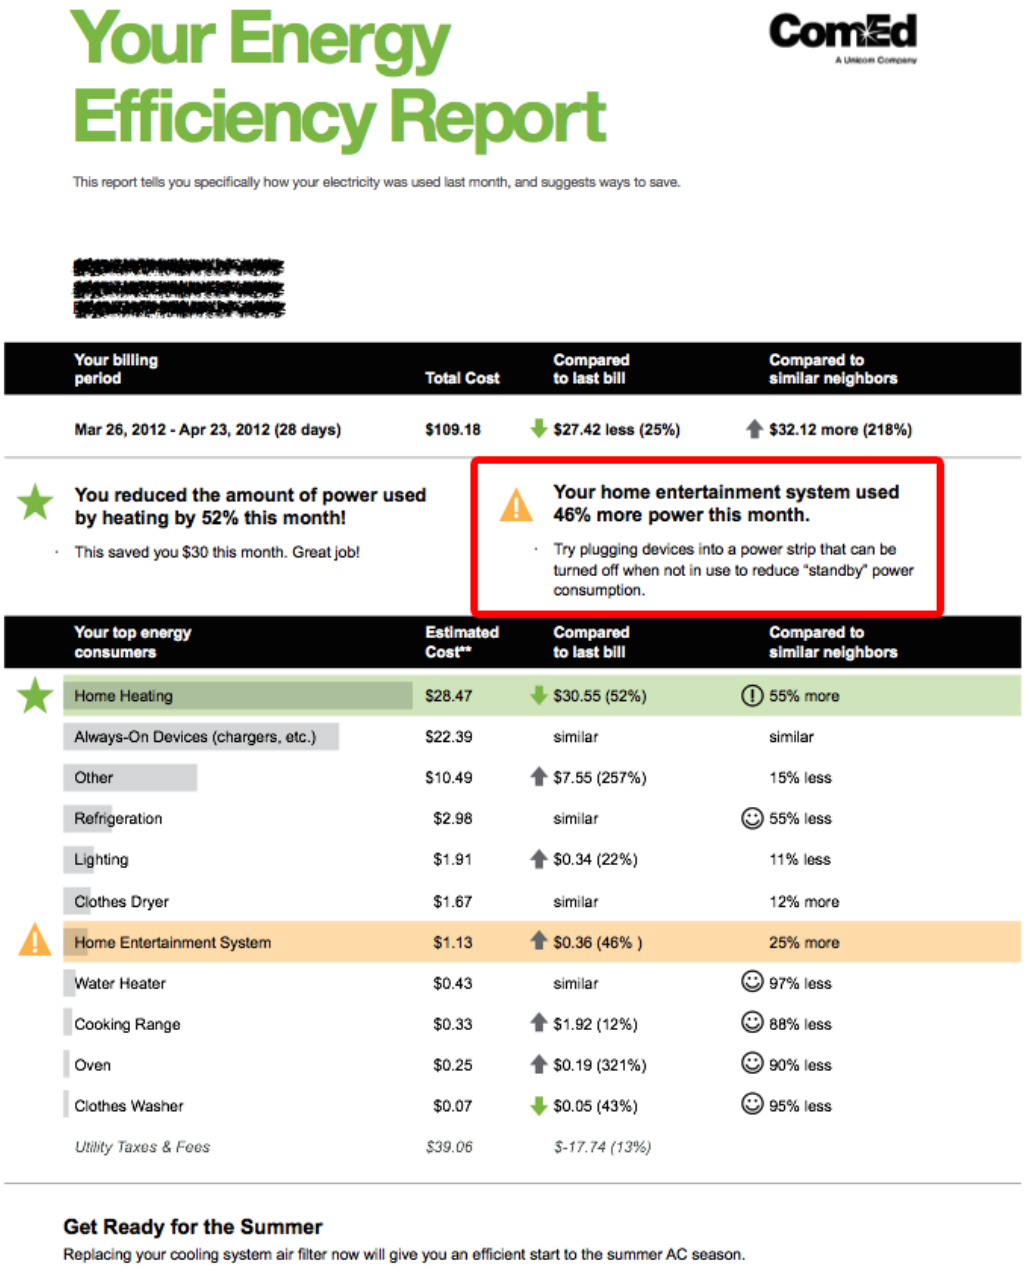

Your Energy Efficiency Report

ComEd
A Unicom Company

This report tells you specifically how your electricity was used last month, and suggests ways to save.

[redacted]

Your billing period	Total Cost	Compared to last bill	Compared to similar neighbors
Mar 26, 2012 - Apr 23, 2012 (28 days)	$109.18	↓ $27.42 less (25%)	↑ $32.12 more (218%)

You reduced the amount of power used by heating by 52% this month!

- This saved you $30 this month. Great job!

Your home entertainment system used 46% more power this month.

- Try plugging devices into a power strip that can be turned off when not in use to reduce "standby" power consumption.

Your top energy consumers	Estimated Cost**	Compared to last bill	Compared to similar neighbors
Home Heating	$28.47	↓ $30.55 (52%)	! 55% more
Always-On Devices (chargers, etc.)	$22.39	similar	similar
Other	$10.49	↑ $7.55 (257%)	15% less
Refrigeration	$2.98	similar	☺ 55% less
Lighting	$1.91	↑ $0.34 (22%)	11% less
Clothes Dryer	$1.67	similar	12% more
Home Entertainment System	$1.13	↑ $0.36 (46%)	25% more
Water Heater	$0.43	similar	☺ 97% less
Cooking Range	$0.33	↑ $1.92 (12%)	☺ 88% less
Oven	$0.25	↑ $0.19 (321%)	☺ 90% less
Clothes Washer	$0.07	↓ $0.05 (43%)	☺ 95% less
Utility Taxes & Fees	*$39.06*	*$-17.74 (13%)*	

Get Ready for the Summer

Replacing your cooling system air filter now will give you an efficient start to the summer AC season.

Figure 1.2: Fine grained appliance level energy usage information can further be used to provide highly targeted personalized feedback. In this example smart bill, the message highlighted with red outline recommends that the homeowner change their dishwasher; and in doing so the expected savings will pay for the new dishwasher itself in a few years time.

tricity bill enhanced with disaggregated appliance level energy usage as well as personalized feedback. Of particular interest is the message highlighted with a red bounding box that suggests that the homeowner switch to a modern energy-efficient dishwasher which will save them $180 per year.

The challenge with such feedback systems is *how* to sense which electrical appliance consumes how much energy to provide such fine grained appliance level feedback. Recent advances in metering technology has allowed many of today's feedback applications to source their energy usage data from smart meters. Though these meters provide a closer temporal coupling between energy usage and feedback (down to 10 minute intervals), the focus is still on the aggregate consumption power consumption, i.e., whole-home usage instead of appliance specific consumption. Appliance level power consumption promises to redefine the way homeowners, policy makers and appliance manufactures think about how energy is consumed in the home.

1.2 The Value of Disaggregated Energy Feedback

The numerous benefits of disaggregated energy feedback can be categorized into those for homeowners, utilities, policy makers, researchers and appliance manufacturers. With regards to benefits for an end consumer or homeowner, the effects of providing real-time and appliance-specific feedback on their electricity consumption has been investigated in more than fifty studies. Although limited and more extensive research is required, the findings have been coherent and have suggested the greatest savings from appliance level feedback. For a detailed review, see [19]. Research in environmental psychology has uncovered profound misconceptions about energy usage in the home. In an interview with 52 households in Germany, Mettler-Meibom *et. al* when asked to estimate energy consumed by each appliance and compared to the actual usage, they found that the consumers vastly underestimated the power consumed by heating while over estimating for appliances like stove and lighting. That is, energy that was 'visibly used' used was often overestimated [35] . Kempton *et al.* conducted a survey of 400 families in Michigan and found that the average consumer incorrectly believed that they could save twice as much energy by reducing lighting than by reducing hot water usage from their electric water heater.

Unfortunately, the increasing trend towards automated appliances that require little to no manual turning on or off, such as heating and cooling systems have brought convenience at the cost of knowledge of energy usage. Not only do people make systematic errors in estimating energy use, but they act on this misinformation. For example, by bringing about changes in short-term behavior like turning lights off when leaving a room and underestimating technical innovation solutions such as deciding to upgrade an inefficient appliance, consumers find that their actions have little or no observable effect on their energy bill and as a consequence they tend to reject the whole notion of energy conservation. Actionable and timely feedback can remedy this problem, both aspects that appliance level feedback promises to capture.

Once a disaggregation system in place, from a policy makers and utilities perspective, the knowledge of appliance level (in)efficiencies and consumer's usage behavior patterns, behavioral approaches such as incentive programs, community drives and media could be much more effective by tapping into personalized recommendation and targeted campaigns similar to how the advertising field uses demographics and other indicators. This knowledge is also critical for development and evaluation of evidence backed energy policies designed to reduce capital expenditure on additional capacity. The effectiveness of such policies and changes to utility infrastructure could be measured in terms of new metrics not possible before - such as in terms of specific appliance categories and knowledge of which individual households or even communities use a particular set of appliances. Thus recommendations or programs could be designed targeted at specific groups that are expected to yield the best results, akin to personalized music recommendation by online music service Pandora. For appliance manufactures and governments, disaggregated appliance data could provide comparisons on how an appliance performs under controlled test conditions to measurements in actual homes, allowing development of better energy usage models.

From a research perspective, easy to install sensors that provide specific appliance usage can spur innovation. For instance, appliance usage can be linked to activity monitoring for elder care applications or simply building computers that better understand context through estimating what activity a homeowner is performing. For example, detecting freezer door open and close followed by the microwave being used could suggest reheating a frozen meal,

while detecting stove usage may suggest cooking. Early fault warning is another area of research that could benefit from monitoring individual appliances and detecting anomalies and warning the user about potential failures or safety issues (for example, detect electrical arcing, a potential fire hazard). Monitoring appliance usage and providing disaggregated energy usage can transform the way we think about energy and activity sensing in the home. However, the challenge is in how to do so in a way that is low-cost, easy to install and require minimal installation.

1.3 Sensing for Energy Disaggregation

The two current methods for disaggregation are *distributed sensing* and *single-point sensing*. The key difference in the two approaches is the number of sensors or measurement units that need to be deployed, which further dictates their software complexity, ease of installation (whether they require installation by homeowner or a professional), cost and how accurate these systems are. Distributed sensing requires a sensor to be installed at each appliance or device to measure its usage making installation difficult and time consuming. In contrast, single-point sensing aims to provide the same level of information as distributed sensors without the need to install one sensor per device.

With ElectriSense, a single-point sensing technique I have significantly advanced the state-of-the-art in disaggregated energy sensing by inventing and studying a new approach that leverages electromagnetic interference (EMI) produced by modern consumer electronics to sense and detect appliance usage. Unlike past NIALM approaches (detailed in Chapter 2), my approach requires only a single plug-in sensor minimizing the cost for both the hardware and installation as it does not require a professional electrician for installation. Additionally, for many appliances, this approach provides new and complementary information (detailed in Chapter 3, Section 3.2) about the the appliances as well as new information that can help infer an appliance's operating state. For example, inferring whether the vacuum cleaner is running on a low speed or high, and if the TV is set to receive over-the-air transmission or over a HDMI connection (detailed in Chapter 4).

1.4 Thesis Contributions & Statement

I have invented and studied a new sensing technique called ElectriSense that requires minimal instrumentation for sensing electrical appliance usage in the home. The core contribution of my work is in leveraging a signal that is otherwise considered "noise" by mainstream approaches. This noise, Electromagnetic Interference (EMI) is produced by most consumer electronics and in this dissertation I show how it can be re-purposed as a signature or fingerprint to detect and classify the use of electrical appliances in real-time. The applications for this sensing technique are myriad, however I primary demonstrate how it applies to sensing electrical appliance usage for energy disaggregation purposes in the home. I show that these EMI signatures allow for extracting machine learning features that are exclusive and complementary to features from prior approaches for inferring appliance usage. In particular, these are past approaches that make use of features or patterns from aggregate power consumption and human behavior. By combining EMI signatures with such past approaches I argue that this work significantly advances the current state-of-the-art in energy disaggregation algorithms.

In addition I demonstrate that the new EMI sensing approach provides a signal that is rich in information to enable new applications that go beyond just estimating whether an appliance is on or off. Specifically, I demonstrate that EMI can be used to infer the internal operating states or modes of an appliance, which can be used to enable new applications, better energy disaggregation models and build better human activity models. I discuss two such applications in this dissertation - application to novel user interaction that allows inferring how a human is interacting with a CFL or non-touch LCD; and inferring TV viewing habits of a user. Since both these applications improve human activity modeling and how they interact with electrical appliances, it has profound implications for energy disaggregation. Lastly, I have developed a high performance ground-truth data collection hardware-software framework that allows for collecting appliance-level energy usage data in a naturalistic and out-in-the-wild setting. I have used this framework to collect a first of its kind research dataset that I will make available to academic research communities around the world. The details of this system are discussed.

I present the following thesis statement:

ElectriSense enables low-cost and practical sensing of (1) electrical appliance usage in real-time using a single plug-in sensor. Applied to energy disaggregation, it provides (2) new and exclusive machine learning features that lead to a more robust energy disaggregation algorithm. In addition, temporal EMI sensing can be leveraged to infer (3) an appliance's internal operating states, which further enhances energy disaggregation by allowing the system to (4) better model human activity and their interaction with electrical appliances.

In this dissertation, though I focus on disaggregated energy feedback as an application of EMI sensing, it should be noted that this is just one of the many applications that the general approach of using EMI signatures can enable. In Chapter 4 I briefly discuss how I have applied this technique for human proximity detection and gesture recognition using unmodified CFLs and LCD monitors.

Chapter 2

BACKGROUND & RELATED WORK

Field surveys and self-report have traditionally been the most common way to understand how energy is consumed by end users, but as described in Section 1.2, consumer self-report could be slow and error prone. In contrast, automated methods can provide very rich datasets that can potentially scale to many thousands of homes. However, currently no commercially available disaggregation systems can automatically measure energy usage down to individual appliance, be accurate and easily deployable. The need for such a system has spurred immense interest amongst the research community since the early 1990s and more recently government labs that study and model home energy usage and the commercial sector that aims to monetize on personalized feedback and targeted marketing, resulting in a wide variety of approaches.

The two current methods for disaggregation can be categorized into *distributed sensing* and *single-point sensing*. The key difference in the two approaches is the number of sensors or measurement units that need to be deployed, which further dictates their software complexity, ease of installation (whether they require installation by homeowner or a professional), cost and how accurate these systems are. Accordingly, this chapter is divided into two sections, one to discuss the application of distributed sensing to energy disaggregation, and another for background and related work pertaining to non-intrusive load monitoring. ElectriSense falls in the latter category, however it is introduced and described in detail in Chapter 3.

Section 2.1 describes the one sensor per appliance model for achieving energy disaggregation. Understanding the limitation and challenges with this sensing model will allow the reader to appreciate the single point sensing or NIALM approach and the contribution that this dissertation makes.

Section 2.2 discusses the background and related work in single point sensing solutions,

the various research and commercial technologies available to monitor energy usage and their limitations. This section also briefly describes the set of features used for disaggregation and the fundamental challenges that ElectriSense overcomes with its novel use of EMI sensing.

2.1 Distributed Sensing

A common approach for obtaining appliance-specific usage data is to install wireless plug sensors on individual appliances and construct a home wide network. Installation generally entails unplugging the appliance, plugging the sensor in to an electrical outlet and plugging the appliance into the sensor. Thus, the sensor is in-line with the appliance. These sensors can either form an ad-hoc network, or more simply broadcast to a central base station that collects and aggregates data for processing and display. The PlugWise [44] and Tendril Volt II [52] are examples of such commercial wireless plug sensors that integrate with a web portal that processes and displays power consumption per plug sensor. Kill-a-Watt is a popular plug sensor that does not feature any wireless or telemetry functions [25]. The intended usage is to use the built in display to observe how much energy a particular appliance consumes.

Though device level sensing is conceptually straightforward, it requires time consuming and expensive installation and maintenance. Cost is high for such sensors ranging anywhere from $30-$100 per sensor and many require a specialized wireless hub. More than the cost, the key drawback of this approach is the tedious installation. Appliances that tend to consume the most energy (for example, a water heater) are hard-wired or difficult to reach (for example, behind a dryer unit) and thus present a challenge. It should be noted that the plug sensors would need to be installed throughout the home, on perhaps *all appliances* to avoid missing 'invisible energy consuming' or 'surprise' devices. As mentioned before, consumers underestimate energy for automated appliances that do not 'visibly' consume energy, such as a water heater, which also happens to be a big energy consumer. It is the energy usage of such appliance that a disaggregated energy feedback system aims to reveal to the consumer, and with installation of a limited number of plug sensors due to cost and installation barriers, the decision of selecting which appliances to instrument is made by the very consumer that misunderstand energy consumption. Plug sensors are thus prone to

missing critical energy saving opportunities.

To overcome some of the installation problems of plug sensors, a more recent approach in distributed sensing has been *indirect sensing* where sensors do not need to be plugged in-line with the appliances. Instead, these are placed close, on or near an appliance. See [22,28,43] for such sensing techniques that leverage light changes, sound, magnetic flux changes and radiated EMI to infer when an appliance is operation. of course, these techniques are sensitive to other changes in the environment and may not be robust in the long run, however they show promise. Additionally, they operate on battery to overcome the installation barriers, but at the cost of increased maintenance effort. Also, these techniques are yet to mature to infer different operating states of an appliance. For instance, a washing machine goes through multiple cycles generating different light, sound and magnetic signatures, each of which have to be calibrated. Despite the shortcomings the general technique of indirect distributed sensing shows promise and could be an integral part of a hybrid solution combining NIALM, direct and indirect sensing.

Before discussing single point sensing, a key advantage of plug sensors worth noting is that their calibration entails simply providing a label to each sensor, unlike sophisticated machine learning and training required with NIALM approaches. Additionally, these sensors often offer control. That is, the ability to turn the appliance on or off remotely - a feature single point and indirect distributed sensing lacks. A hybrid approach could potentially be optimal allowing plug sensor installation only on devices where control provides substantial benefit while disaggregating entire home's energy using an easy to install NIALM system.

2.2 Non-Intrusive Load Monitoring

Single-point sensing aims to provide the same level of information as distributed sensors without the need to install one sensor per device. Instead, a single sophisticated sensor is used that can infers appliance usage for the entire home. Since this approach is non-intrusive from an installation perspective, George Hart in the early 1990s coined the phrase Non Intrusive Appliance Load Monitoring (NIALM). There has been substantial research in the field of NIALM that spans different technologies and hardware complexity but primarily the focus has been on identifying new features and development of learning algorithms. Despite

the variety of available NIALM methods, the fundamental components in each are the same. First, they all have a set of features or signatures specific to appliances that are selected and characterized using machine learning techniques. Second, data acquisition hardware that captures the intended features, and finally, a set of algorithms that detect and classify appliance usage based on the features.

In rest of this section I provide a background of the various NIALM techniques and approaches organized by data acquisition hardware's sampling rate, which closely follows the type of feature set selection. Within each section, I will describe both technological contributions made over the past years as well as research in machine learning that focused on energy disaggregation as their application.

2.2.1 Low Frequency Sampling Based Features

Low frequency sampling hardware are primarily intended to provide stronger temporal coupling between usage and feedback, but not necessarily designed with NIALM in mind. For example, smart meters are designed to report aggregate power every 10-15 minutes, which is much more "real time" as compared to traditional meters they replace that had to be read manually and feedback from the utility was once per month or two. In addition to the smart meters, there are several commercial devices that an end consumer could install at one location, i.e. single sensor, which provide more frequent feedback. These devices are typically installed in the breaker panel or the main electrical junction box of a home. Some examples of these devices include the Energy Detective (TED), Flukso and AlertMe.

The reason these are classified as low frequency sampling hardware is that they report real power and voltage from anywhere between one sample per 15 minutes (smart meters) to one per second (TED, Flukso and AlertMe). Since the fundamental frequency at which power is delivered to homes is 60 Hz in the United States (50 Hz in Europe and Asia), such low frequency sampling hardware can only capture coarse grained or *macro-features*, such as changes in complex power, phase, power profile over time, and changes in real power.

Pioneering work in NIALM methods performed by Hart et al. leveraged changes in a home's *total real and reactive power as features* [24]. As the authors report, this approach

could only be applied to appliances that (1) turned on or off, which include most household appliances like a light bulb or toaster oven; and (2) those that can be represented with a finite state machine (FSM), for example, appliances that go through a particular set of states such as a washing machine or dishwasher. This approach is ineffective for always-on appliances like an alarm clock and continuously varying consumer devices like a hand drill or a modern television. The algorithm described by Hart et al., first involves normalization of the measured power to reduce susceptibility to variations in the voltage. This is necessary since the features itself are step changes in real and reactive power that could seem low or high for the same appliance with variations in the voltage. Hart proposed the following normalization, with 120 being the nominal AC supply voltage in the US:

$$P_{norm} = \left(\frac{120}{V}\right)^2 P_{measured}$$

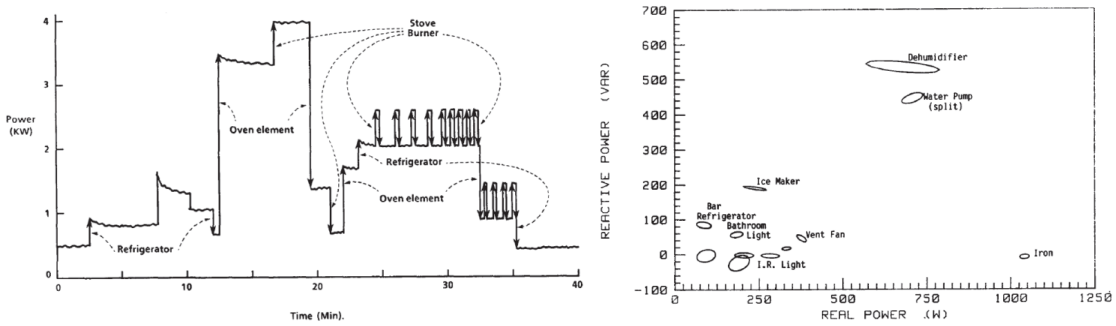

Figure 2.1: (*left*) Power vs. Time showing step like changes when appliances are turned on or off. (*right*) Clusters of appliances in complex power space. Image courtsey of Hart *et al.*

.

Second, the algorithm performs an edge detection to identify the times and sizes of all step-like changes. Figure 2.1 (left) demonstrates this step. Next real and reactive power is used by a clustering algorithm that reduces the problem to a 2D space as visualized in Figure 2.1 (right). As evident from the figure, for the chosen set of appliances, this approach clusters the appliances well. With cluster of step changes generated, a cluster-pairing method is used next to match the negative and positive clusters of similar magnitude

together. These pair form the on/off of a particular appliance which is further verified that the two clusters largely alternate to give an on/off sequence. Clusters that do not match are processed through an anomaly resolution scheme to associate these with either existing clusters or generate new clusters as per the maximum likelihood. Finally, based on the power consumption and statistics of operation duration, these clusters are assigned appliance names and power consumption levels to produce a disaggregation. Similar to Hart *et al.*, more recently researchers have investigated the use of support vector machines (SVM) to detect and classify the edge events caused by appliances [26].

The edge event detection method presented above works well for appliances in the on/off category but can suffer problems for appliances that go through multiple states and those with variable loads. For example, changing a setting on the washing machine could require a new FSM model. Additionally, two appliances turned on or off in short duration of each other could potentially cause the edge detector to trigger only once or confuse the FSM. This is primarily dependent on the averaging window used to smoothen the signal, which if larger would miss events that happen close to each other; and a window too large would risk increasing false positives in edge detection. The cluster-pairing algorithm used is based in an assumption that the 'on' power change closely matches the 'off' power change of a device. Unfortunately, this is not always true as many appliances regulate their power draw when operational. For example, a modern television could go through power draw changes of as much as 60%. Finally, appliances that consume similar power may not be separable with this technique [31].

Recently in [30], Kolter et al. have devised an approach that uses only power signal as the feature and makes use of additive factorial hidden Markov models, an extension to HMMs where the state factors into multiple independent chains and the output is an additive function of the hidden states. The approach seems robust, however since it was the first such application of additive factorial HMMs to disaggregation it is not clear how it scales to use in practice. For instance, the paper validates the inference against only circuit branch level power usage and not down to each appliance. Additionally, only nine appliances were considered. A typical home can have anywhere from 30-40 appliances including lighting, and two open questions remain. First, how computationally expensive this approach becomes,

and second how accurately it disaggregates given the large number of appliances. The use of additive factorial HMMs however show promise. These could potentially be used to disaggregate by circuit to assist other algorithms for further refinement.

Additional macrofeatures were used in Norford *et al.*'s follow up work in [37] that extends Hart's approach to remedy some of its shortcomings. The authors observed that in commercial settings that the length of the transient events, in particular the appliance start-up was particularly long (on the order of minutes), and that the reactive power change was relatively small causing problems with edge detection and clustering. To compensate for these, Norford *et al.*'s approach uses the *time varying shape of the power curve* as an additional feature. In [12, 13] Cole *et al.* extend Hart's work and add two additional feature, namely, upward spike in power and the follow-up slope. The spikes are observed for a brief period when an appliance turns on, and the slope results from it getting to a steady state following the spike. This is most commonly observed for appliances that have large capacitors, such as pool pumps where the inrush current is greater as a result of capacitors charging than that in the steady state.

Baranski *et al.* have demonstrated a genetic algorithm based NIALM classifier that requires no training and detects frequently occurring patterns from the power consumption profile sampled at one hertz. The fundamental idea is to identify repetitive appliance structures without any a priori knowledge of appliances. Though not particularly mentioned by the authors, I believe this technique could also capture repetitive human behavior patterns and not just those from automated appliances like a refrigerator. Since a genetic algorithm is employed, it is not guaranteed to produce an optimal solution.

Though Hart's approach already uses low sampling hardware and simpler features when compared to more modern techniques as discussed in the next section, research has also been conducted to produce disaggregation with even simpler features that do not use real and reactive power. In [45], Powers *et al.* use sparsely sampled power data at a rate of one observation per 15-minutes and apply a rule based classifier that fits the power usage to proprietary load profiles of appliances. The load profiles are limited as the approach only considers large loads in a commercial setting, thus the level of disaggregation is quite different from that in a home. Additionally, the 15-minute sampling rate inherently limits

the variety of appliances that could be disaggregated. For instance, a toaster oven could turn on and off within that time and go undetected.

Recently, many commercial entities have undertaken energy feedback and disaggregation. Opower [38] and Bidgely [7] are two such examples. They source their data from smart meters or TED like devices. Thus, the primary features available to them are power step changes and in some cases real and reactive power at a rate of no more than one sample per second. Due to the proprietary nature of their algorithms, it is not clear how accurate the disaggregation is and what techniques are being leveraged.

2.2.2 High Frequency Sampling Based Features

The limited sampling rate, much lower than the fundamental 60 Hz AC frequency has the disadvantage that features such as harmonic distortions due to non-linear loads, power quality, power factor and higher frequency EMI could not be used as features. Most researchers in the NIALM domain concede that the macro-features from low sampling rate alone are not enough for highly accurate systems that can cover the growing range of consumer electronics. High frequency sampling for producing micro-features can itself be divided into two classes. Majority of the prior work described here limits the sampling to a few kilohertz, however the approach I investigated as part of this dissertation work samples at as high as 2 MHz. Additionally, with the exception of the approach taken by Patel *et al.* [41], all systems described in this section pertain to sampling power and current waveforms, however both Patel's approach and that presented in this thesis leverage measurements of the voltage line.

The need for higher sampling can be better understood by realizing that two electrical devices, such as a fans and a hand drill, despite having similar power consumption profiles could have very different 60 Hz harmonics, which can be leveraged as a distinguishing feature. Laughman et al. in [31] took this approach, where they employ the 3rd harmonic as an additional feature in addition to those proposed by Hart.

Further extension to the use of harmonics can be found in [32], where Leeb et al. use distinct features in the current waveforms by isolating harmonics and analyzing the spectral envelope over a fixed duration during the device start-up for transients. The spectral envelop

is formed using the first-N coefficients of the FFT of the current waveform. The task of transient event detection was accomplished by searching for a time pattern of segments that show significant variation. Whether a change is significant or not, in other words the threshold, is learned during the training phase where the event detector employs a change-of-mean detector to segment a transient that best represents a particular appliance. Once transient shapes are recorded, it forms a database that is searched for every N samples of measurement data. Any match to the the known library of start-up features is used to classify the unknown loads. Leeb et al. were able to show that this feature space is less susceptible to overlapping categories and is therefore able to separate many devices with similar load characteristics compared to approaches based on aggregate power consumption. For example, two motors with similar real and reactive power consumption can exhibit significantly different start-up features making them easy to differentiate.

Though this approach of using spectral envelops can differentiate between similar loads and detect variable power appliances, the method suffers from several drawbacks. First, it is not clear how this approach reacts to changes in the electrical infrastructure, i.e., does adding new appliances to the home significantly change the features or not. Secondly, excessive training is required that can make deploying such a system in practice difficult. Lastly, the performance has not been studied and characterized in a practical and naturalistic setting.

In [50] Srinivasan *et al.* have developed an approach that extends the use of harmonic as transient features to include continuous measurement of harmonics. This makes the features robust. Since transient features extraction must depend on a transient event detector, like any other decision system, it can miss events. A continuous measurement approach on the other hand can potentially minimize such missed feature extractions. However, continuous monitoring also implies that multiple appliances could be operational at the same time, thus either the system needs to be trained for all possible combinations of all appliances turning on/off or signal component separation is required. The authors chose the former approach and trained a Neural Network (NN) with all 256 possible combinations for 8 appliances to demonstrate their system's performance. Though the performance of most appliance usage detection was in high 90% range, it is not clear if such a system is feasible in practice due to the large number on/off combinations for even a modest home with 20-30 appliances.

Neural Network based disaggregation approaches have also been investigated by researchers in [2, 3, 17, 33, 34, 46], however each approach assumes an extensive training dataset often requiring power profiles of appliances in isolation and over longer periods of time to capture the duty cycles that they exploit. Unfortunately, though a promising approach, it reduces the practical applicability of such system.

2.2.3 Voltage Domain High Frequency Sampling Features

All of the above described approaches require that the sampling hardware be *in-line* with rest of the home as the current waveform needs to be observed. Described in this subsection are approaches based on features derived from the voltage domain. It should be noted that sampling the voltage domain can be done from any available electrical outlet in the home, since the hardware needs to be in-parallel to the electrical network.

Using voltage-domain measurement for electrical device disaggregation at first seems counter-intuitive. The incoming power feed to a home is often assumed to be a well-regulated 60 Hz pure sine-wave AC source. This may be true at the point of generation but it is not true within a home. Instead, appliances conduct a variety of noise voltages back onto the power wiring of the home. In an attempt to limit interference among devices, the Federal Communications Commission (FCC) Part 15 and Part 18 rules restrict the noise voltage each device may conduct back onto the power line. Researchers have found, however, that devices that comply with the FCC's limits still yield measurable noise signatures that are easily detectable using appropriate hardware. These signatures occur at a broad range of frequencies, not just 60 Hz and its harmonics.

Pioneering work in leveraging noise on the voltage domain for appliance usage detection was presented in [41] by Patel *et al.* In their approach, the authors leveraged transients generated by mechanically switched incandescent, heating, and motor loads to detect and classify electrical events. This work showed that transient noise is generated whenever a mechanical switch is turned on or off. Any mechanically switched device, such as a light switch, induces a transient pulse when making or breaking an electrical circuit. This is the result of a rapid series of openings and closings due to the physical nature of the switch

itself. The arcing and consequently broadband noise, produced as a result typically lasts only a few microseconds and consists of a rich spectrum of frequency components ranging from a few KHz up to 100 MHz. As this broadband pulse travels through the electrical infrastructure, it excites the step response of electrical wiring's transfer function, resulting in a unique signal as observed at the point of sensing.

Several problems with this transient approach motivated me to develop ElectriSense, the sensing technique introduced in this thesis and presented in Chapter 3. First, transients are an extrinsic property of mechanically switched loads, making their characteristics unpredictable and hence requiring *strict* supervised training for each physical device. Thus, signatures learned from one device cannot be applied to another even if they are similar. This also poses a challenge for identifying mobile devices, as the transients generated by a particular device change as the device is moved from one location to another. Secondly, Capturing and analyzing transient noise is a computationally expensive process, and requires that the system continuously capture and analyze every transient noise event for reliable detection. In addition, there is no way to coherently integrate these transient events to improve detection likelihood as they occur infrequently and are relatively weak due to their broadband distribution of signal energy. Lastly, and most importantly, since the measured broadband pulse is dependent on the step response of the home's electrical wiring, or in other words, the home's electrical wiring acts as a filter that characteristically changes the pulse - *any change* in a home's wiring could potentially alter the observed broadband pulse requiring re-training of the machine learning models. For example, plugging in a Christmas tree may alter the wiring's transfer function thus potentially reducing the practicality of this sensing approach and its application to energy monitoring.

Another key motivation for development of EMI based sensing technique is the fact that most modern consumer electronic appliances are moving towards having a "soft switch", i.e. unlike a mechanical switch they make use of a software driven push button that cycles the power to the appliance electronically (for example via the infrared remote control of a flat-panel TV). In such devices, the indirect activation of the device by software driven electronic switch minimizes the transient generated at the moment of activation.

Chapter 3

SENSING & FINGERPRINTING ELECTROMAGNETIC INTERFERENCE

Whereas the previous section provided background into appliance load monitoring and non-intrusive sensing techniques, this chapter details my work in EMI sensing. In addition to the core electrical event detection and classification, also discussed are implications of EMI sensing to energy disaggregation in the home.

3.1 ElectriSense: Electrical Event Detection & Classification

ElectriSense [23] uses a single plug-in sensor to detect and identify disaggregated electrical appliance usage in a home by using the incidental electromagnetic interference (EMI) generated by modern consumer devices as an identifying signature. That is, it can detect *when* and *which* electrical device a user turns on or off in real-time. As detailed in the previous chapter, previously Patel *et al.* in [41] introduced a single-point sensing method for automatically identifying the usage of resistive and inductive electrical loads. ElectriSense, on the other hand is a fundamentally different approach and focuses on sensing "'electrical noise" that is continuous as opposed to being a transient. This noise is primarily produced by appliances that make use of a switched mode power supplies (SMPS) and motor based appliances like vacuum cleaners.

Devices that rely on SMPS have become increasingly prevalent because of their higher efficiency, smaller size, and lower cost compared to traditional power supplies. For eample they are commonly found in most modern consumer electronics (e.g., PCs, LCDs, high efficiency washers, etc.), in addition to common household compact fluorescent light (CFLs) bulbs. Manufacturers are increasingly employing SMPS in their products to meet minimum efficiency requirements (e.g., the Department of Energys Energy Star program [51]). For example, in one of our deployment sites, all lights and most appliances used SMPS. In contrast to Patel *et al.*s approach, which sensed voltage transients from the activation/de-

activation of resistive and inductive electrical loads, I examine the continuous EMI signals generated by all SMPS based devices due to their reliance on asynchronous high frequency oscillators for operation.

I have found through experimentation that the EMI generated by SMPS has a highly repeatable frequency-domain signature that can be sensed and identified throughout a typical home during device operation. Perhaps more importantly, these signatures are largely specific to a devices circuit design and maintain consistent properties across homes. This is in contrast to much of the work in infrastructure-mediated sensing [11, 21, 40–42], which involve sensing techniques that require *per-home calibration*. Moreover, because we rely on a continuous noise signature, we can identify devices that do not generate transients such as those with soft switches or solid-state relays and transient suppressors.

In the most general sense, electrical noise on a power line refers to any signal other than those intended for delivering power in the home. In the US, power is delivered at 120V at 60 Hz, however one can find signals at many other frequencies on the power line albeit at very small voltages. In normal operation, such noise does not interfere with the operation of electrical appliances. As described in more detail in this section, many modern consumer electronic devices themselves produce noise as a result of their operation.

3.1.1 Theory of Operation

The electrical noise present on a power line when a device is operational is called conducted electro-magnetic interference (EMI), which can be classified into two types: transient and continuous. Transient noise is characterized by the short duration for which it can be observed, generally few tens of nanoseconds to a few milliseconds. Continuous noise on the other hand can be observed for as long as the device is operational. Both transient and continuous noise can either be concentrated within a narrow frequency band or spread over a wider bandwidth (also called broadband noise). A compact fluorescent light bulb (CFL) is an example of a device that generates continuous noise, which is conducted over the power line due to its physical contact with the power line. Since a home's electrical distribution system is interconnected in parallel at the home's circuit breaker panel, conducted EMI

propagates widely from a given device throughout the electrical infrastructure.

Continuous noise is usually intrinsic to the device's operation and internal electronics. Appliances like grinders, fans and hair dryers that make use of a motor create voltage noise synchronous to the frequency of AC power (60 Hz in the USA) and its harmonics (120Hz, 180Hz, etc.) due to the continuous making and breaking of electrical contact by the motor bushes. In contrast, modern SMPS based electronic devices generate noise that is synchronous to their power supply's internal oscillator.

In contrast to traditional linear power regulation, a SMPS does not dissipate excess power as heat, but instead stores energy in an inductance and switches this stored energy in from the line and out to the load as required, thus wasting much less energy. The key to a SMPSs smaller size and efficiency is its use of a power transistor to switch the stored energy at a high frequency, also known as the switching frequency. The switching frequency is much higher than the 60Hz AC line frequency because at higher frequencies the inductors or transformers required are much smaller [47]. Typical SMPS operate at tens to hundreds of KHz. The switching waveform is adjusted to match the power requirements of the appliance it is powering.

A CFLs power supply employs the same fundamental switching mechanism to generate high voltages necessary to power the lamp. The switching action, which is the cornerstone of a SMPSs operating principle, generates a large amount of EMI centered in frequency around the switching frequency. This phenomenon can be understood by modeling a simple DC-DC SMPS circuit that uses the same fundamental switching topology (See Figure 3.1).

The large inductor L_PowerLine mimics the power line inductance. The SMPS is plugged into the power line. To measure the conducted EMI, we place a voltage probe V on the power line, which is analogous to having the single sensor plugged into the power line with a SMPS based device operational somewhere else. The switching frequency Fc for the model is governed by the PER (period) parameter of the V_Switching component. We arbitrarily set it to 10 KHz. Figure 3.1 (right) shows a frequency domain plot of the noise at probe, which clearly shows that the power supply emits EMI, which is conducted over to the power line and is most prominent at the switching frequency fc (10 KHz here) and its harmonics. This is the same behavior that we observe when a SMPS based appliance is turned on in

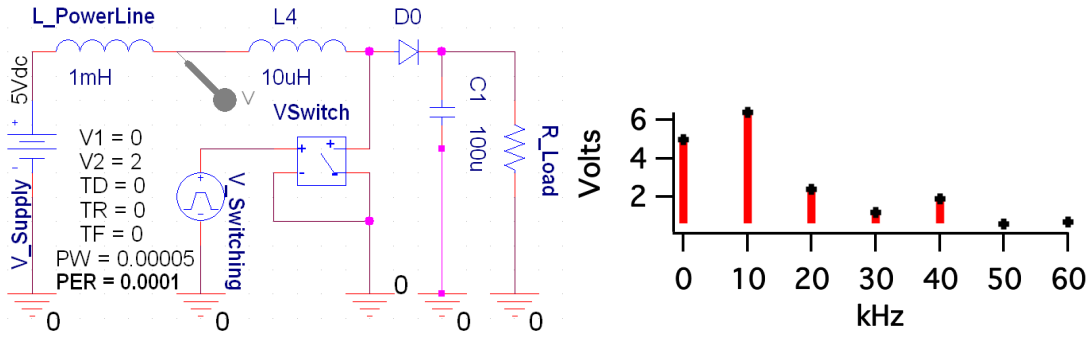

Figure 3.1: (*left*) Circuit model of a SMPS with placement of the voltage probe. (*right*) Frequency domain analysis at the voltage from probe showing EMI at 10 KHz.

the home.

In the US, the Federal Communications Commission (FCC) sets rules (47CFR part 15/18 Consumer Emission Limits) for any device that connects to the power line, which dictates the maximum amount of EMI a device can conduct back onto the power line. This limit is 66 dBuV for frequency range between 150 KHz to 500 KHz, which is nearly -40dBm across a 50 ohm load. The ElectriSense data acquisition system is sensitive enough to capture noise from -100 dBm to -10 dBm across a frequency range of 36KHz 500KHz.

Figure 3.2 shows a frequency domain waterfall plot showing appliances being turned on and off. As is evident from the graph, when the device is turned on we see a narrowband continuous noise signature that lasts for the duration of the devices operation. Also note that the noise center is strongest in intensity and then extends to lower and higher frequencies with decaying intensity, which can loosely be modeled with a Gaussian function having its mean at the switching frequency. This behavior can be attributed to the error tolerance of the components that make up the switching circuit core, as well as the characteristics of the power supply's load. If all the components were ideal, we would see a single narrow signal peak at the switching frequency. The error tolerance of SMPS components also allows for distinction between otherwise identical devices, such as a variety of units of the same model of CFL bulb. Finally, the power line itself can be thought of as a transfer function (difference in the inductance between the sensing source and the appliance) and provide

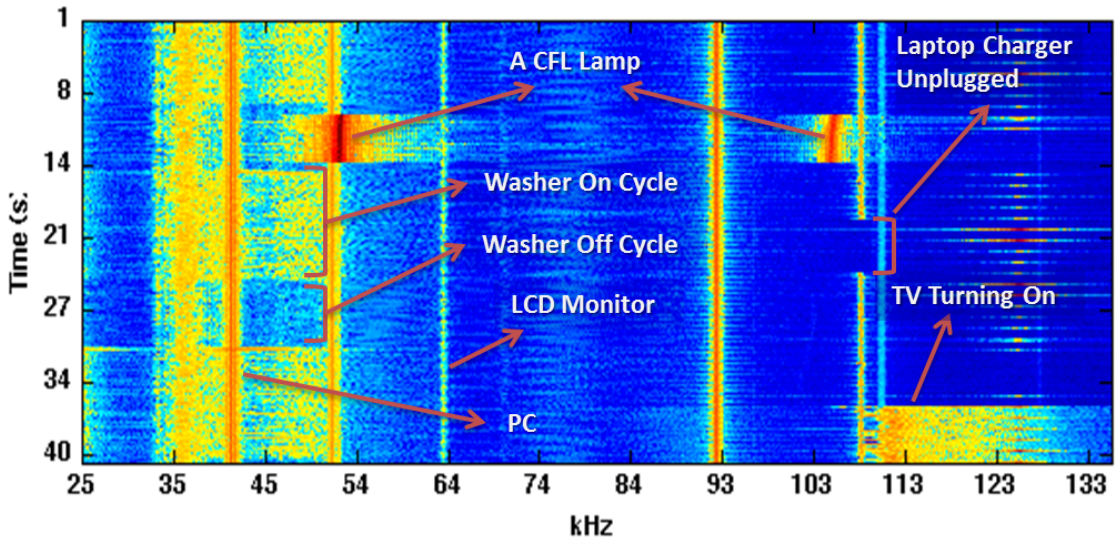

Figure 3.2: Continuous voltage noise signatures of various devices during various periods of operation. Colors indicate amplitude at each frequency with blue to red as low to high. Observe that the EMI is observed for the duration of an appliance's operation.

additional discrimination among multiple similar devices.

Dimmers also produce continuous noise due to the triggering of their internal triac switches, which can be used to detect and identify incandescent loads they control. In contrast to the narrowband noise produced by SMPS, a dimmer produces broadband noise spanning hundreds of KHz, which could be modeled as a Gaussian having a very large variance.

3.1.2 Prototype System Implementation

The prototype system consists of a single custom Power Line Interface (PLI) plug-in module that can be plugged into any ordinary electrical outlet in the home (Figure 3.3). The output of the plug-in module is connected to a high speed data acquisition system based on the Universal Software Radio Peripheral (USRP) that digitizes the analog signal from the plug-in module and streams it over a USB connection to a data collection and analysis PC running GNU Radio, which in real-time samples and conditions the incoming signal.

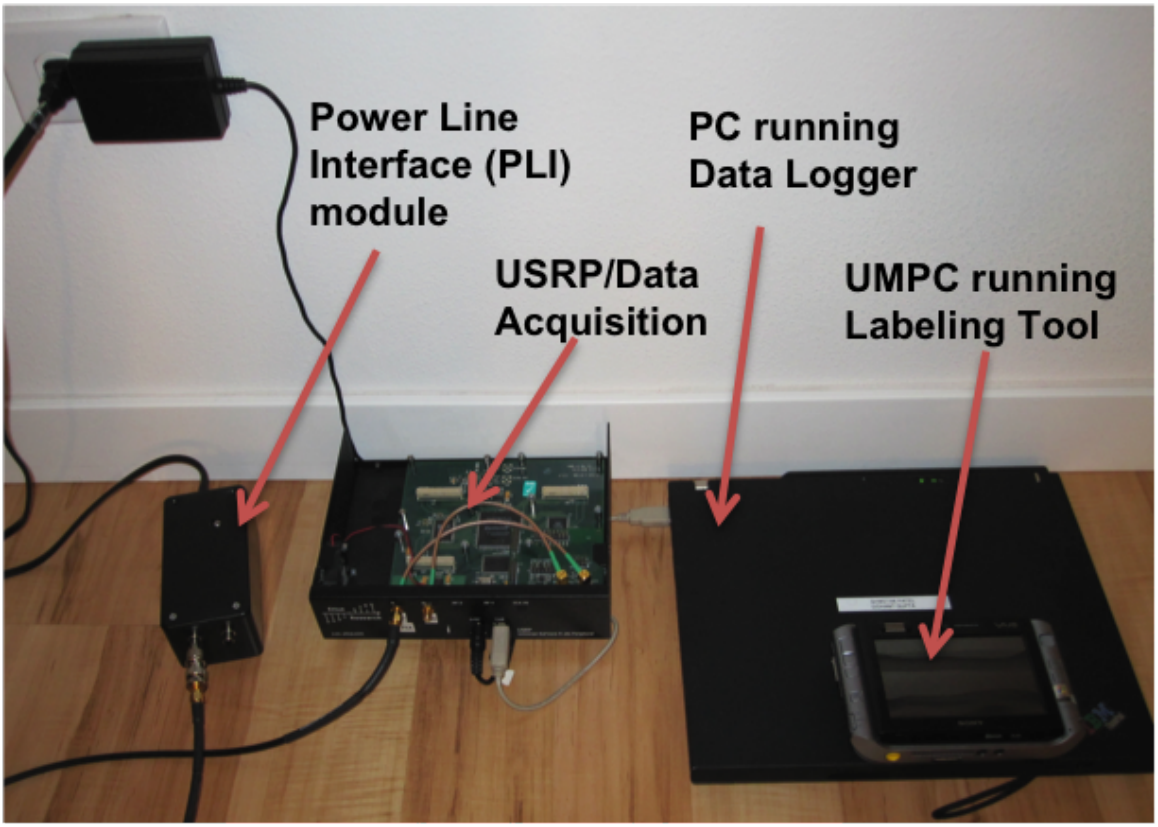

Figure 3.3: Prototype system consists of a single plug-in smodule, acquisition hardware and the supporting software

The ElectriSense algorithms then watch for an event and extract features that are used to identify and classify the device causing the event.

Though I tested the prototype system on a 120V, 60 Hz electrical infrastructure, this approach can easily be applied to electrical infrastructure utilizing different frequency and voltage rating with little change to the hardware and no change to the software. For homes that have split phase wiring (i.e. two 120 V branches that are 180-degrees out of phase), the coupling between the two phases allows for monitoring at a single location and capture events on both.

In the following sections, I describe the various components of our prototype system (Figure 3.4) and present in detail the hardware design and software algorithms used for

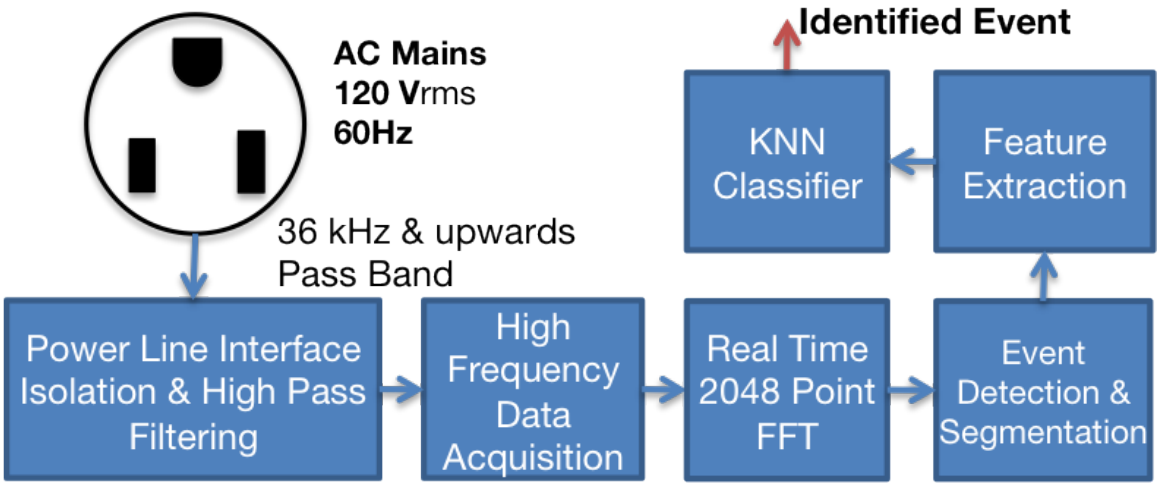

Figure 3.4: Block diagram of major components of our system.

event detection, feature extraction and classification.

3.1.3 Proof of Concept Hardware

To capture and analyze the electrical noise on the power line, I built a custom power line interface (PLI) module for the analog frontend. It is necessary for this module to filter out the AC line frequency (60 Hz in the U.S.) so that the spectrum analyzer or any analog-to-digital device is not overloaded by the strong 60Hz frequency component. The PLI consists of a high pass filter, which has an essentially flat frequency response from 50 KHz to 30 Mhz. The 3 dB corner is at 36.7 KHz resulting in a wide enough band to look at the complete range of any conducted EMI. Figure 3.5 shows the schematic of the PLI. The filter design also includes a 10 dB attenuator so that a constant 50-ohm load is presented at the input of the data acquisition hardware, irrespective of the signal frequency or the AC line conditions. For safety and isolation from the line voltage, high voltage capacitors are required. It should be noted that the polarity shown should be strictly followed, i.e., the line and neutral lines should not be connected in reverse and the isolation capacitors should be of AC-line rated polyester film type for safety.

The filtered signal is then fed into the USRP, which acts as a general purpose analog to

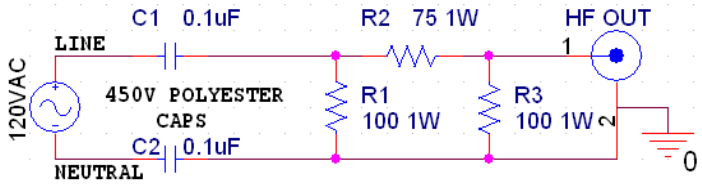

Figure 3.5: Schematic of the plug-in power line interface module.

digital converter sampling at a rate of 1 MHz set through the software. The digitized signal is streamed from the USRP to our processing software over USB.

3.1.4 Signal Processing, Event Detection & Feature Extraction Algorithm

The incoming time domain signal stream from the USRP is buffered as 2048-point vectors and FFTs of these are computed to obtain the frequency domain signal. The 2048 points are spread equally over the spectral width of 500 KHz, which yields a resolution of 244 Hz per FFT bin. The FFT vector or frequency vector is computed 244 times per second, which is then fed into our event detection and extraction software.

I found that most SMPS devices generate noise peaks that are 8 dB to 60 dB above the baseline. The baseline noise in a home varies unpredictably between -90 dBm to -70 dBm across the entire spectrum with a period of a few Hertz. Since the variability of the baseline noise is high, we must average the incoming frequency vector over time to obtain a stable baseline. We use a sliding window average with a window size of 25. Using a window that is too small results in an increase in the false positives whereas a large window size increases the shortest duration between near simultaneous events that is needed for the system to detect them as separate events.

When the system first starts, it computes an average of 25 frequency vectors and stores it as the baseline noise signature (Figure 3.6 *left*) Thereafter, a new window is calculated every 25 frequency vectors, and a difference vector with the baseline noise signature is computed. When a device is actuated and new noise in the frequency domain is introduced to the power line (Figure 3.6 *center*), the difference vector reflects this change, thus segmenting the event (Figure 3.6 *right*).

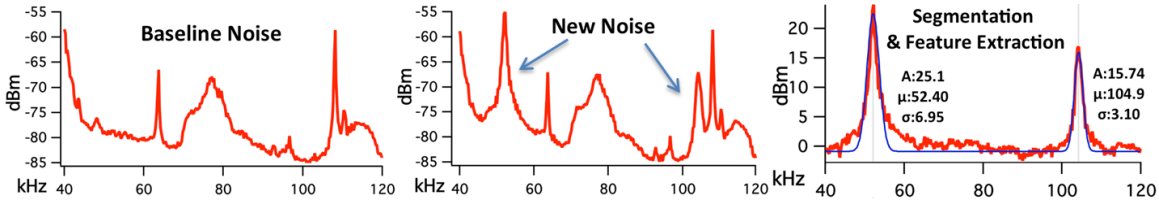

Figure 3.6: (*left*) Background noise observed on a particular power line. (*center*) A new device is turned on, producing EMI that introduces new signals to the power line. (*right*) After background subtraction the new signal features are extracted. The resulting Gaussian fit and its features amplitude (A), mean (μ) and variance (σ) are also shown.

The event detection algorithm scans the difference vector to find any values that are greater than a predefined threshold. It should be noted that this is a global threshold that is set once and works across different homes. Through deployments in homes, it was found that 8 dB above the noise baseline was a sufficient power threshold. Since the vectors in a window are averaged, if the window only partially overlaps with an event, the detection algorithm may still correctly detect it, but the difference vector will reflect a smaller magnitude. To mitigate this problem, when an event is detected, a new difference vector is calculated using the next window cycle. After the detection of an event, the baseline noise vector is updated appropriately so as to reflect the new noise floor of the power line.

The feature extraction algorithm finds peaks above the threshold using the difference vector and fits a Gaussian function to extract the mean, amplitude and variance parameters (Figure 3.6 *right*). The change in amplitude can be positive or negative depending on whether the device is turned on or off. These noise signatures are inverse of each other for opposite state transitions.

A feature vector for the suspected event is then created using the parameters of the center frequency, which is generally the global maximum frequency component. Other peaks may also be present as harmonics.

ID	*Style/Built/Remodeled*	*Size/Floors*	*No. of Test Devices*	*No. of Events*
H1	Apartment/1985/NA	750 sq. ft./ 1 flr.	10	135
H2	Single Family/2003/NA	3000 sq. ft./ 2 flrs.	15	203
H3	Single Family/1974/2009	1200 sq. ft./ 2 flrs. + basement	13	170
H4	Apartment/1910/NA.	450 sq. ft./ 1 flr.	7	108
H5	Single Family/1960/NA	1700 sq. ft./ 1 flr.	13	198
H6	Single Family/1926/2003	2800 sq. ft./ 2 flrs. + basement	20	404
***H7**	Apartment/2009/NA	657 sq. ft./ 1 flr.	16	1358

Table 3.1: A summary of the homes showing the style, size, number of appliances we tested and the number of events. *Longterm 6 month deployment.

3.1.5 Performance Evaluation

To validate the learning approach, I conducted experiments in seven different homes. I collected data from one house for a longer period, spanning six months and from multiple homes for a shorter period (spanning a single day). This allowed me to show the general applicability of ElectriSense to a diverse set of homes as well as the longterm temporal stability of the sensing solution. Table 3.1 shows the summary of the homes used in my evaluation.

Data Collection Procedure

The prototype system was packaged such that it could be rapidly setup in a home. The laptop, data acquisition hardware and a wireless router were pre-configured, connected properly, and setup on a rolling cart. For each home, I picked at random an available electrical outlet that had two sockets and plugged the PLI module into one and used the other to power our laptop and the USRP. The laptop and the USRP of course generate their own EMI but this noise is subtracted from the baseline. After the installation, I made a note of every appliance, electronic device, and light fixture that incorporated a switching power supply. This included incandescent lights that were driven by a dimmer switch in addition to any light fixtures with CFL bulbs. For dimmers, I only collected events at 0% and 100%

dim levels, because of the challenge of accurately and repeatedly setting intermediate dim levels manually. More analysis on dimmers is presented later in the discussion section. The collected labels were then fed into my ground truth labeling software. I then went through the home in two phases.

In the first phase, I actuated each appliance on and off five to six times individually to ensure that I captured an isolated signature for each appliance. Every time the system detected an event (i.e., when I turned on a device), the features were automatically extracted and sent to the labeling softwares queue running on an ultra-mobile Sony Vaio UX computer (UMPC). As events were being captured, I could label those events using the UMPC. This process was carried out for all of the devices in the home.

For second phase of data collection in each home, my goal was to collect data under a naturalistic setting allowing for overlapping and simultaneous events. To do so, (1) I asked the homeowner to perform certain activities (on an average three activities per home involving 2-4 devices each), such as watch TV, prepare a meal, etc. and (2) I actuated random appliances and/or logical groups of appliances such as turning on the DVD player, digital cable box, a gaming console and then a TV to simulate a TV watching experience through a universal remote. Events generated from these tasks were manually labeled.

In addition to the devices already found in each of the test homes, I collected data from a laptop power adapter, two CFL lamps and a camera charger that I took to each one of the deployment sites. To simulate a naturalistic use of each device, I moved them around in the home and plugged them into any available outlet. This allowed me to analyze the stability of these noise signatures across different homes.

The data collected by the labeling software was time stamped, labeled and stored in an XML database. Since I collected more features than required by the prototype, the XML database allows me to easily filter and parse the data with help of an XQuery to generate output data in a format that can be directly imported and processed by the Weka machine learning toolkit.

Long-Term Deployment

To test the temporal stability and long-term feasibility of the sensing technique, I de-

ployed the prototype system at one of the homes for over six months. Events were collected and labeled throughout this entire period manually using a slight variant of the labeling tool that was used in other homes. Every time an event was detected, the logging tool would put the extracted features in a queue and generate an audible beep. The home occupants were instructed to label the events with the electrical appliance they actuated using the labeling tool.

The labeling tool was designed to have a highly streamlined interface, requiring only the selection of device labels from an onscreen selection list. If the user did not label an event in the queue for more than six minutes, it was labeled as unknown and purged from the queue. This feature allowed the home occupants to handle any erroneously detected events and ignore events when the occupants were unable to attend to the labeling tool promptly. Though a convenient feature, this also meant that we lost labels on actual events that the occupant missed or decided to ignore. However, the intent of this experiment was to gather as many ground truth labels as possible for an extended period of time.

Data Collected

I collected a total of 2576 (Table 3.1) electrical events from seven homes. The largest number of events came from lighting, which tended to be either CFL- or dimmer-based. Most of the other detectable devices were common consumer electronic devices, such as LCD or LED TVs, gaming consoles, PCs, power adapters, etc.

Appliances such as dryers and electric stoves did not appear to generate events in some of our test homes. Generally such devices are large resistive loads and hence do not emit any high frequency noise. I did observe continuous noise events from a washer in H7. Out of all the test homes only H7 had a modern Energy Star compliant high efficiency washer that generated continuous noise, which was from its electronically controlled DC motor powered through a SMPS. Also, for most large appliances the use of a SMPS is considered negligible compared to the overall power consumption for the appliance.

Classifying Specific Devices within the Homes

To determine the classification accuracy of electrical device actuations in a home, I

Home	*10-fold CV (As is)*	*10-fold CV (Combined)*	*Minimal Training Set*
H1	96.29	96.29	93.46
H2	89.65	89.65	83.73
H3	96.47	96.47	93.52
H4	97.23	97.23	93.81
H5	85.35	92.42*	90.53
H6	84.41	91.83*	84.61
H7	92.85	92.85	85.13
Total %	**91.75**	**93.82**	**89.25**

Table 3.2: The performance using 10-fold cross validation and minimal training for classification for each home. The results of combining the same devices that are spatially near each other as a single event are also shown (*).

evaluated the classification approach using two different procedures.

In the first, I evaluated the performance of the KNN-based classifier using a 10-fold cross validation for each home. In the second evaluation, I used a minimal training set (a single training example for each device of interest) in order to simulate a more practical and real-world situation. Table 3.2 shows the overall accuracy for classifying devices within each of the seven homes. Using 10-fold cross validation, we observed an overall average accuracy of 91.75%.

Upon analyzing the confusion matrices for each of the homes, I found that for both H5 and H6, there was confusion between some of the lighting (see Figure 3.7 for the confusion matrix). From the data, I observed that the classifier correctly identified similar light fixtures (i.e, the same model and brand) that were located in different rooms. However, H5 and H6 had rooms where the same models of fluorescent light fixtures were installed spatially near each other (1-2 feet apart), which produced very similar noise signatures. Thus, these particular lights did not have sufficient differences nor were they far apart enough along the power line to differentiate between those lights.

Higher frequency resolution hardware may partially alleviate this problem, but with added costs. Also, for some applications, logically combining or clustering the lighting that

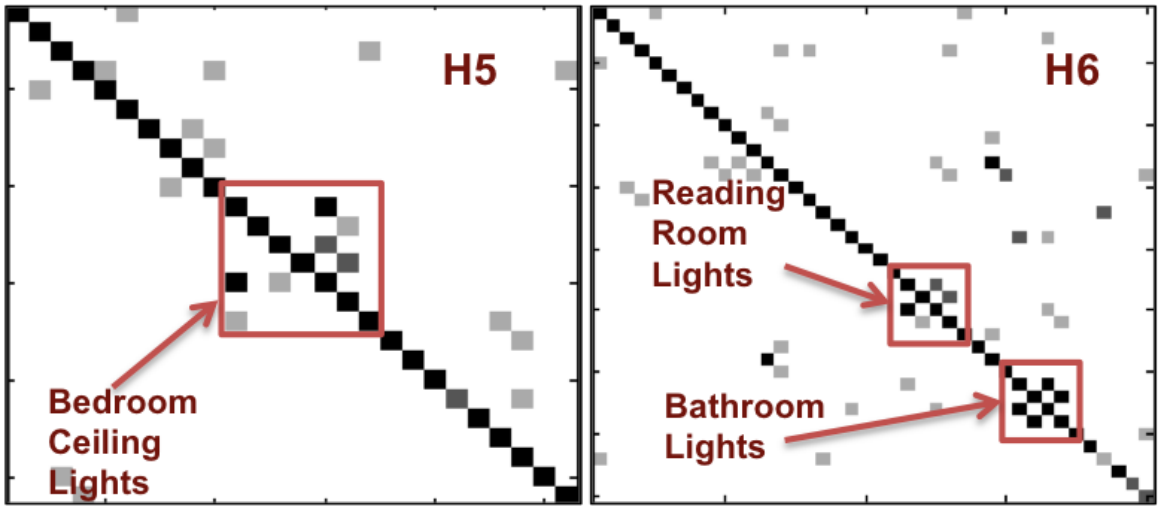

Figure 3.7: Visual confusion matrix highlighting misclassification due to physical proximity of similar fixtures in H5 and H6.

are spatially co-located might be acceptable. Using this latter approach, Table 3.2 shows the performance of the classification scheme after clustering the same model of lights or devices that are 1-2 feet apart from each other as a single event. This approach yields an increase in classification accuracy for H5 (92.4%) and H6 (91.8%) for an average accuracy of 93.82%.

The lower accuracy of H2 was due to classifier confusion between two devices of the same brand that were on the opposite phase of the installation point of the PLI. Part of the second floor exhibited very weak coupling between the electrical phases, which required us to plug in the PLI on that phase and thus causing some of the signatures to look very similar. I later discuss ways to address this problem, such as using two different PLIs, one on each of the two phases or installing a single PLI at a split-phase 240 V dryer outlet.

Since N-fold cross validation is generally optimistic and is not a true measure of the expected classification performance for a real world system, I performed a follow-up analysis using a minimal training set, i.e. a single event signature for each device to model and then apply it to a test set. For example, a homeowner would likely be only willing to provide a few training events for each appliance. Table 3.2 summarizes the classification accuracies

Device	*10-fold Cross Validation %*
Camera Charger	100
Laptop	87.5
23W CFL Lamp	100
12W CFL Lamp	100
Aggregate	**96.87**

Table 3.3: The performance of four devices across different homes using a 10-fold cross validation classification to evaluate stability of EMI signatures independent of home.

when using a minimal training data set. I found that this approach has an accuracy of 89.25%, which could be sufficient to enable a number of applications with high robustness.

Stability of Signatures Across Homes

Since no two homes have the exact same electrical infrastructure and can have drastically different baseline noise present on the power line, I examined the portability of noise signature across homes, which is a prerequisite for allowing signatures from one home to be applied to a similar device in another home.

I performed two experiments that together suggest the viability of applying a learned signal for a device in one home to a similar device in another. In the first, I show that the EMI signal of a device is independent of the home in which it is used, thus proving that the signal is intrinsic to the devices functioning. Second, the signal is consistent to within the variance limits for multiple, but similar devices, that is, the signals from the same brand and model of devices are similar.

My first experiment comprised of collecting a dataset for four preselected devices in each of the seven homes and studying the signal portability through classification results across homes. The average accuracy of the ten trials for a 10-fold cross validation test was 96.87%. For 3 of the 4 devices, the classification accuracy was 100%, strongly suggesting that these devices generated similar signals. This further bolsters the argument that the EMI produced is intrinsic to the device's operation and from a machine learning perspective, is minimally

affected by a home's electrical wiring. Of the four test devices, only one performed poorly. Table 3.3 summarizes these results.

I found that the laptops power adaptor was harder to identify because the extracted feature vector for this device in H6 looked different than the other homes. Upon closer inspection, the noise generated by the laptops adaptor had a harmonic peak that was very close in amplitude to the center frequency. This marginal difference caused the peak detection algorithm (which selects the strongest peak) to assign the first harmonic as the center frequency in H6, thus extracting a different feature vector. A simple approach I have developed to alleviate this problem is by employing an algorithm that classifies an event only if the distance between the event feature vector and the nearest neighbor in KNN is less than a certain threshold. If not so, the algorithm builds a new feature vector from the next strongest peak. Using this new method, classification accuracy for the laptop was 100%.

In my second experiment, I collected data for eight 20 Dell LCD monitors that were of the same model (from our research lab). This LCD model was also coincidently found in H5 and H7. I swapped out the EMI signature for one of the other 9 signatures obtained from a different home or building. For example, I trained the classifier using the signature from H7 and tested it in H5. This test ensured that if the signatures for any of the LCDs were different it would be misclassified as another device. All of the tests yielded 100% accuracy suggesting that similar make and model devices produce similar signature irrespective of the home or building they are in, implying the *feasibility of crowd sourcing to obtain device signatures.*

Temporal Stability of EMI Signatures

For any signature or fingerprint-based classification system, temporal stability is important. The classifier must perform well for months (or ideally, years) without requiring frequent re-training, which necessitates that the underlying features to remain stable over time and deviate substantially.

To show the stability of the noise signatures over time, I chose devices from our long-term deployment dataset, in particular those that met two criteria: (1) devices thatwere

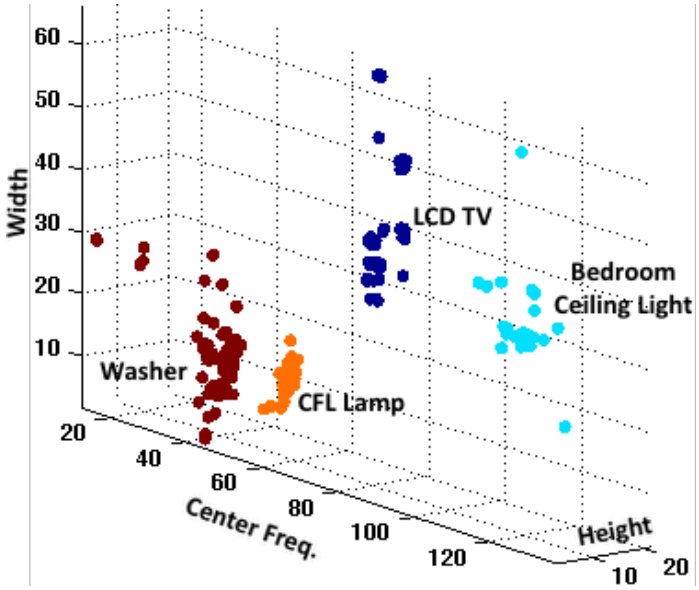

Figure 3.8: Variation of features over 6 months for four devices shown in the feature space. Note that no cluster overlaps.

fixed in their location during the duration of the deployment and (2) those that were not altered in any way, for example light bulbs that were not replaced. For these devices, I extracted randomly selected EMI signature vectors spread over the period of the 6-month evaluation.

Figure 3.8 shows the temporal stability or variation of the signatures over time for four randomly chosen devices by visualizing the feature vectors in the feature space. We observed that the long-term temporal variation was similar to what we observe in the short-term temporal variation in these devices and that none of the clusters overlapped.

To better understand how temporal variation effects the classification accuracy over time, I generated test sets for each device consisting of all events that happened more than one week prior to the last day of the deployment and a training set consisting of events from all devices that happened in the last week. This setup ensured that, if the EMI signatures in the test set deviated more than the distance between the device cluster that the classifier had computed, it would result in misclassifications. I observed 100% accuracy with KNN

classifier, which strongly indicate that the devices' EMI is largely stable over a long period of time. It is important to note that this long-term experiment was straddling the summer and winter seasons.

3.1.6 Discussion and Further Insights

The prototype system deployments and performance results from datasets collected in real-world homes strongly suggested that using EMI for electrical event detection is a promising approach. In this section, I provide additional detail and insights that can shed some light on improving the overall approach as well the limitations and challenges I uncovered.

Multiple Instances of Similar Electrical Devices

Having a number of similar devices is a common occurrence in a home, such as having multiple TVs or, more commonly, lights that all use the same brand of CFL bulbs. This can cause problems, especially if similar devices cannot be grouped into a single group. For example, grouping two ceiling lights in a bedroom may be acceptable, but grouping lights that are in different rooms or floors may not be. There are two potential resolutions to this.

First, in many instances I noticed that the tolerances in components that make up the switching circuitry of a device can introduce enough variability in switching frequency that the mean of the Gaussian fits observed on the power line are also shifted. Figure 3.9 shows a subset of the spectrum observed by our system showing the spectra of the noise generated by four CFL lamps of the same model that were purchased as a pack of four, thus ensuring that they came from the same manufacturing batch. Note that the spectra do not overlap even among the same batch of CFLs.

The prototype hardware is able to discern these subtle features only when observed in isolation, for instance, a line isolation transformer was used to create a noise free power line for this particular experiment. With higher ADC resolution and a larger FFT, this shortcoming may be overcome. Thus, increased resolution may give us better differentiability. The deployable commercial version of ElectriSense that I briefly describe in the next section was built keeping these finding in mind.

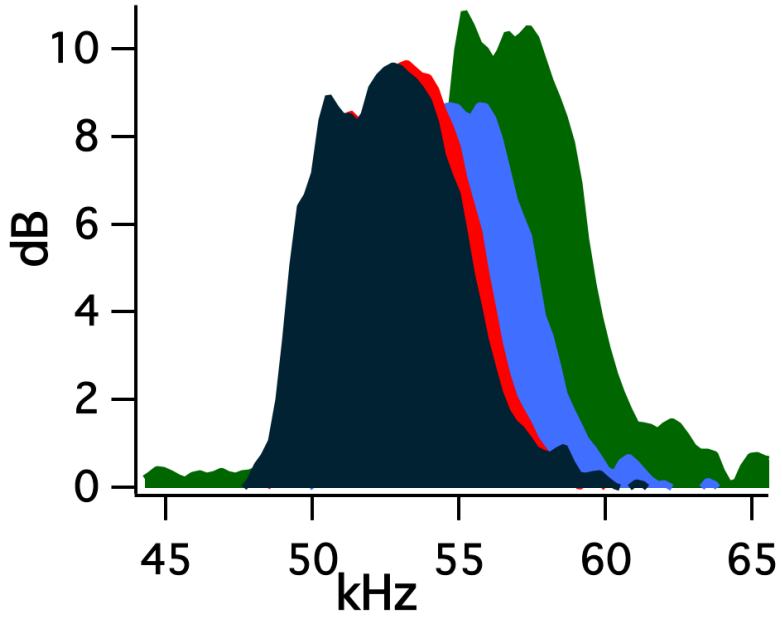

Figure 3.9: Small, but discernable variation in the mean of the EMI peaks for four of the same model and brand Compact Fluorescent Lamps.

Second, as the conducted EMI travels through the power line, it is affected in several ways, but most prominently the signal is attenuated as a function of the line inductance between the source of noise and the point of sensing. Thus, two identical devices generating identical EMI may look different at the sensing source depending on where the devices are attached along the power line, which I observed in our in-home experiments.

To confirm this, I plugged a device in two different locations in a home and logged the raw spectrum data as sensed by the system. Figure 3.10 shows a spectrum of a small section of this data. The difference in amplitude can be used to differentiate between similar devices located in various parts of the home. This suggests that we could determine the number of fixed devices present in the home (i.e., the number of CFL lights in the house or the number of a particular type of TV).

It should be observed that only the amplitude varies prominently as a function of the

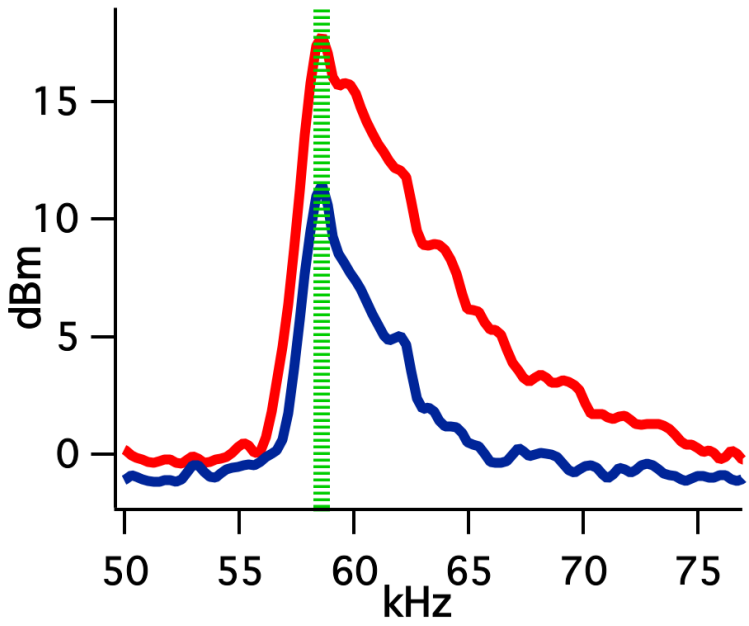

Figure 3.10: Same CFL lamp plugged into different regions of the home producing EMI amplitude variations.

position of the device in the home. Other features such as the mean and variance of the Gaussian remain relatively stable (as also indicated by the long term stability study discussed previously). It is these stable subset features that allow for correctly identify a mobile appliance, such as a laptop adapter despite changes in its position in a home. It also allows reuse of signature from one home to another without re-training for similar devices. Thus, the stable subset features can be used to identify which device (CFL, TV etc.) and the amplitude can be used to resolve the instance (TV1, TV2 etc.) in a particular home.

The probability of two different types of devices having the same Gaussian fit mean and variance is small. For independent EMI signals, it is 1/(FFT Size) if we only consider the probability of two devices having the same mean frequency. In practice, since I use not only the location, but also parameters of a fitted Gaussian, the true probability is likely much lower.

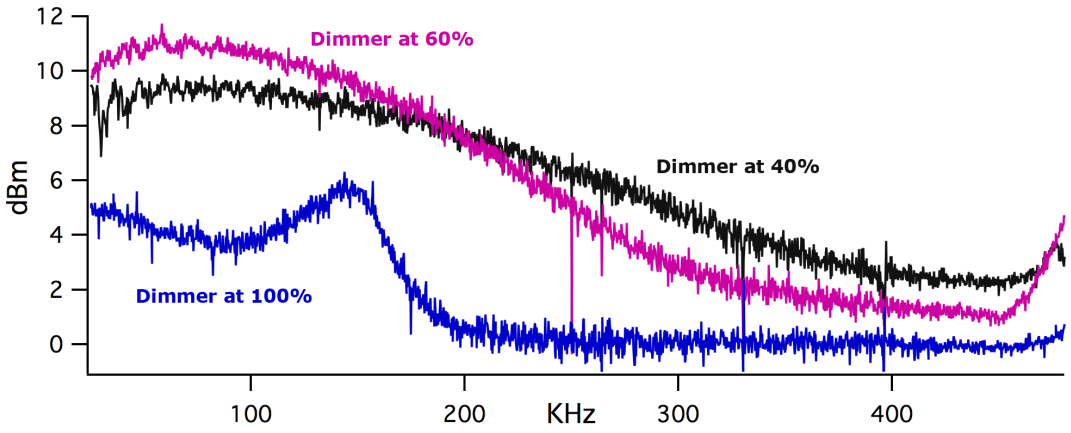

Figure 3.11: Band limited EMI generated by a dimmer shown at various dim levels.

Analysing Dimmers and Triac Based Devices

Unlike SMPS based devices, dimmers generate a wide band signal from their internal triac switches. Since the prototype feature extraction algorithms are designed around a Gaussian fit, I model the broadband noise emitted by a dimmer with a Gaussian distribution, even though a band-limited uniform distribution would be more appropriate. Thus, I found that for dimmer controlled devices the observed Gaussian fits with very large variances. Figure 3.11 shows the EMI signal generated by a dimmer at various dim levels.

The difference in the signature generated at various dim levels opens up the possibility to identify as well as infer the dim of such devices device. Since it is not possible to train at different levels, one possible resolution could be to build a model for the noise characteristic and how it is affected by the dim level or conduction angle of the triac.

Additional Features Beyond Gaussian Parameters

The performance and robustness of a classification algorithm is only as good as the features being used, so it is desirable to explore and extract more advanced features from the underlying signals. During my experimentation I observed multiple potentially useful

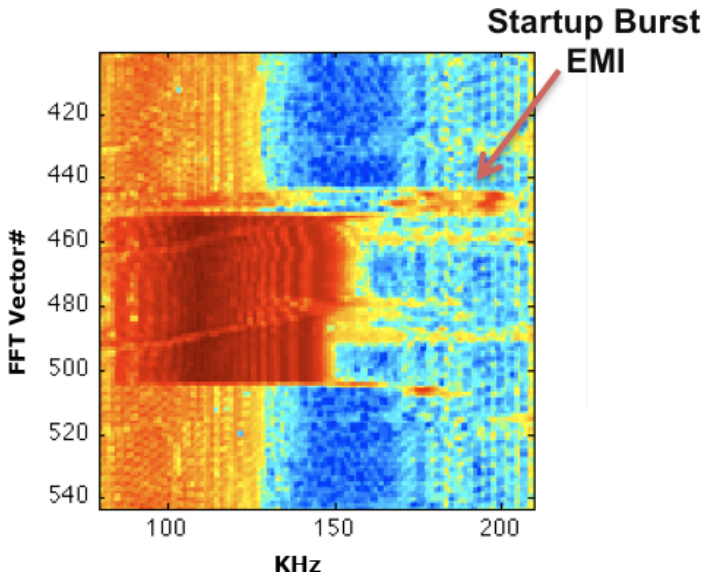

Figure 3.12: Startup burst of EMI signal generated by CFL lamps on ignition.

features that could be used in future classification work. Several devices produce a characteristic EMI pattern that lasts for a short time when they are switched on, i.e. they produce narrow band transient EMI while the SMPS is starting up. These transients are lower frequency than those described by Patel *et al.* [41]. Figure 3.12 shows a short burst of EMI that most CFL lamps produce when first powered up. The short burst comes from the ignition circuitry in a CFL, which is required to warm the lamp up for operation. Similarly, other devices such as modern TVs and DVD players that have multiple power supplies and supporting electronics also produce start up noise signatures that may provide additional information for determining the category of the devices.

In the current implementation, the algorithm design assumes that the mean or location of EMI peaks do not change when the device is operational. This may not be true for certain devices. I have come across a particular brand of a LCD TVs, where the switching frequency of its power supply is a function of the screen brightness, thus causing the mean of EMI peaks to shift as the content on the screen changes. The current algorithm will report a new event every time such a change happens, which could overwhelm the system. An alternative approach, as discussed in the next chapter is to lock onto the time varying

noise peak and extract temporal features or templates.

Certain motor based devices such as a washing machine or a dishwasher generate low frequency periodic noise patterns in their motor controllers, which could be used to identify the device and its state. For example, in H7, I observed that the clothes washer in its wash cycle produced intermittent noise at roughly 0.1 Hz in contrast to constant noise while in its spin cycle. Such features can be used to build a finite state machine or a statistical model for more detailed classification.

Thus, as domain experts we can both discover and predict new features given the underlying physical properties and electrical circuitry of a particular electrical appliance. For instance, energy efficient LCD monitors by definition are energy efficient because they *adapt* their power consumption to how the LCD is being used. In other words, we expect the EMI from a LCD TV to vary over time. In chapter 4 I explore this further and demonstrate that many of the modern TVs exhibit this, which opens a whole new dimension of features pertaining to appliance's operating state.

Simultaneous Electrical Events:

The current implementation can detect near simultaneous events as close as 102 milliseconds (ms), that is, two events that occur more than 102 ms apart can successfully be detected as separate events. This is based on the sampling frequency and averaging window size. If simultaneous events happen in shorter than 102 ms duration, they are detected as a single event, and the features extracted are a collection of features from multiple devices. Thus, the underlying features remain intact, but are reported as one event. A new classification approach may be able to separate out these compounded features to identify individual devices, however the difference vector approach investigated here is inherently limited to resolving between near simultaneous events defined by the sampling rate and the averaging window size.

Phase Coupling:

During my in-home deployments, I observed strong signal coupling or crosstalk across the two electrical phases in most homes, including larger homes like H5 and H6. However,

for H2 I observed that for some parts of the home, which were on the opposite phase from the sensor it was very difficult to detect events. Either installing two PLIs (one on each phase) or installing the PLI to an available 240 V outlet where both phases are present (typically for a dryer) can address this problem. Installing two PLIs would allows the system to capture events from both phases, however this also increases the chances of similar looking signatures for two similar devices. This can trivially be addressed by knowing which of the two PLIs detected the event.

Reducing Training Effort:

Since ElectriSense employs a fingerprinting based approach for device identification, it requires a training process to learn the parameters of various devices in a home. Like any such approach, this is generally accomplished by having the user go through actuating each device at least once, i.e., generate a training dataset to train the machine learning classifier.

Fortunately, my approach benefits from the portability of EMI signatures across homes, which allows end users to share their device signatures, for instance, through crowd sourcing. In other words, an end user may be able to access an online database of popular consumer electronic device signatures submitted by other users who have gone through a calibration process. Additionally, since the EMI that I sense is an engineered signal (it is a result of a particular circuit design), it may be possible to generate such signature databases without physically actuating a device or having physical access to it. This can be realized by mining information from a devices FCC compliance report, which lists the frequency and raw magnitude in dBuV for various noise peaks it emits. An example compliance report for a PC power supply is shown here [16]. Mining the datasheets of the internal integrated circuits and oscillators found in consumer electronic devices are other sources for this information.

3.2 Implications for Energy Disaggregation Algorithms

In this section I briefly highlight the complementary and exclusive features that can be extracted to sense appliance usage using EMI. Though the merits of EMI based features over power-usage based features were discussed in detail in Chapter 2, it is worth noting that

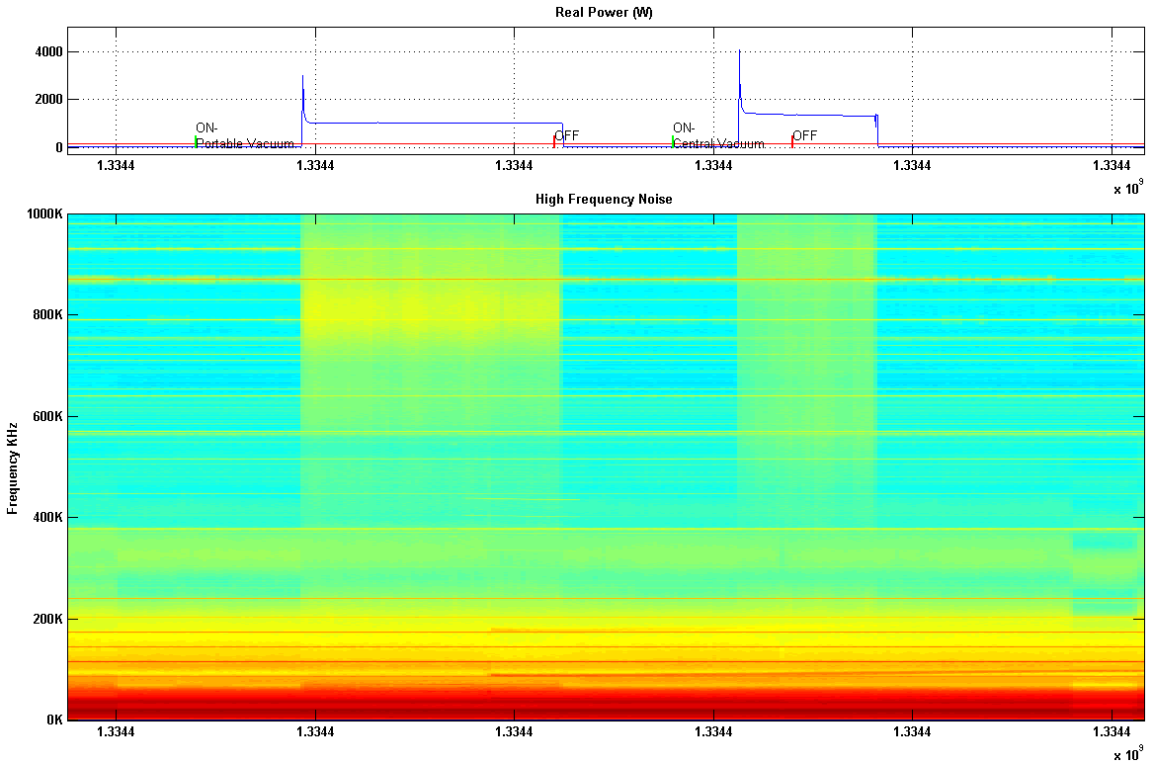

Figure 3.13: EMI sensing complements prior power based appliance usage detection. Shown here are two vacuum cleaners being operated briefly. Note the change in power as well as EMI that can be used together for a more robust event detection and feature extraction than either one independently.

not only does EMI based sensing add new features allowing for more robust appliance usage sensing, but also adds ability to sense appliances and electrical activity not previously practical. For instance, EMI sensing is independent of the power usage of an appliance. Thus, unlike power-based approaches that often threshold at about 150-200 Watts to avoid false positives, appliances that consume arbitrarily small power can also be detected. This has an immediate and direct implication for energy disaggregation algorithms as they fundamentally depend upon being able to resolve appliance state for as many appliances as possible with a high degree of accuracy.

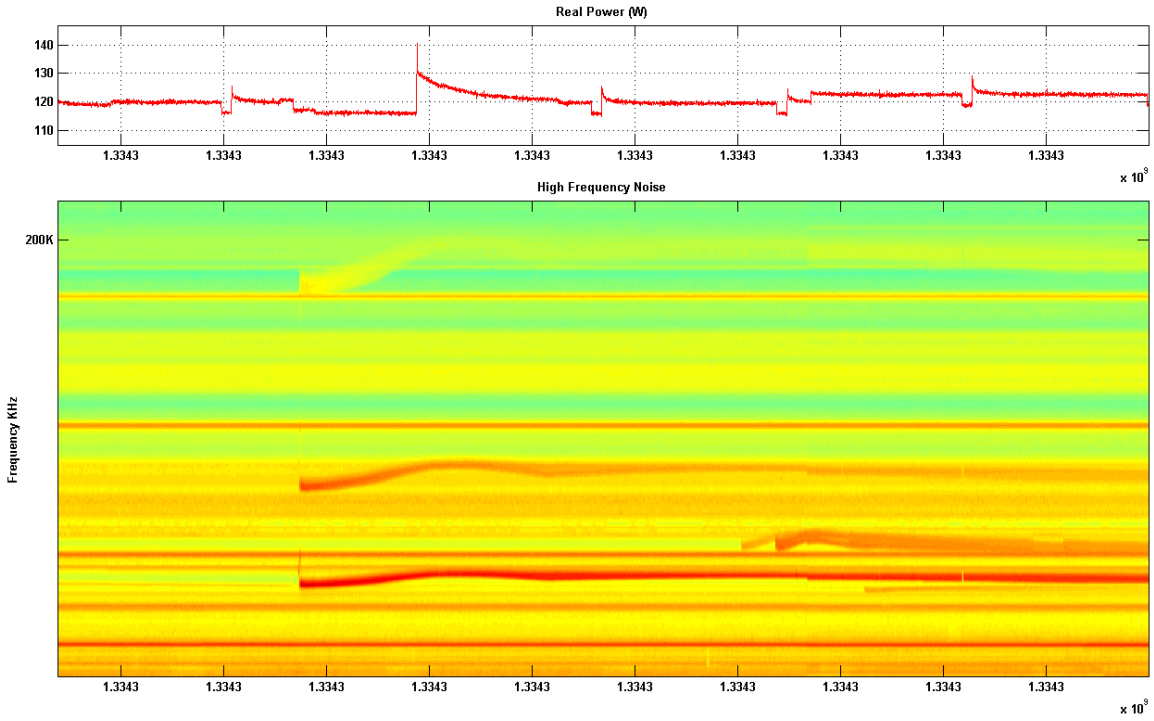

Figure 3.14: A strong EMI signal can be observed despite the appliance consuming only 8.2 Watts. Reliably detecting such small power changes using power-based disaggregation algorithms proves to be a challenging task in practice. On the other hand, the strong EMI complements the prior approaches and makes electrical event detection robust. Also note that the device's power consumption changes over time (step downs and ups) reducing to 1 watt, but the EMI remains detectable.

3.2.1 Complementary EMI Features

Detecting presence or absence of EMI on the power line complements existing energy disaggregation approaches that detect a positive or negative step change in current or power to infer an appliance turning on or off respectively. Figure 3.13 shows an example of two different vacuum cleaners turning on and off. Though their power consumption profiles may be similar, complementary features extracted from EMI could be used to better disambiguate between the two. In addition, there is one key distinction: sensing presence or absence of

EMI is independent of the magnitude of the power and appliance consumes. Specifically, EMI sensing allows for detecting electrical events even if the step change in the power is small. That is, a 5W mobile phone adapter could produce EMI in magnitude similar to that of a 1500W PC. This allows for building a more robust electrical event detectors that is not fundamentally limited by detecting changes in power that are sufficiently higher than the variability of power commonly observed in a typical residence. This variability can be anywhere from 100W-200W, forcing algorithm designers to set a threshold high enough to mitigate false positives at the cost of being not able to detect appliances with lower power consumption. Figure 3.14 highlights how a strong EMI is observed for an appliances that causes a small 8W step change. Also note that the device's power consumption changes over time (step downs and ups) but the EMI remains detectable.

3.2.2 Exclusive EMI Features

While analyzing EMI from different appliances, I discovered that certain properties of the EMI strongly correlate with the mode or state the appliance is operating in. This phenomena is discussed in detail in Chapter 4, however from a machine learning perspective it implies features that are exclusive to EMI sensing. One feature characteristic of compact fluorescent lamps was discussed previously and is shown in Figure 3.12, that it produces a transient EMI burst at a frequency higher than the center frequency when it is turned on. In addition, I noticed that as the CFL warms up to a steady temperature, the center frequency drifts over time by as much as 0.5 KHz - 1 KHz (see Figure 3.15). Such device specific features give an insight into the operation of the device that goes beyond just sensing whether an appliance is on or off. Leveraging such appliance specific models and features can potentially make existing energy disaggregation algorithms more robust.

3.2.3 Combining ElectriSense with Real-time Power Consumption Data

ElectriSense by itself only provides information about which appliances in the home are being used, however for an end-to-end energy disaggregation system this information needs to be correlated with real-time power consumption of a home. Specifically, when ElectriSense

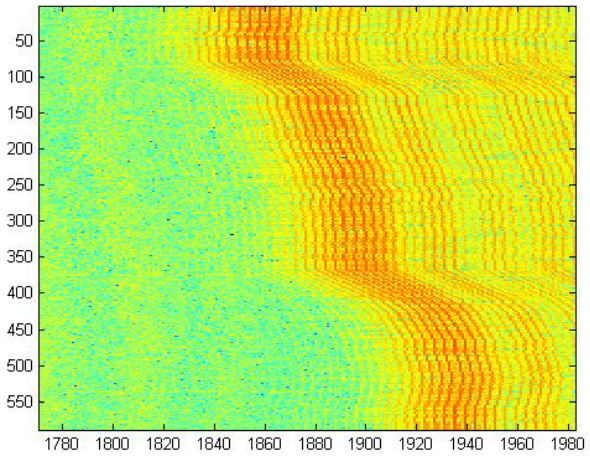

Figure 3.15: Center frequency of the EMI from a CFL drifts by 0.5-1KHz over time as the lamp warms up to finally reach a steady state. Sensing such appliance specific behavior allows extracting machine learning features that can make disaggregation algorithms more robust.

detects an appliance turning on, the change in power consumption could be used to track how much energy that particular appliance uses.

There are a number of commercial electrical sensors available, with the most popular and inexpensive being the current transformers (CTs). CTs uses a transformer based current sensor that is installed inside a breaker panel (see Figure 3.16 (left)) and requires hiring a trained electrician because it involves placing the sensor around the main electrical feed. Smart meters on the other hand are easy to install because the utility company installs them. However, this also means that not every home has a smart meter as the initiative to install smart meters is still in its infancy. Finally, smart meters typically only provide interval data (i.e., every 10-15 minutes), vary across manufactures, and are not currently easy for end users to obtain data from.

To alleviate these problems, I have designed an end user installable whole home power consumption sensor that *does not* require any electrical experience or professional installa-

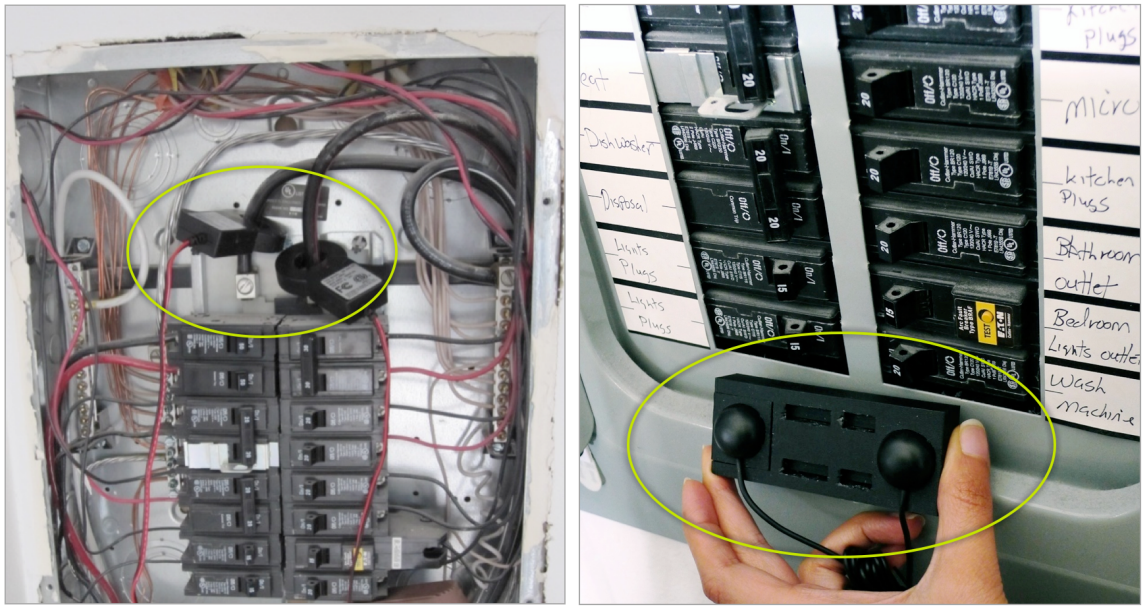

Figure 3.16: (left) Conventional real-time power monitoring requires current transformers to be installed *inside* the breaker panel. Such a task requires professional installation. (right) In contrast, Contact less Current Sensor (CCS) I developed measures real-time power consumption using a sensor that is installed on the *outside* of the panel and is end-user deployable.

tion (see Figure 3.16 (right)). This sensor provides real-time data, and can be used in any home where there is a breaker panel for the living space, which is required by the national electrical codes. This sensor, called the Contactless Current Sensor (CCS), senses the magnetic field induced by the 60 Hz AC current flowing through the conductors inside a breaker to precisely measure the power consumption. Since it uses magnetic fields, unlike sensors like the TED, there is no requirement to open the breaker panel for installation. Instead, the user simply attaches the sensor to the *outside* panel with adhesive.

More specifically, this current sensing approach involves computing the power consumption by inferring the current being drawn through the main lines coming in the home at the breaker panel. Since most homes in the U.S. have split, two-phase electrical service, we need to detect the current at both legs. Standards in breaker size and code guidelines provide

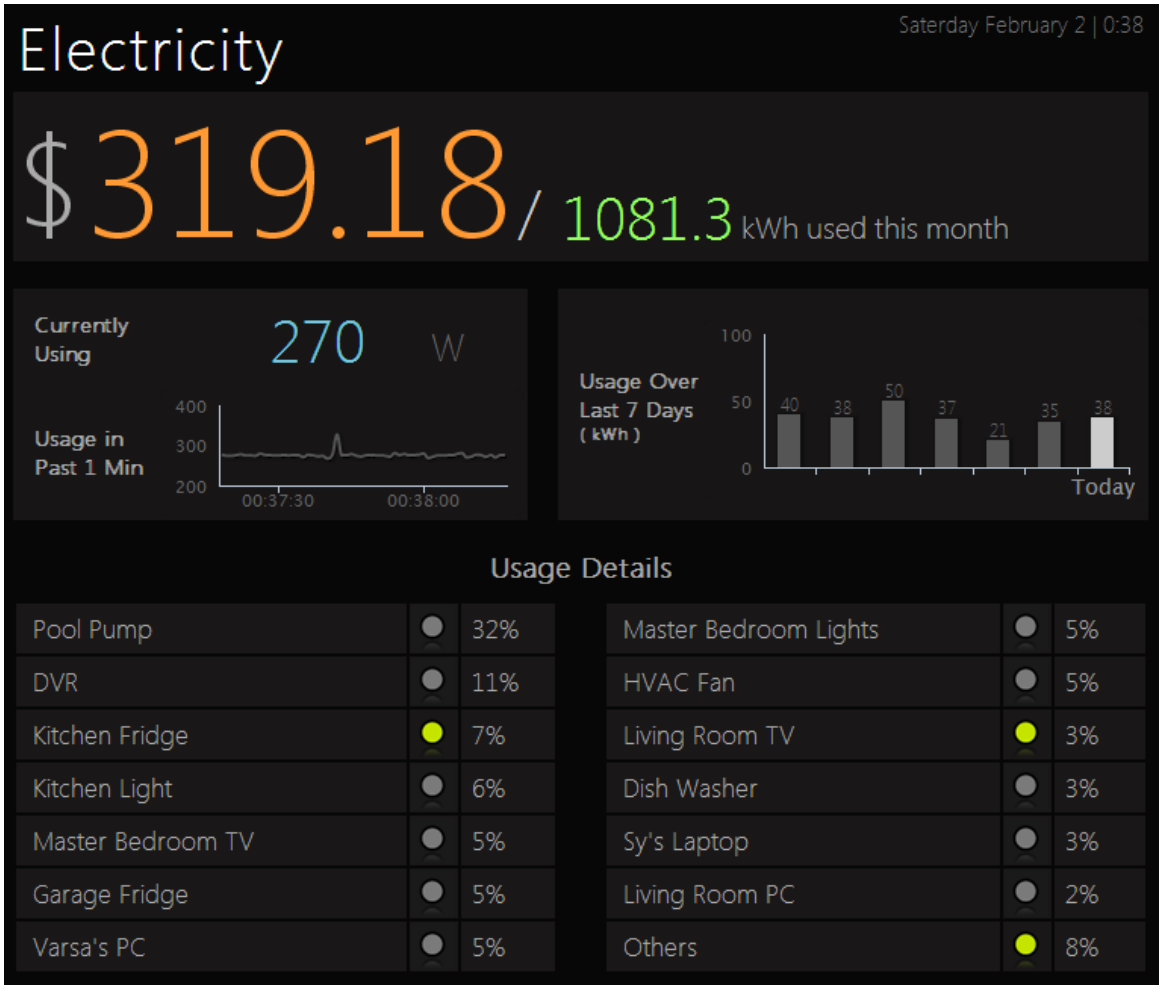

Figure 3.17: ElectriSense combined with Contactless Power Consumption sensor shows how much and where energy is being consumed in real-time.

some structure to how the electrical feeds connect to the lugs on the bus bar in the breaker panel. For example, the wires for the two electrical phases come in parallel to each other when connected to the bus bar. The field generated from the two legs or phases allows us to estimate the current flow through each leg separately, which radiates a few centimeters from the wire and even through the layer of sheet metal. Further details and performance evaluation of the Contactless Current Sensor can be found in my published work [39].

In a recent prototype system, I have combined the real time power consumption data

from the CCS with ElectriSense to assign power consumption to appliance usage. Figure 3.17 shows the output of this system, which not only reflects real time energy usage, but also disaggregates it by the various appliances in the home. Appliances being currently used are indicated with a 'green light' on the interface, along with what percentage it contributed the home's net energy usage.

3.3 Discussion & Overview of Contributions

Appliance level energy usage feedback systems can have a profound impact on informing, educating and motivating people to save energy. However, the sensing technology that can enable such systems have not yet been realized in a way that are accurate, easy to install, low cost and robust. In this chapter, I presented the next generation of electrical appliance usage sensing system and through real-world deployments showed that it is a practical and reliable technique. In doing so, I have significantly improved the current state-of-the-art in energy disaggregation by developing a novel sensing approach that leverages signals, or electrical noise produced as a result of a device's operation. This noise or Electromagnetic Interference can then be used as fingerprints to detect when a particular appliance is being used in the home.

When compared to the state-of-the-art in disaggregation algorithms and approaches, ElectriSense adds a new dimension to the feature space that was previously unexplored. It also offsets and complements fundamental limitations of prior power-based disaggregation approaches as detailed below:

- Current or power based approaches (like those based on using step change in power) often fail to resolve *compound events*. That is, two or more appliances being operated simultaneously poses a challenging problem since the power signatures from multiple appliances add up. Recent advances in use of HMMs have improved the detection, however the signal is inherently 1-dimensional. EMI based features from ElectriSense not only add an additional set of features but presents a much richer and higher resolution feature space. For instance, multiple appliances could be operating simultaneously, yet their EMI could be different. Section 3.1.6 discusses how the probability

of two signals overlapping is inversely proportional to the FFT size. It should be noted that the probability is even lower when all of the features, namely mean, variance and amplitude are taken into account.

- Similar to the point made above, the power step changes and harmonics of the same are inherently limited to 1-d feature space. This can be enhanced to some extent by using real and reactive power, however EMI sensing adds *multiple new dimensions* to the feature set. In particular, higher frequency resolution allows for resolving even a crowded EMI feature space, something not afforded by the 1-D real and reactive power feature spaces.

- Current or power based approaches have a practical limitation to the smallest change in power magnitude that triggers event detection. In other words, thresholds used to detect when an appliance has changed state has a practical lower limit gated by the normal variations of power observed in a typical home. This variation stems from variable load appliances such as cooking stoves, TVs, computers, heaters etc. As discussed in Chapter 2, typically this variation is 100-200W. Thus, it becomes difficult in practice to detect an appliance turning on/off or changing state, which results in a change smaller than 200W. In contrast, detecting electrical events using EMI sensing is *independent of the power a device consumes.* This is why ElectriSense can reliably detect appliances like a camera charger or a laptop charger that typically consume only 40-90W.

- Appliances that consume *variable power* also pose a challenge to reliable classification with features based solely on step power changes and/or its harmonics. For example, a hand drill or a food blender could be powered on in a number of different modes as well as depending on the task they may consume varying power. This not only makes it difficult to train a machine learning classifier due to the myriad states but also impacts classification accuracy. On the other hand features based on EMI are largely immune to varying power, making detection of such appliances more robust.

- Always-on appliances like power adapters, DVRs, security systems etc. are harder to detect using power step-change approaches because they may never turn on/off like other appliances. In the current ElectriSense prototype's algorithm, though this also poses a challenge for EMI sensing, however, by instead analyzing the number of static EMI signatures, an estimate of the number of background appliances and their type could be inferred. This opens up new opportunities to sensing always-on or more commonly the "vampire power" appliances.

Chapter 4

GOING BEYOND SENSING APPLIANCE ON/OFF EVENTS: USING EMI FOR OPERATING STATE ESTIMATION

In this chapter, I present an advancement of the EMI sensing technique developed in Chapter 3 that allows for sensing not only whether a particular appliance is on or off, but also estimating what state it is operating in by capitalizing on the unique temporal patterns in the EMI it generates as those states change. I investigated these temporal patterns by studying three classes of appliances in detail, namely compact fluorescent lamps, televisions and liquid crystal displays. The results and findings are presented as case studies in this chapter. Also discussed are preliminary observations on how estimating operating states may extend to other classes of appliances. Finally, implications of estimating an appliance's operating state on energy disaggregation is presented.

4.1 Overview of Operating State Estimation

I define *Operating States* as different modes that a particular appliance could operate in. For example, a washing machine with a wash, rinse and spin cycle can be categorized as having three operating states. The number of operating states depends on (1) the granularity with which these states are interpreted by a researcher, (2) on the intended application and (3) the appliance's capabilities. For instance, a finer granularity of states for the aforementioned washing machine could be the three sub-rinse cycles that make up the overall rinse operational state.

For the purposes of energy disaggregation, I have investigated sensing operating states that are user selectable or are a result of direct user interaction with an appliance. That is, these are states that an appliance goes through depending on how the user operates it. In our washing machine example, the sub-rinse cycles would not be considered distinct operating states unless they are user selectable, however the rinse cycle itself would be considered because the user can selectively enable or disable it. For certain appliances there

may not be any direct selectable state, however distinct states may exist as a result of user's direct interaction with the appliance. Like in the case of a PC, as a user interacts with it, the resulting CPU load could be categorized as low, medium or high.

Sensing such operating states could be vital to accurate energy disaggregation. As detailed later in this chapter, I found that the EMI from laptops and PCs vary over time as a function of the CPU load. Tracking this change could enable energy disaggregation algorithms to adapt to varying loads and be more robust. This is especially useful when power variations from an appliance is similar to or lower than the average variations of a home under normal circumstances. In addition, for compound events when multiple appliances are operational simultaneously with varying power consumption profiles, it is relatively straightforward to track variations in EMI and resolve the power profiles for individual appliances.

I contend that being able to sense operating states of appliances in a non-intrusive manner not only adds new machine learning features to the set of those already discussed, but also opens up an immense opportunity to better understand and model appliance usage in an environment. In addition to more robust energy disaggregation models it can also inform and educate home owners of potential power-saving modes or alternative appliance-usage scenarios. For example, an operating state aware energy monitoring system could suggest to the user that they decrease their TV's brightness by 10%, which saves energy but has little effect on the TV viewing experience. This is only possible if the brightness level could be sensed (discussed in section 4.2.2). Similarly, the system could suggest that the user operate their clothes dryer in a different mode depending on the hyper-local weather measurements and in-house thermostat settings.

Prior art in sensing appliance's operating state through NIALM approaches has been limited to predicting appliance failures and degradation [5, 14, 48, 49]. In particular, this body of work focuses on NILM-based analysis that can identify poor appliance performance based on sudden changes in power profiles and AC harmonic content. More recently, direct sensing approaches like [22] that make use of an EMF sensor placed close to a device shows potential for being able to sense an appliance's operating state, however such an approach will require one sensor per appliance. To the best of my knowledge, the non-intrusive

approach presented in this chapter is the first to sense fine-grained (like wash cycle for a clothes washing machine) appliance operating states.

The key insight behind inferring an appliance's operating state through its EMI is that *changes in how an appliance operates manifests itself as variations in the EMI it produces.* Though the exact process responsible for such a change varies from one class of appliances to another, the underlying effect is on the oscillator that is responsible for producing the EMI.

4.2 Case Studies of Common Appliances

The observed EMI variations and the underlying operating state that they map to could vary from one appliance to another. This is because depending on the function of the appliance, how its power varies in relation to those function, its SMPS design and overall circuitry, the EMI produced could correlate differently to different operating states. Though specialized models are needed for each class of appliance to map EMI variation to its various operating states, the way a device's circuit behaves under varying conditions is well understood. This limits the types of EMI variability and mappings of the same to the internal states.

In order to understand these mappings and develop a generalization I chose to investigate three classes of appliances in detail: compact fluorescent lamps, televisions and LCDs. In this section I discuss these classes of appliances as case studies where EMI variation was used to infer internal operating states and then applied to human activity and behavior sensing. The choice of these appliances was motivated by the fact that they represent one of the most ubiquitous consumer electronics found in modern households. In addition, while CFLs represent a simple device with seemingly no operating states, TVs on the other hand have complex internal states. This allowed me to investigate the electrical properties that govern the EMI variations using CFLs and apply those to TVs and LCD monitors. The findings from investigating these appliance classes has a profound impact on understanding how the general approach of leveraging variations in EMI could be extended to other electrical appliances.

4.2.1 EMI Variations of Compact Fluorescent Lamps

CFLs are one of the least complex appliances that generate EMI. They are composed primarily of a switched mode power supply that generates a high voltage to ionize the gases in the tube portion of the lamp, which in-turn produces light. I found that the EMI generated by a CFL can vary over time. More specifically, the EMI varies in response to changes in the impedance of the CFL bulb circuit due to variations in the local electric fields surrounding the bulb. These electric field variations are caused by the movement of grounded conductive objects (like humans) in the space near the bulb. When the bulb is illuminated, the ionized gases within the bulb are conductive and they form the sensing electrode of a capacitive sensor. Thus, sensing these EMI variations allows for sensing any local capacitive changes, such as those caused by human proximity. In other words, CFLs can be re-purposed as human proximity and motion sensors for modeling human behavior that requires no modification to the CFL bulbs, the lamp fixtures, or the human body, and many CFL fixtures can be used as sensors using only a single ElectriSense device plugged into any available electrical outlet.

I also observed other interesting properties of the CFL and its built-in power supply that could be used as a more general purpose distributed sensor. Specifically, I found that the oscillator in the CFL is sensitive to ambient temperature changes and these changes manifests as variations in the EMI. An important implication of this approach is that each CFL in a home can be thought of as a low-cost general-purpose sensor that is inherently networked since all bulbs are connected to the power line. Consequently, sensing variations in the EMI from CFLs has allowed me to investigate their use in modeling human behavior, their proximity to CFLs and detect motion in the home.

Since the CFL lamp by itself has no operating states, I chose to experimentally verify if changes in local capacitive fields and ambient temperature could be inferred from the EMI by defining *user induced operating states* as a variety of basic gestures performed by a user in proximity to a CFL. These included: (1) hover: the passing of a hand near a CFL-equipped lamp, (2) lamp shade touch: touching the lamp shade with one or more fingers, and (3) bulb touch: touching the bulb with one or more fingers; in addition to artificially varying

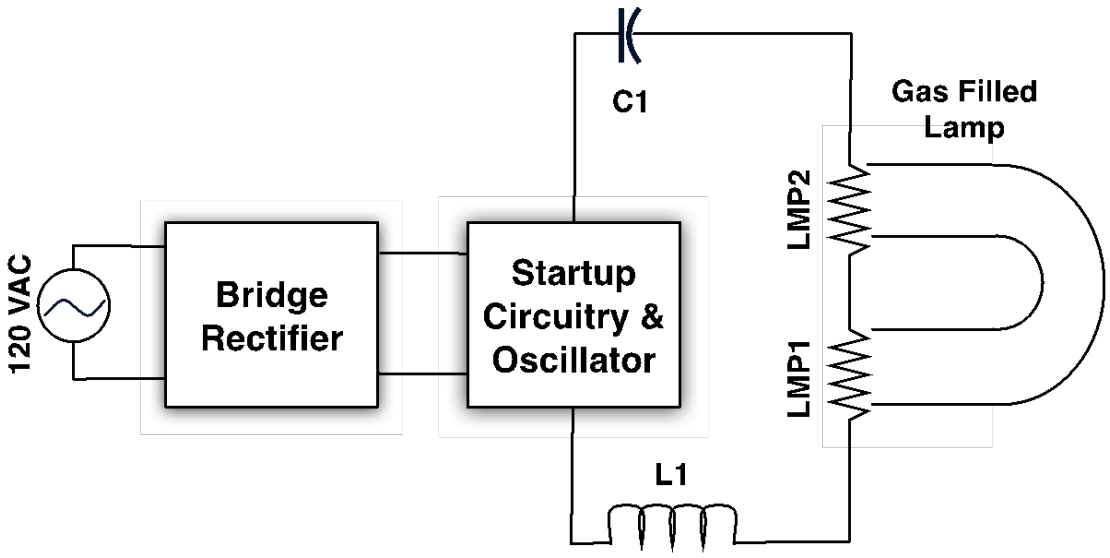

Figure 4.1: Simplified block diagram of a CFL. L1-Lamp-C1 forms a series resonant circuit. Human proximity causes the effective capacitance to change and detune the oscillator.

the ambient temperature around a CFL.

4.2.1.1 Variations in CFL's Line-Conducted EMI

When a human is in proximity to a CFL, the conducted EMI changes in amplitude, fundamental frequency, and the distribution of harmonic energy among the harmonics. By detecting these changes in amplitude and shifts in frequency, it is possible to infer when a human is in proximity to the lamp, and determine what kind of interaction took place. To better understand how the EMI changes as a result of human proximity, it is essential to understand the basic operation of a CFL. Figure 4.1 shows a highly simplified block diagram of a typical CFL.

A CFL has two main components: a gas filled tube (e.g., the bulb), and an electronic circuit called a ballast. In a CFL, the ballast is integrated with the bulb into a single unit; however, in other styles of fluorescent fixtures the ballast is separate from the bulb. The

ballast is connected to the line voltage (120V/ 60 Hz AC in the US), which is rectified and stepped up by a high frequency switching power converter to around 300V AC, which drives the fluorescent tube. The reason for using high switching frequencies, typically around 40 KHz, is the increased efficiency of the power converter circuit as well as the smaller size of components, thus making the entire CFL unit both compact and more efficient. This high frequency can be achieved by a variety of methods from simple L-C resonant oscillators in cheaper lamps to microcontroller-controlled oscillators in others. Irrespective of the mechanism to generate the high frequency, the purpose of the resonant circuit is to stabilize the current through the fluorescent tube as the bulb warms up, and then maintain equilibrium in lamp current and voltage. Thus during operation, the oscillator is at resonance with the L1-Lamp-C1 series network as shown in Figure 4.1.

This equilibrium can be disturbed in many ways, including changes in the supply voltage, current, or failure of a component. Additionally the bulb's equilibrium state is disturbed by capacitive and/or ambient temperature changes in and around the lamp. This state change can be detected in the bulb's EMI signature.

The bulb itself serves as a sensing electrode because the CFL's electronic ballast generates a strong electric field inside the bulb to ionize the internal gases and produce light. However, this process also produces a strong electric field around the periphery of the bulb, because the ionized gases act like a single equipotential electrode. The field that results is similar in essence to electric fields that are explicitly generated for capacitive proximity sensing [54], with the difference that, the former is a *by-product* of the way CFLs functions and latter is intentional.

By moving a hand in the vicinity of the CFL, the electric field surrounding the bulb is disrupted and there is an increase in the effective capacitance in the lamp circuit since the human body is essentially a capacitor to ground. In a traditional capacitive proximity sensor, this change in capacitance results in a change in displacement current, which is sensed directly. In contrast, the displacement current being coupled from a CFL results in a small imbalance in the oscillator and bulb equilibrium, resulting in a shift of the CFL's fundamental resonant frequency by 15%. A balanced oscillating circuit tends to suppress even-order harmonic products, while unbalancing it tends to enhance odd-order products.

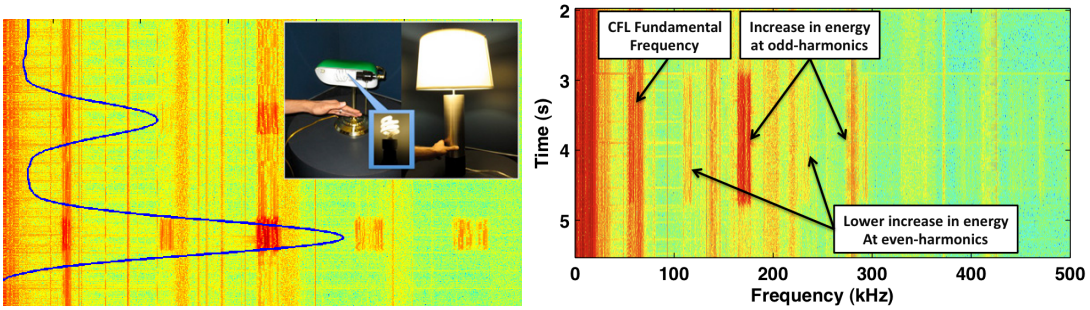

Figure 4.2: (*left*) A user bringing their hand closer and then touching the lampshade. The blue line shows the energy in the spectrum change as a result. (*right*) Spectrogram showing change in observed EMI as a result of human proximity to the CFL bulb, which causes a large capacitive change resulting in increased energy.

Figure 4.2 shows how the energy at odd-harmonics increases when a bulb is touched creating an even more significant imbalance than that created by human proximity. It should be noted that multiple similar CFL bulbs could be operating simultaneously in an environment or on a common electrical circuit without effecting the ability to detect the change in EMI energy. This is because, the harmonic energy only increases when the capacitance changes, thus two lamps could be on, but neither could be generating increased energy at the odd-harmonics.

When a user brings their hand close to the CFL lamp, it causes an increase in energy at the harmonics, but the change is much smaller and harder to visualize on a spectrogram. These changes are only reliably detectable after significant filtering and post processing as described in the following section.

4.2.1.2 Signal Processing Pipeline & EMI Tracking Algorithm

The hardware to capture EMI from the powerline is identical to that used in Section 3.1.2, however I developed a new signal processing pipeline to capture not just the presence or absence of EMI but also its variability over time. The magnitude in dB of each FFT vector as output by the ElectriSense prototype hardware is fed into this new signal processing

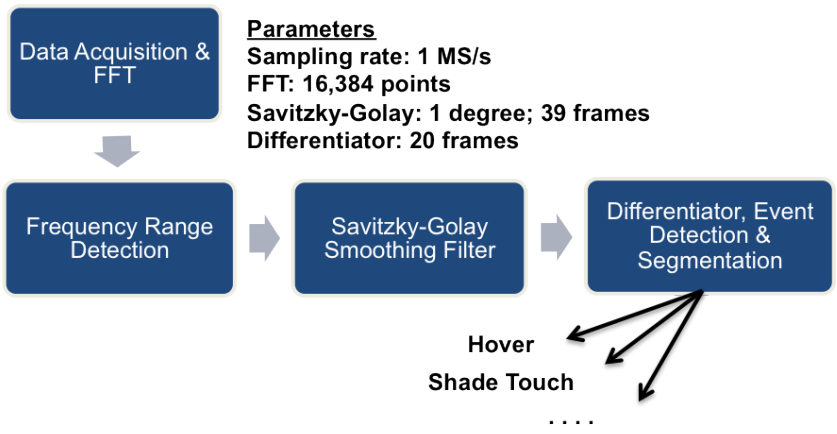

Figure 4.3: Stages of the signal processing chain that condition the input signal, apply differentiators, detects and segments events.

chain (see Figure 4.3), for processing and detection of events.

In this context, I define events as instants in time when a human is in proximity to a CFL and performing a gesture and causing a change in CFL's operating state. Since the input signal is real valued and we are only interested in the magnitude, the magnitude FFT vectors are symmetric around baseband, and hence half of the bins are redundant. These are truncated, reducing the number of points in the FFT vectors to 8192 across a spectral width of 500 KHz.

Frequency Range Detection

As described earlier, human proximity to a CFL appears as increased energy in harmonics of the fundamental frequency at which the CFL produces EMI. Thus, the first step is to identify the fundamental and harmonic frequencies. Though most CFLs have a switching

frequency between 40 KHz and 120 KHz, it can vary from one brand to another and thus setting a global range for all lamps is not practical. If the range to monitor is set too wide, the aggregate noise power from other sources reduces the overall signal to noise ratio. This makes the detection of changes in harmonic energy less reliable. Thus, a potential solution is to setup multiple narrow ranges specific to each CFL lamp and monitor each simultaneously.

My earlier prototypes required that these frequency ranges be set manually for each lamp. I later discarded this manual approach in favor of a simple one-time calibration step, which involves switching each lamp on and off twice. By turning the lamp on and off, the system can detect the fundamental and its harmonic by looking for presence and absence of EMI on the power line.

Actual measurement of the frequency of the fundamental is made only during the second on/off sequence, since during first interval the CFL bulb could be warming up, during which the frequency changes until the bulb reaches thermal equilibrium with its environment.

Once the frequency range for the desired bulb's 3rd harmonic is identified, a sum across this entire band is computed. For each FFT vector in time, a sum of energies in this range of frequencies is computed. Figure 4.4 (left), shows such a running sum plotted over time for a lamp. When the energy in this range increases as a result of a gesture or human proximity, the sum also increases. The reason a sum over a frequency band is computed instead of tracking individual frequency bins over time is because even when the CFL is in equilibrium, its fundamental frequency shifts a few KHz over time due to temperature changes. Tracking a range of frequencies (in tens of KHz) is immune to such small shifts.

Smoothing the Summed Energy

As evident from Figure 4.4 (left), the tracked sum over time is quite noisy and cannot be directly used for reliable event detection. In addition to EMI noise from other appliances, occasional broadband noise also plagues the signal. The source of such broadband noise could be from flicking mechanical switches [41], noise from dimmers, or other switch mode power supplies. Prior to taking the sum in a range of frequencies, I discard FFT vectors where the sum across the entire spectrum is more than 4 standard deviations greater than

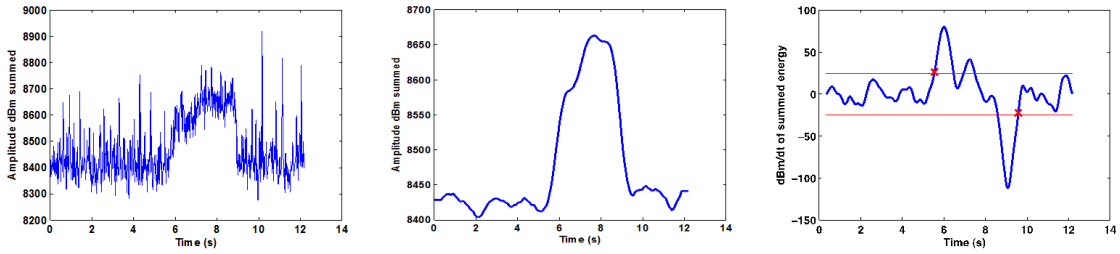

Figure 4.4: Summed energy in band over time shows increase in energy as a result of a hover (left), filtered summed energy (middle) and first derivative waveform along with start and stop of the hover event as detected (right).

the running mean over 15 FFT vectors.

To smooth the summed energy time series data, I experimented with several filters, beginning with the standard averaging FIR filters and a moving window Gaussian smoothing filter, but found that not only did they not remove the high frequency noise, but also "smoothened out" the peak from the gesture event itself, making it flatter and more difficult to detect. I found that the Savitzky-Golay smoothing filter with a degree of 1 and frame length of 39 to be most effective (see Figure 4.4 (center)). Not only does it remove the high frequency noise, but it also maintains the underlying shape of the peak. I apply two passes of the same filter on the summed data, which minimizes noise and makes the event detection more reliable.

Event Detection and Segmentation

I make use of a moving window first order derivative filter with a frame size of 20 to obtain a signal that captures any abrupt changes in the signal. Since small abrupt changes also cause the derivative filter's output to generate a peak, it is necessary to set a rejection threshold. That is, we only want to consider changes in summed energy that are abrupt as well as large in magnitude. I chose a threshold value of 25 dB/dt (dt = 20 frames, 0.327s), which has good sensitivity to detecting hovers and other gestures while still maintaining a low false positive rate. The ROC curve characterizing the tradeoff between true positive

rate and false positive for different threshold values is described later. It should be noted that this single threshold works across all people and all lamps that I tested. Figure 4.4 (right), the red lines show the threshold.

When a hand approaches the lamp there is an increase in the summed energy signal producing a large positive spike from the derivative filter. As the hand moves away from the lamp, the summed energy signal decreases back to its baseline level, producing a comparable negative spike in the derivative filter. I leverage this phenomenon in finding the beginning and ending of an event.

I also leverage the expectation that a large positive derivative should be directly followed by a large negative derivative to prevent the algorithm from over-segmenting a single event. This is necessary because during the event the summed energy signal is larger than normal, resulting variations in the signal that are also more intense for example, when a user's hand naturally moves closer or farther while performing a gesture. These subtle variations can also result in shifts in the derivative signal. An example is shown in Figure 4.4 (right). The signal rises above the positive threshold, falls below it and rises again (due to variation in signal during the event). We wish to ignore the second positive threshold crossing to avoid overly segmenting a single event. This is easily achieved by building hysteresis into the threshold values.

Events that surpass the derivative threshold must also meet a set of specific constraints. First, a positive derivative, which surpasses threshold, must be followed by a negative derivative value that surpasses threshold at least after 600 ms, but no more than 5 seconds. Second, the peak derivative value of a positive spike and magnitude of the negative spike must be within 30% of each other. This ensures that dissimilar spikes are not matched. Third, when two or more consecutive positive spikes are followed by a negative spike, the process is repeated for each spike, in chronological order. If a match is made, the remaining positive spikes are discarded and the algorithm moves on, begins searching for another positive spike after the matched negative spike. It should be noted that the maximum duration only limits the time for each gesture, not the time between them or the total duration of a complex gesture. For example, one can perform three hovers one after another forming a single complex gesture (at an application level) that lasts 6 seconds and each hover

will be detected. Points marked with an "x" in Figure 4.4 (right) show the start and end of a segment as found by this search algorithm. These segments are then outputted as events.

4.2.1.3 Observable Signals and Operating States

In this section, I detail the EMI variations that I observed from a CFL and also suggest the potential uses of it as a human activity "sensor". Unlike ElectriSense in Chapter 3, since not all of the minute changes in the EMI are visually apparent on a spectrum analyzer, my signal-processing pipeline discussed in previous section was critical in identifying some of these signal variations.

User Induced Operating State: Hover Gesture

As described previously, when a human is in proximity to a CFL lamp it can cause changes in capacitance of the resonating circuit resulting in variations in the observed EMI. This allows the system to sense hover gestures. Thus, when a person passes their hand close to a CFL lamp, it produces a detectable change in the energy levels of the EMI. Figure 4.5 shows how the EMI changes when a person hovers their hand close to a CFL. Observe how the energy in the 182 — 189 KHz band increases as a result. The small bumps are the slight movement of the person's hand under the bulb.

User Induced Operating State: Touches on Lamp Shade, Bulb, and Base

Much like a hover, touching the lampshade or the bulb also causes an increase in the harmonics energy. Since CFL bulbs are not dangerously hot like incandescent bulb, especially when installed in an open fixture, I decided to experiment with touching the bulb as well, which also causes the maximum capacitive change due to physical contact. To ensure safety, I measured the temperature of the bulb and found it to be around 40 C for a 13W CFL.

A hover event produces the smallest capacitive coupling, followed closely by touching the lampshade and bulb touch produces the most coupling. An interesting signal observed with touch gestures is that the amplitude of the signal is proportional to the surface area of

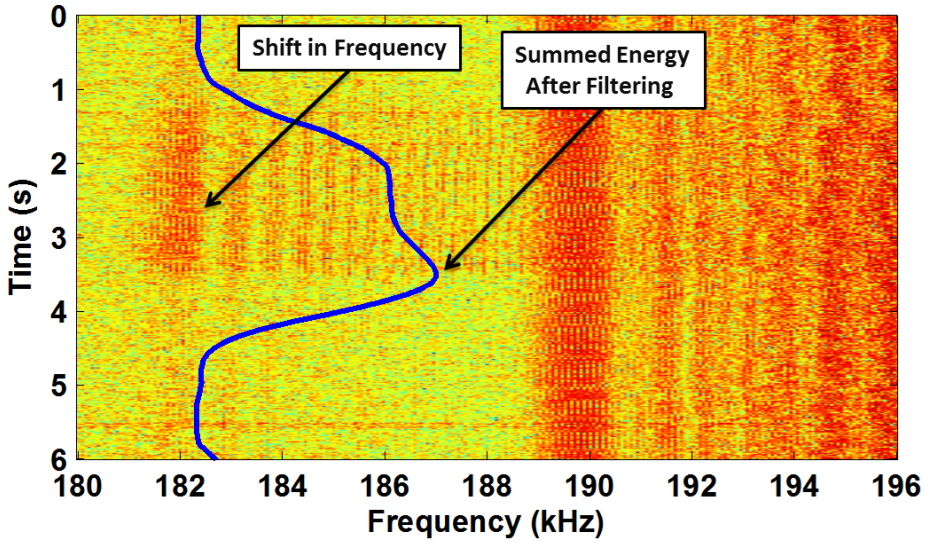

Figure 4.5: Summed energy in band (blue line) tracked over time overlaid with the spectrogram showing EMI from CFL. Notice the bump in blue line and increase in harmonic energy at 182 KHz in the spectrogram as a user performs a *hover* gesture or induces a hover operating state for 2.5 seconds.

the contact. That is, the energy amplitude of the signal observed is much higher when the bulb is touched with three fingers than when touched with one. Figure 4.6 shows an overlay of energy tracked over time for one-finger and three-finger touch on top of the spectrogram. Notice the amplitude difference between the first touch and the second. Also, for metallic lamp fixtures, it is even possible to detect touches on the base of the lamp, which produce a similar response to the lampshade being touched.

User Induced Operating State: Changes in Proximity

The change in proximity of a user's hand to the CFL, that is, whether the hand is moving away from the lamp or towards the lamp can be detecting by observing the slope of the signal. When the hand is moving towards the lamp, instead of seeing an abrupt increase in energy, a gradual increase is seen since the amplitude depends on how far the user's hand from the lamp is. Figure 4.7 shows an instance where the hand is moved towards the lamp

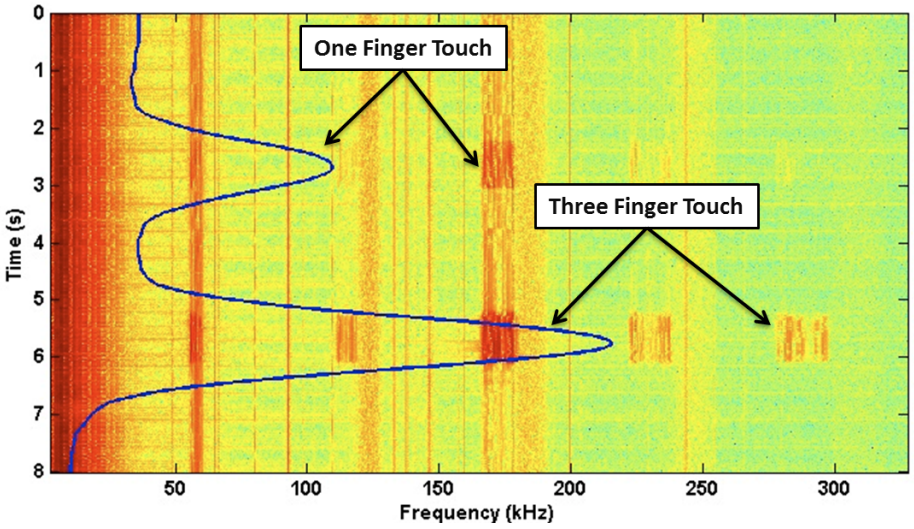

Figure 4.6: Notice the difference in amplitude of energy between a one finger touch on bulb and a three finger touch. Blue line is filtered sum of these energies over time.

and then abruptly removed. From our experience, the hand can be as far as 30 cm from the CFL bulb and be detected.

Detecting Ambient Temperature Change

In addition to variations in EMI as a result of capacitive changes from touches and human proximity, I also observed changes in EMI when the temperature of the CFL bulb or the ambient temperature around it changes. Instead of an increase in harmonic energy, changes in temperature cause the frequency to shift up to tens of KHz. Tracking the fundamental frequency over time can give an indication of changes in ambient temperature around the lamp. If calibrated, the true ambient temperature could also be measured.

I observed that when the ambient temperature around the lamp is increased, it causes the frequency at which the CFL's EMI is observed to shift to a lower frequency. To verify this phenomenon, I performed two experiments. In both experiments, I put a CFL bulb on a tray along with a thermocouple based digital temperature sensor. I let the bulb stay on for 10 minutes to reach room temperature (temperature of environment at time was 27 C).

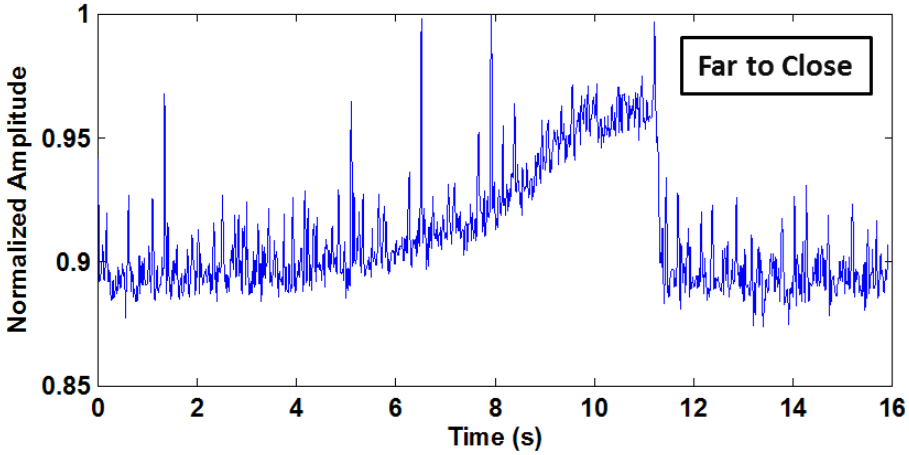

Figure 4.7: Amplitude of the observed energy depends on the proximity to the lamp. Amplitude gradually rises when a hand is brought closer over time.

In the first experiment, I used a heated chamber whose inside temperature was around 65 C. I then introduced the tray with the bulb and the temperature sensor into the chamber. The second experiment was similar, but instead of a heated chamber, a small refrigerator, with an internal temperature of 18 C was used. Figure 4.8 shows the change in frequency as a result of the heated chamber experiment. The change in the CFL's fundamental frequency tracks closely with the change in ambient temperature.

With these preliminary experiments I wanted to confirm that EMI changes as the ambient temperature around the lamp changes. This change in frequency with temperature opens up the possibility of using CFLs in a home as distributed temperature sensors. This technique could also be used to monitor large or abnormal increases in lighting temperature to detect potential safety hazards. I also observed that the CFL as a temperature sensor has an incredibly fast response time, which can be used as an air movement or air convection sensor. For instance, this can be used for detecting air drafts, opening windows or doors, or potentially air movement created by people passing by. For energy disaggregation purposes, sensing ambient temperature or changes in thereof adds another data point to sense potential HVAC activation, opening of windows or in general to sense temperature variations for thermal modeling.

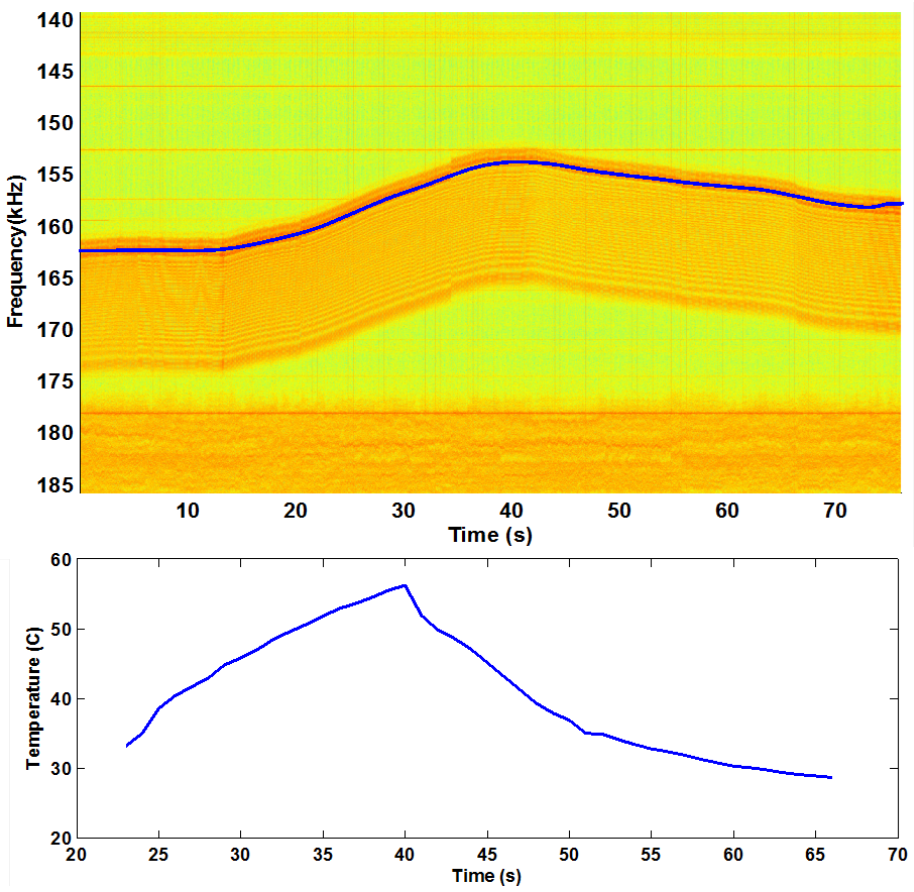

Figure 4.8: Shift in frequency (top) follows similar trend as ambient temperature (bottom).

4.2.1.4 Experimental Validation of CFL's Operating States Estimation

To verify how well the different user-induced operating states (gestures) could be inferred given the variations in the EMI and its applicability across various people, different brands of CFLs and lamp fixture styles, I conducted experiments with 10 participants (7 males, 3 females). These experiments were conducted in a simulated home environment in our lab with 6 lamps. I also conducted feasibility experiments with a subset of those participants and lamps in 2 actual homes.

I collected data for a variety of different lamp fixtures (see Figure 4.9) that included lamps with metal shade (3), glass shade (1), plastic shade (1) and no shade (1). The

Figure 4.9: Lamps that we used for our experiments. Two lamps (top left, top middle) were chosen to be similar for comparison. All lamps were on during the experiment.

CFL bulbs used also varied in brand: GE (1), Energy Star certified GreenLite/Sylvania (1), Commercial Electric (1), ProLume (1) and a generic brand (2) from a local home improvement store. This collection of bulbs we used also varied in wattages from 12W to 26W (mean 16W).

Data Collection Procedure

The simulated home setup included 6 lamps placed randomly throughout the lab space. The lamps were powered through separate residential style power source (a 60 Hz split single phase transformer). It should be noted that despite this isolating transformer, EMI from a large number of devices running in the building were observed, which made this have at least comparable EMI levels that would typically be found in a home environment. My own data collection and monitoring equipment produced EMI, as well as the many computers and measurement equipment that operated in the lab space.

Each participant in the experiment induced 5 different operating-states by performing 5 different gestures, with each gesture repeated 4 times. The gestures included: a hover, touching the lampshade with two fingers, touching the lampshade with hand, touching the

glass part of the bulb with one and three fingers. These 20 gestures were performed on each lamp one after another. To ensure consistency in where participants touched the lamp, I marked a spot on each lampshade. For one lamp that did not have any lampshade, instead of the two gestures involving touching the lampshade, I instead had the participants perform a hover to the left of the lamp and to the right of the lamp.

I built a data collection tool that allowed me to label the gestures as the users performed them. For each lamp, the software randomly ordered the sequence in which the gestures were required to be performed, in addition to randomizing the order in which each lamp was tested. This was done to minimize any temporal biases. Additionally, to minimize any effects from when the CFL lamps were turned on relative to start of the experiment, I turned the lamps off and only turned them on 5 minutes prior to each experiment. The entire experiment with each participant was repeated twice, on separate days to validate the stability over time. Each experiment took 45-50 minutes. The entire experiment ran for a period of two weeks during which I collected a total of 2400 events (120 per person, per session).

Feasibility in Homes

I randomly chose two lamps, and performed the same set of experiments in two different homes to confirm that the approach extends to an actual home. The homeowners were asked to continue their daily routine, thus lights and devices were being turned on and off as we collected the data.

Analysis of Hit Rate Accuracy

I chose a threshold of 25 dBm/dt (dt = 0.3277s) after plotting a ROC curve that characterizes the tradeoff between true positive rate and the false positive rate over the entire dataset (see Figure 4.10). The choice of 25 dBm/dt was based on maintaining good sensitivity (91.2%) towards detecting events while having a low (6%) false positives rate. All the results presented here on use this threshold.

Table 4.2.1.4 and Table 4.2.1.4 summarizes the hit rate for event detection and segmentation that was obtained for each lamp. Since hover events have the lowest energy change

Hover												
Lamp	**P1**	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9**	**P10**	**1st run**	**2nd run**
L1	1	0.75	0.75	1	1	1	1	1	1	1	0.95	0.9
L2	1	1	1	1	1	1	1	1	1	1	1	0.95
L3	1	1	1	1	0.75	1	1	1	1	1	0.975	0.9
L4	1	0.75	1	1	0.75	0.25	1	0.75	0.75	1	0.825	1
L5	1	0.25	1	1	1	0	1	0.5	0.25	1	0.7	0.85
L6	1	1	1	1	1	1	1	1	1	1	1	0.9
Avg.											**0.908**	**0.917**

Table 4.1: Average hit rate per user (P1-P10) for user-induced hover state's 1st run.

	Bulb Touch One Finger		**Bulb Touch Three Fingers**		**Shade Touch One Finger**		**Shade Touch with Hand**	
Lamp	**Mean 1st run**	**Mean 2nd run**	**Mean 1st run**	**Mean 2nd run**	**Mean 1st run**	**Mean 2nd run**	**Mean 1st run**	**Mean 2nd run**
L1	1	1	1	1	0.65	0.75	0.75	0.85
L2	1	1	1	1	0.725	0.775	0.8	0.85
L3	1	1	1	1	0.825	0.875	0.975	0.9
L4	1	1	1	1	*	*	*	*
L5	1	1	0.975	1	1	1	1	1
L6	1	0.975	1	1	1	0.975	1	1
Avg.	**1**	**0.996**	**0.996**	**1**	**0.84**	**0.875**	**0.905**	**0.92**

Table 4.2: Average hit rate for touch gestures across different participants.

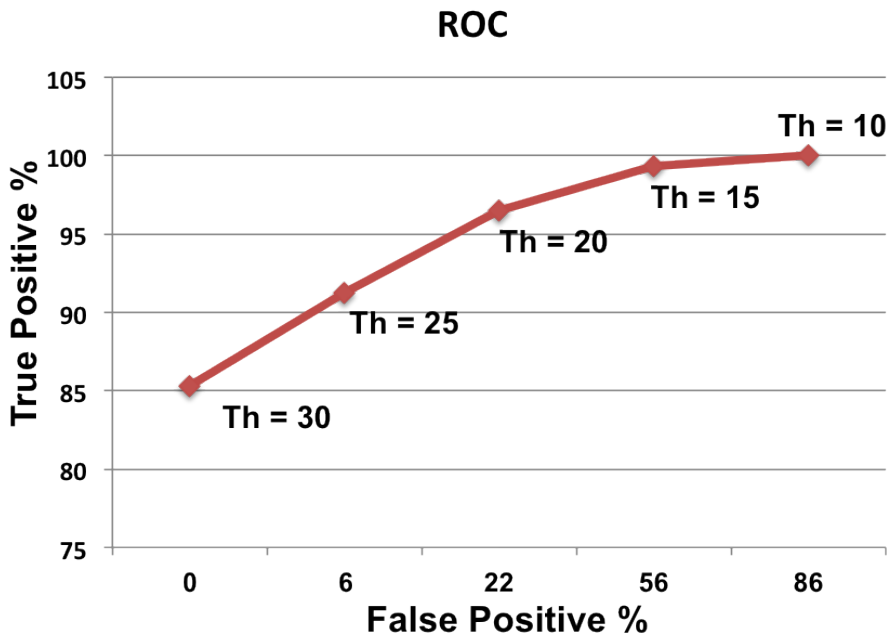

Figure 4.10: ROC showing tradeoff between the true positives rate and false positives rate for different threshold values.

relative to other touch-based gestures and are more challenging to detect and segment, I show the hit rate for hover events from the first session across lamps for each user. Means for each session for all gestures are also shown. Overall, hovers had an average hit rate of 90.8% for the first session and 91.7% for second. A key observation is that the mean hit rates were calculated are across all lamps and all users, suggesting the robustness of using a CFL as a "capacitive sensor" and inferring the resulting operating states using EMI.

Touches on bulb itself were easily detectable (100% and 99.6% between the two sessions), which can be attributed to the large energy change that passes the various event detection and segmentation checks. The average energy for bulb touch with three fingers events were 82.7 dB/dt (44.3 for one finger touches), roughly three times the set rejection threshold

I expected lampshade touches to perform well, but relative to both hover and bulb touches some of the lamps exhibited a lower hit rate. In particular, I found that L1 performed poorly (see Table 4.2.1.4). This lamp has a plastic lampshade (black frame with

Lamp	Hover	Bulb Touch One Finger	Bulb Touch 3 Fingers	Shade Touch One Finger	Shade Touch with Hand
L1	6.671	25.937	62.828	4.86	5.038
L2	6.155	43.521	77.853	8	8.837
L3	5.343	43.036	83.028	4.6	6.428
L4	6.23	40.114	69.43	*	*
L5	9.402	55.577	104.51	14.9	16.776
L6	6.594	45.201	85.4486	9.12	10.353
Avg.	**6.53**	**44.35**	**82.78**	**9.12**	**10.353**

Table 4.3: Average energy change (dB/dt) for each lamp and gesture. Notice the difference between the energy values for a hover vs. bulb touches vs. shade touches. Shade touch energy for L2 is low due to the grounding effect of the metal surface.

white shade in Figure 4.9). Lamp L3, which has a metal frame and a glass shade performed better than L1. The best performers were lamps L5 and L6, both of which have a metal body and metal shade. Clearly, the metal and glass shades are far more conductive, and hence easier to detect (see Table 4.2.1.4). In the case of L4, which did not have a lamp shade around the bulb and we found that there were no noticeable differences if the hand gestures or hovers occurring on the top, right or left of the illuminated bulb.

Interestingly, metal fixtures actually made the entire metal surface conductive and thus a touch gesture could occur anywhere on the stem or base. For example, a lamp with a stainless steel frame can be touched anywhere along the body, thus expanding the interactive area on the lamp. If a metal lamp was placed on another metallic surface, then that surface would also become a potential touch surface to some extent. Other metallic lamps exhibit this same behavior, where the entire lamp surface can act as a capacitive touch sensor. This was particularly evident in L2, which is the same style lamp as L6, but it performed worse because it was placed on a metallic table and appeared always grounded. The metal surface produces a grounding effect on the signal, thus touches to lamp had noticeably small increase in energy. When we placed L2 on a non-conductive surface, we found that the effect went away and the increase in energy was comparable to L6.

Overall, I found that EMI variations could reliably infer both touch and proximity

gestures with a CFL. Much of the variations in performance were not due to the person or differences in CFL, but the style of the lamp itself (metal, plastic, glass, etc). This implies that even though the CFL bulbs act consistently, the various lamp styles can add some variation to what operating states can be reliably sensed.

	Hover	**Bulb Touch One Finger**	**Bulb Touch Three Fingers**	**Shade Touch One Finger**	**Shade Touch with Hand**
Home 1	0.875	1	1	1	1
Home 2	0.875	0.875	0.75	0.875	0.875
Avg.	**0.875**	**0.9375**	**0.875**	**0.9375**	**0.9375**

Table 4.4: Hit rate from the two in-home experiments.

Table 4.2.1.4 summarizes the results of the in-home experiments. I found that the initial evaluation shows promising results for the approach extending to a home environment. The results are actually similar to those found in the lab. Given the lab scenario was likely noisier in terms of additional EMI noise on the powerline, these results seem plausible.

4.2.1.5 Discussion

Overall, both the in-home experiments and lab experiments strongly indicate that temporal variations in the EMI could be reliably used to detect and differentiate between different user-induced CFL operating states. The core experiments also reveal that the EMI signal is temporally stable. A single segmentation algorithm was developed to isolate and detect the various operating states or gesture events across all of the different participants and various brands of CFLs. I also found that multiple CFLs (even those of similar brands) operating simultaneously can still be distinguished using this approach. There is enough of a difference in the operating frequency between two similar bulbs, that the odd harmonics are also distinguishable.

Although the methods presented in this case study primarily considered the 3rd harmonic, I found that there is an increase in energy not only at the 3rd but also the 5th and 7th harmonics. The algorithm can potentially be improved by tracking energy in multiple

higher order harmonics. This could also alleviate a key limitation, which is the potential for a strong interfering device to mask the smaller energy changes from the CFL. Because the fundamental frequency of the CFL changes based on the ambient temperature (the temperature sensor effect), the location of the odd harmonics also changes. Thus, to infer both a temperature induced and a capacitive proximity induced operating state, both the fundamental and its odd harmonics need to be tracked while it is operating. In Section 4.4 I introduce an algorithm capable of such dynamically tracking of EMI.

4.2.2 EMI Variaitons of Modern Televisions

Unlike CFLs in the previous section, televisions are comparatively complex electronic appliances and the operating states they exhibit are complex and larger in number. For example, some of the many different operating states of TVs could be the type of input (HDMI or Component) selected, the brightness of the screen, transitional states that occur when channels are switched etc. Keeping this complexity in mind and to further develop an understanding of how EMI varies as a function of an appliance's operating states, I chose to investigate one of the most ubiquitous classes of consumer electronic devices: modern TVs.

4.2.2.1 Overview

I found that for TVs, the EMI variations follow the screen brightness closely, so much so that given a particular time-varying EMI trace, the properties of the visual content being played on the screen can be inferred. In other words, as a TV's operating state changes, such the changes in the screen brightness as a result of changing visual content, the resulting EMI variations correlate strongly. On further investigation, I found that this happens because modern televisions make use of energy efficient power supplies (such as SMPSs or their derivatives). For such power supplies to be efficient, they inherently *adapt* to the power requirements of the appliance they power. That is, the power supply adapts to the load patterns of the appliance, which in effect results in the EMI to vary as a function of these load patterns. In case of TVs, as the screen becomes dark, it requires less energy to which the power supply responds by going into a lower power mode, and when the screen needs

Television Name	**Technology**	**Date of Fabrication**	**Price (USD)**
Panasonic-42-A	Plasma	Sept 2010	$1099
Panasonic-42-B	Plasma	Sept 2010	$1099
Samsung-58-A	Plasma	Sept 2010	$2599
Samsung-58-B	Plasma	Sept 2010	$2599
Samsung-32	LCD	Oct 2007	$896
Sharp-42	LCD	Sept 2008	$1399
Sharp-32-A	LCD	Dec 2009	$499
Sharp-32-B	LCD	Dec 2009	$499

Table 4.5: A summary of the tested TVs showing make, technology, date of manufacture, and price for each device.

more energy it does the opposite. Thus, depending on the screen content, the power supply cycles through a particular pattern of low and high power modes that directly correlates to the dark and bright scenes of the content. As the power supply's operation changes, it manifest as varying EMI over the powerline. I confirmed this behavior further by correlating EMI variations with power variations.

To further build on this finding, we performed an experiment with multiple TVs playing a variety of screen content and verifying whether the EMI variations correlate to the screen content. In particular we asked the following research question: *"Given the EMI is it possible to infer the TV's operating state, and more precisely infer what content is being played on the screen?"*. This section details our experiments in pursuit of answering the aforementioned research question.

For breadth and depth, we obtained a collection of eight TVs spanning three manufacturers (Panasonic, Samsung, and Sharp),[1] two technologies (LCD and Plasma), and three sizes (32", 42", and 58"); see Table 4.5. We repeatedly played video sequences on each of these TVs (to cause change in operating states) while recording the EMI that these TVs

[1]We do not have any reason to believe that our results are specific to these manufacturers.

produced as measured on the powerlines. We obtained initial measurements in a clean laboratory setting, and then collected further measurements in unrestricted home environments with diverse collections of other electromagnetically noisy electronics attached to the same powerline.

We found that all but the Sharp 32" TVs produced a significant amount of stable, robust operating state information over the powerline via EMI. As an example, for the Panasonic 42" TVs in a laboratory setting, we found that we could correctly match a random 15 minute EMI trace from a movie to a database of EMI signatures totaling 1200 movie minutes 96% of the time (the remaining 4% were below our matching threshold, not incorrectly matched); the hit rate drops to 92% in noisy home environments. Broadly, we sought to investigate the following questions:

- Do TVs of different makes, models, and technologies produce repeatable EMI patterns, when going through the same operating states given repeated screen content?

- Do different operating states produce differentiable EMI for a given TV?

- Are EMI patterns consistent across multiple instances of TVs from the same model family?

- Can we determine if someone is watching a particular video by matching an EMI sample to a database of EMI signatures?

- Can we match EMI recorded in a lab setting to EMI extracted in homes, especially when those homes might have a diverse array of other consumer electronics connected to the same powerline?

We found that we are able to answer all these questions affirmatively, in turn reinforcing the hypothesis that an appliance's operating could be reliably inferred given the variations in the conducted EMI it produces.

4.2.2.2 System Description

The hardware setup used to collect EMI from the powerlines in various home and lab scenarios was identical to that used in Section 3.1.2 with the alteration that we modified the powerline interface R-C filter to have a lower high pass cutoff. This was necessary since the original design employed a filter with 36 KHz cutoff and having observed several TVs with EMI at much lower frequencies, we designed a new filter by altering the R-C network to have a cutoff at 6 KHz. The experimental infrastructure was highly automated; it could automatically stream video content to TVs while simultaneously logging EMI signals to allow for later analysis.

Lab vs Home Setup

Over the course of our experimental analysis we collected EMI data from lab and home settings. During lab recordings, we connected the TV being tested and our PLI sensor to the output of an isolation transformer (model Tripp Lite IS500HG). The isolation transformer reduces electrical noise and EMI present on the power line and presents a cleaner AC power at it's output [2]. Generally, isolation transformers are used to protect against electric shock, or to suppress electrical noise in sensitive devices such as high-end audio systems.

Conversely, in home recordings, we did not use the isolation transformer but rather connected the PLI directly to any available electrical outlet. This meant that we were no longer guaranteed a clean electrical background, and our target EMI could now be affected by the noise produced by any of the devices on the powerline.

4.2.2.3 Selection of TVs and Movies

To evaluate to what extent EMI produced by modern TVs varies, we procured a set of eight TVs that differ along a number of axes including: manufacturer (Panasonic, Samsung, Sharp), display technology (LCD or Plasma), size (32", 42", 58"), and date of manufacturing (2007–2010).

[2] We recognize that a "Line Impedance Stabilization Network" LISN device would have provided a more optimally stable and repeatable signal; however, we expect a LISN to only improve the already high matching accuracies achieved with our acquisition hardware.

Genre	Movie
Action	(1) Lord of the Rings: Return of the King, (2) Star Wars V: Empire Strikes Back, (3) The Bourne Ultimatum, (4) The Matrix
Animation	(5) Wall-E, (6) Shrek 2, (7) The Lion King, (8) Aladdin
Comedy	(9) Office Space, (10) Meet the Parents, (11) The Hangover, (12) Wedding Crashers
Documentary	(13) Planet Earth: Fresh Waters, (14) Food Inc., (15) An Inconvenient Truth, (16) Top Gear (Season 4;Episode 7)
Drama	(17) The Shawshank Redemption, (18) American Beauty, (19) Titanic, (20) Requiem for a Dream

Table 4.6: List of twenty movies that we included in our dataset selected such that they span various genres.

Table 4.5 summarizes our hardware selection. We opted for an equal distribution of LCDs and Plasmas – the two dominant technologies on the market. Furthermore we deliberately selected three pairs of duplicate TV models to enable analysis of EMI signature similarity across identical hardware. This table also introduces a naming scheme that is used in the remainder of this section. Under this convention, television names include manufacturer, size, and model-pair information. For example the name Samsung-32 indicates that the Samsung TV is 32" and that we have only one model instance; similarly, Panasonic-42-A is a 42" Panasonic TV which is the first of a pair (as indicated by the "A").

Next we needed to choose what video signals to send to the TVs while recording EMI. For experimental feasibility, we opted to create a 20 movie database. To make our movie selection systematic, we chose four top-rated films within each of five distinct genres (ratings were extracted from the Internet movie database, `imdb.com`). Since these movies had variable running times, we limited our analysis to the first 60 minutes of each DVD video stream. Table 4.6 lists the movies that constitute our library. Looking ahead, the results from studying TV EMI results suggest that the methodology is applicable far beyond a collection of 20 target movies.

4.2.2.4 Data Collection Procedure

We performed recordings on each of the 8 TVs in a lab environment where we could carefully regulate the electrical conditions. This data accounted for 13 days of EMI traces, and provided the foundation of most of our analysis. Next, we selected a single TV (Panasonic-42-B) and deployed our system to three home locations to gather data in electrically uncontrolled settings.

Lab Recording:

During lab recordings, we plugged a given TV and the PLI into an isolation transformer to remove interference from other electrical devices in the building. We then streamed our database of 20 movies twice[3] while logging EMI in the 6 to 250 KHz range.

Home Recording:

To evaluate EMI based information leakage under a more naturalistic setting, we also performed data collection in three homes with distinct sizes, neighborhoods, ages (1906, 2003, 2009), and styles (apartment, suburban house, multi-family home) on a subset of our movie database (12 movies, 15 minutes each). Table 4.7 shows the specifications for these residences. During the home data collection, we did not prevent occupants from actuating electrical devices and we also did not use the isolation transformer. Also, we plugged in our sensor in a arbitrarily chosen electrical outlet that was available. Thus, we captured not only the EMI from the TV but also any number of active electrical devices operating in the home under natural conditions (e.g., lamps, kitchen appliances, computers, other televisions). Many times during our home data collection EMI from other devices overlapped with the EMI from the TV under investigation. In addition, Home#2 used power line communication (PLC) during recordings, however we did not observe significant signal interference since home PLC signals are typically transient and narrow-band in our region of interest.

[3]Two runs of our movie library were recorded on each TV to investigate the repeatability of EMI signatures.

ID	Style/ Year Built	Size/Floors
H1	Single family Home/2003	3000 sq. ft/ 2 flrs
H2	Apartment/2009	657 sq. ft/ 1 flr
H3	Multi-family Home/1906	800 sq. ft/ 3 flrs(*)

Table 4.7: A summary of the homes showing the style, size and number of floors. (*) 800 sq. ft refers to only one apartment in a converted multi-family home.

4.2.2.5 Signal Conditioning and Automated Signal Identification

To facilitate analysis we needed to devise a method for extracting meaningful signals from the FFT data captured using our hardware. We used the search algorithm detailed later in this chapter in Section 4.4 for automatically identifying the range, location and type of EMI produced by a device and extract the appropriate time varying signal through peak tracking. The two types of EMI we observed were energy density and frequency shift types. The former causes the amplitude of the EMI to vary while the latter causes the center frequency at which the EMI is produced to vary.

Using the search method, we were able to characterize the EMI signatures for all of the TVs in the lab setting (Table 4.8). In addition, the search method proved invaluable when we analyzed EMI from home recordings as different residences had unique levels of dynamic background noise which overlapped with tracking ranges we found in lab; our methods were always able to find frequency regions minimally influenced by the other powerline noise and thus enabled us to recover the TV EMI signal (due to harmonics).

4.2.2.6 Analysis and Results

We began by experimentally establishing that individual televisions produce repeatable EMI traces when the same video content is played multiple times. Building on this result, we found that the EMI signals produced by multiple TVs of the same model are also highly correlated given identical video inputs. We further used our collection of 20 movies to study how well could we match the EMI collected from one movie with a database of previously

Television Name	Freq. Range	Signal Type
Panasonic-42-A/B	1 to 60 KHz	Energy Density
Samsung-58-A/B	45 to 55 KHz	Frequency Shift
Samsung-32	10 to 50 KHz	Energy Density
Sharp-42	60 to 90 KHz	Frequency Shift
Sharp-32-A/B	35 to 40 KHz	Energy Density

Table 4.8: Frequency regions and EMI signal types for our 8 TVs.

collected EMI measurements from all 20 movies. In other words, given the EMI, how well could the TV's operating state (in this case the content being played on the TV) be inferred. Next, we studied the feasibility of this approach in three home environments, where additional signals sources are active on the power line. The data from homes also allows us to reason about the possibility of matching EMI traces gathered in an electrically clean setting to those collected from noisy environments (i.e., matches are possible between lab and home EMI).

Signal Repeatability in Lab:

The first question we sought to answer was whether repeated video content played on a target TV would produce consistent EMI. To test this, we computed the cross-correlation of a movie's EMI trace between multiple recording sessions. The left column of Table 4.9 contains the cross-correlation statistics when this computation was applied to every movie in our database and repeated for each TV, using only the first 15 minutes of each trace. These levels of high cross-correlation suggest that for most televisions multiple runs of identical video content produced consistently similar EMI. The exceptions, which we consider in more depth later, are the Sharp 32" TVs. Conversely, dissimilar video content tended to produce highly uncorrelated EMI traces. The latter conclusion was drawn from the results of our followup experiment which analyzed the potential for similarity between EMI traces from different video streams. More specifically, we recorded the highest levels of cross-correlation possible while matching a movie against all other films in the database (itself excluded); our measurements are reported in the right column of Table 4.9.

While the above findings strongly indicate operating state estimation using EMI for TVs, it's not clear whether this approach is generalizable across different TVs, but of the same model. The question arises: are operating state estimation models applicable from one TV to another? We explore that question here by assessing how similar the EMI is between two TVs of the same model. We performed this test by replicating the methods above, but instead of comparing EMI signals from multiple recording sessions on a single TV, we analyzed the correlations between EMI signals between single sessions of EMI recorded on TVs within the same model family (i.e., Panasonic-42-A vs Panasonic-42-B, Samsung-58-A vs Samsung-58-B, and Sharp-32-A vs Sharp-32-B). As before, we computed the EMI similarity of identical content (Table 4.10, left column) as well as distinct content (each movie against all others, Table 4.10, right column).

The experiments suggested that, with few exceptions, EMI from identical video signals are highly repeatable while distinct video sources yield independent EMI signals. These results were shown to hold within individual TVs as well as across pairs of TVs within a model family. In both cases, the Sharp 32" TVs produced outlier data points, reflecting a significant decrease in EMI consistency relative to the averages of other TVs. We attribute this to the fact that these devices had the least dynamic EMI (minimally modulated by screen content changes) within the televisions we tested. We tested the power consumption of the Sharp 32" TVs and found that their power draw changed by only 1W (175W to 176W) when playing a software generated custom video, whereas the Sharp 42" TV exhibited a 24W (190W to 214W) change using the same video signal. Thus we attribute the small dynamic shifts in EMI of the Sharp 32" to the power consumption characteristics of the circuitry. There is also a possibility that this particular model incorporates a fixed frequency power supply.

Identifying Movie Using Varying Length of EMI:

Our initial experiments suggested that extracted EMI signals are consistent within TVs as well as between identical models using cross correlations computed over 15 minute segments of content. That is, given a 15-minute varying EMI signature, it is possible to estimate the TV's operating state. We next evaluated how varying the length of this EMI signature

TV	Cross Correlation Same Content (%)		Cross Correlation Different Content (%)	
	Average	**Standard Deviation**	**Average**	**Standard Deviation**
Panasonic-42-A	98.18	+/-2.77	60.36	+/-12.83
Panasonic-42-B	98.99	+/-0.98	59.51	+/-12.97
Samsung-58-A	97.47	+/-1.97	53.29	+/-10.22
Samsung-58-B	96.71	+/-2.73	53.37	+/-9.70
Samsung-32	98.26	+/-5.21	65.56	+/-16
Sharp-42	97.30	+/-8.26	63.45	+/-12.84
Sharp-32-A	60.03	+/-8.12	56.41	+/-5.10
Sharp-32-B	60.91	+/-8.39	56.39	+/-5.79

Table 4.9: Analysis of EMI signal similarity given identical and different video content within a TV - averaged over the 15 minute intro segment of all movies.

TV	Best Cross Correlation Same Content (%)		Best Cross Correlation Different Content (%)	
	Average	**Standard Deviation**	**Average**	**Standard Deviation**
Panasonic-42-A and B	96.82	+/-4.67	59.93	+/-12.85
Samsung-58-A and B	94.96	+/-9.23	59.30	+/-10.09
Sharp-32-A and B	77.21	+/-18.84	53.31	+/-12.96

Table 4.10: Analysis of EMI signal similarity given identical and different video content between TV pairs (averaged over all movies). For this we used randomly selected 15 minute streams from each 60 minute movie

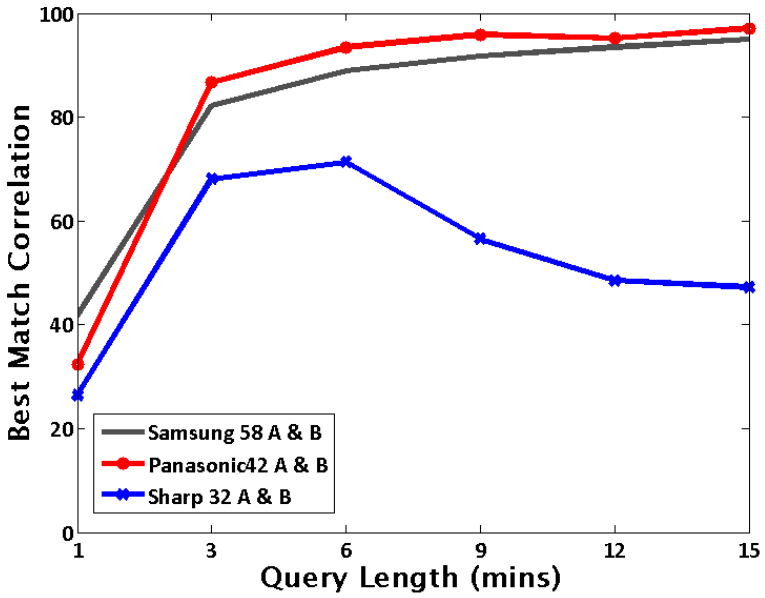

Figure 4.11: Cross Correlation results at various query lengths.

impacts signal repeatability and in effect ability to correctly identify the operating state.

We designed a search procedure that would take as input a full EMI trace of a movie and extract 10 query segments (of variable lengths) from multiple (randomly chosen) starting indexes along the 60 minute stream. Each query segment would then be matched against the EMI signals of the entire movie database (using sliding-window cross correlation), and the highest correlated match would be returned. This technique is also referred to as matched filtering in signal processing communities. We applied this procedure to find matching signals between TVs of the same model; query data would be extracted from the EMI library of the first TV in the pair, and the database signals would come from the other (e.g., query = EMI from Panasonic-42-A, database = EMI from Panasonic-42-B).

We performed a series of runs using query lengths ranging between 1 and 15 minutes (1, 3, 6, 9, 12 and 15 mins) and in each instance computed the cross correlation between the query and its best match (averaged across 10 samples for each query length). The matching results for the three TV model pairs are show in Figure 4.11. Even short length queries tend

to find highly correlated matches in the database. In particular, for TVs with repeatable EMI signals, once the query length reaches 6 minutes the correlation reaches 93.7%.

Interestingly, for the Sharp-32-A and B TVs, longer queries led to degraded performance. We attribute this to the weak EMI signatures of these TVs (as previously mentioned) which are susceptible to noise whenever the EMI signal is not being modulated along its entire dynamic range. However, the Sharp-42 performed very well with just short queries.

Home Results:

Having found evidence of significant EMI variation from TVs within a lab setting, we next turned to evaluating our approach in a more natural environment. We deployed a TV in multiple homes and see if we could match queries from home EMI to lab EMI datasets. In choosing which television to utilize for our residential recordings, we initially restricted the candidate pool to TVs for which we had duplicate models (Panasonic-42-A&B, Samsung-58-A&B, and Sharp-32-A&B). We rejected the Sharp 32" TVs based on their weak EMI signatures. Given the choice between the Panasonic 42" TVs and Samsung 58" TVs, we selected the smaller TVs because they were significantly easier to transport.

We set up our system (Panasonic-42-B, PLI, and logging equipment) in three different homes (see Table 4.7) and in each context recorded a smaller version of our database (3 hours total - first 15 minutes of the first 12 movies). Next we matched queries extracted from home EMI (Panasonic-42-B) against the EMI signature database collected in the lab (Panasonic-42-A). We found that the hit rate for 15 minute queries drops from 96% in the lab environment to 92% in the home environment, strongly supporting our argument that even in environments with other EMI present, our tracking and EMI variation extraction algorithm works well.

4.2.2.7 Discussion

Both the lab and in-home experiments strongly indicated that temporal variations in the EMI from modern TVs could reliably be used to infer their operating states. In particular, the amount of information in this varying EMI is high enough to even reliably predict

TV screen content. With this promising findind in mind, we additionally performed a preliminary survey of other common home electrical appliances to assess if their EMI also varies as a function of their operating state. We found that laser printer generated increased electromagnetic interference while warming up and during the printing process that engages the motors. In the case of the personal computer and laptop, we observed increases in EMI amplitude while running a collection of intensive CPU/memory/disk benchmarks. A DVD player's (with internal FM Radio) EMI was affected by volume changes, and the game console produced different EMI signals depending on whether it was in the main menu or in the process of rendering a game.

To summarize the finding in this case study: we conducted an extensive study of operating state estimation using EMI from eight modern TVs spanning two technologies (LCD, Plasma), three sizes (32", 42", 58"), and three manufacturers (Panasonic, Samsung, Sharp). We found that these TVs generally produce stable, robust EMI signals on the powerline that correlate with the video being displayed. We experimentally evaluated our approach with 20 movies in both a clean laboratory setting and electrically noisy home environments with other active devices connected to the powerline and found strong evidence suggesting that the temporal patterns in EMI from TVs could reliably be used to estimate their operating states.

4.2.3 EMI Variations from LCD Driver Circuitry

In addition to TVs, I have also investigated general-purpose LCDs like computer monitors and found that *another* source of EMI variation is that from the display driver circuitry. Similar to the capacitive coupling of humans to CFLs as described in section 4.2.1, proximity to a LCD can result in capacitive coupling with an LCD's driver circuitry resulting in EMI variations. This section details my investigation and findings of estimating operating state in the form of detecting a human touching the display of LCD monitors.

In this subsection, I describe the theory of operation and present a feasibility study demonstrating the ability to detect and classify 5 different operating states in the form of gestures (see Figure 4.12) across 11 users and 8 LCD displays (on both desktops and

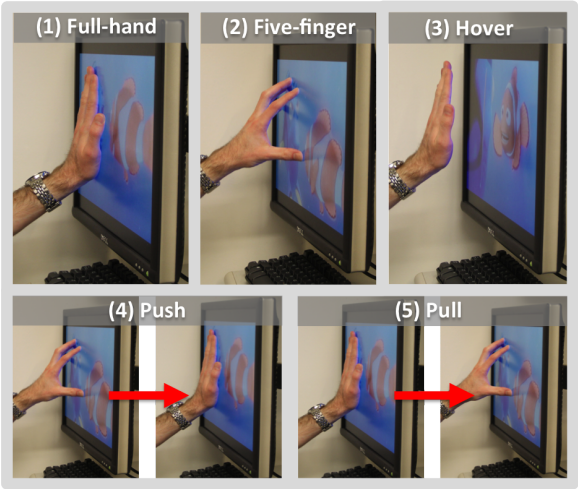

Figure 4.12: 5 gestures or induced operating states that can be inferred from the time varying EMI a LCD's internal driver circuitry produces: (1) full-hand touch, (2) five-finger touch, (3) hover, (4) push, and (5) pull.

laptops).

4.2.3.1 Theory of Operation

In LCD monitors, a backlight produces uniform white light using either a cold cathode fluorescent lamp (CCFL) or light emitting diodes (LED). The white light passes through a polarizer, liquid crystal (LC), color filter, and a second polarizer before being emitted at the front of the display. The intensity of the light is controlled by the strength of the electric field applied to the LC. Pixels are made by closely grouping red, green, and blue colored filters, which visually combine to produce the desired color.

Although the panel is made of a large array of pixels, only a single row of pixels is on at any time, and therefore small thin-film transistors (TFT) are used to enable each pixel. Figure 4.13 shows a small section of an LCD panel array. With the gate voltage applied only to the active row, a field is created on all electrodes in that row. Each row is selected once per frame, and enabled periodically at the refresh rate. We will refer to the rate at which the display switches active rows as the row rate. The row rate is dictated by the refresh

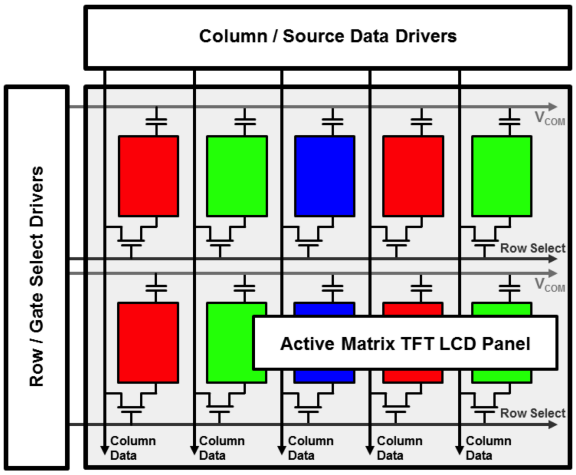

Figure 4.13: Simplified schematic of a small section of the LCD panel and the row and column drivers.

rate of the display (commonly 60 Hz) and the number of rows (i.e., the native resolution), and does not change when the driving resolution is changed.

As explained above, the row select lines and column data lines are changed every row, at the row rate. As a result of this, the row and column drivers consume power in bursts at the row rate. In the same way that current spikes from a digital clock (or an oscillator in a SMPS) couples EMI onto the power line, the current spikes from the row and column drivers result in EMI on the power line at harmonics of the row rate [2]. On some monitors, EMI is also observed at harmonics of half of the row rate. This is because some LCDs group adjacent rows in what is called line-paired inversion. In this case, if we assume that the colors of nearby pixels are typically similar, As explained above, the row select lines and column data lines are changed every row, at the row rate. As a result of this, the row and column drivers consume power in bursts at the row rate. In the same way that current spikes from a digital clock couple EMI onto the power line, the current spikes from the row and column drivers result in EMI on the power line at harmonics of the row rate [18]. On some monitors, EMI is also observed at harmonics of half of the row rate. This is because some LCDs group adjacent rows in what is called line-paired inversion. In this case, if we

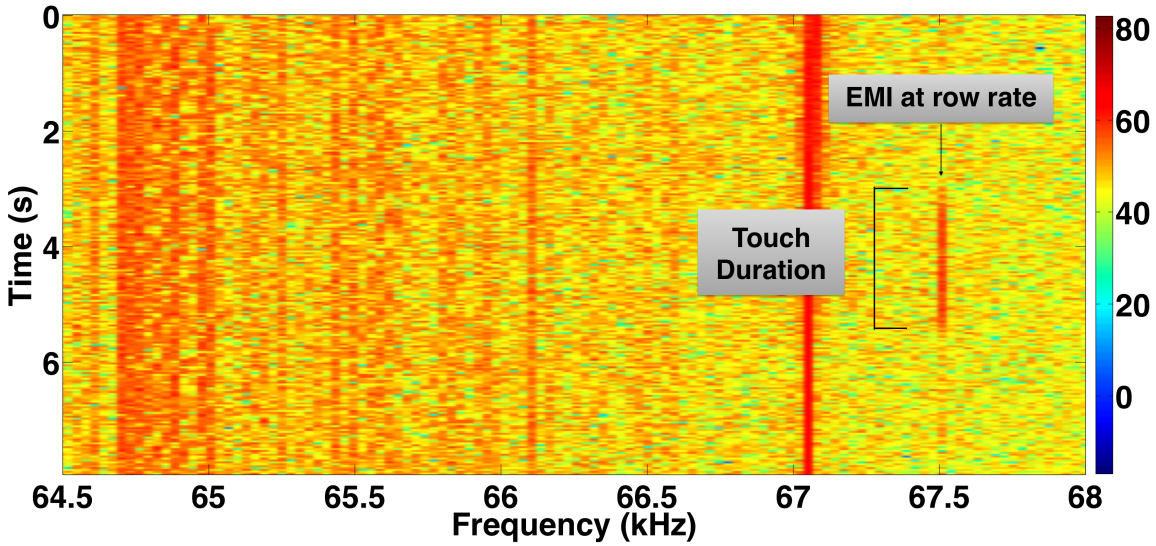

Figure 4.14: EMI seen at the row rate of 67.5 KHz is observed well above the noise level during the user's touch of the LCD.

assume that the colors of nearby pixels are typically similar, then the voltages on the column data lines only change significantly every other row. This will therefore cause EMI at half of the row rate.

Note that since the EMI is produced by multiplexing the rows of a panel, it is thus independent of the backlight technology (i.e., LED vs. CCFL) and independent of the pixel-level electrode configuration (i.e., TN vs. IPS). All such varieties of panels have the same type of TFT array and therefore produce EMI at the row rate in the same way.

Sensing Panel Touches:

Although the row rate EMI on the power lines is typically below the noise level, when a user's hand hovers over or touches the panel, a very large capacitance to ground is added in parallel with the row select lines and column data lines. This added capacitance results in significantly higher power consumption by the row and column drivers, which causes higher levels of EMI at harmonics of the row rate (and half of the row rate in LCDs that use line-paired inversion). This EMI is both conducted onto the power line and radiated onto the power lines by the panel and user. The resulting EMI on the power line can then be

ID	Model	Size / Resolution	Panel / Backlight	Row Rate (KHz)
M1	Asus VW246H	24 in / 1920x1080	TN / CCFL	67.5
M2	ViewSonic VX2035wm	20.1 in / 1680x1050	TN / CCFL	65.3
M3	Samsung 226BW	22 in / 1680x1050	TN / CCFL	64.7
M4	Dell 2007WFP	20 in / 1680x1050	S-IPS / CCFL	64.7
M5	Dell 1703FPs	17 in / 1280x1024	TN / CCFL	64
M6	HP S2231	21.5 in / 1920x1080	TN / CCFL	67.5
L1	Acer ASPIRE5736Z	15.6 in / 1366x768	TN / LED	47.1
L2	Dell INSPIRON1545	15.6 in / 1366x768	TN / CCFL	49.4

Table 4.11: Specifications of the six monitors (M1-M6) and two laptops (L1, L2) used in the experiment.

seen well above the noise level as shown in Figure 4.14. Furthermore, the relative amplitude change of this EMI is a function of the strength of the capacitive coupling between the panel and the hand. This noise is therefore a robust signal for sensing different kinds of touches and hovers on the panel.

4.2.3.2 Experimental Procedure and Analysis

An experiment was setup in a real home environment and a user study conductedwith 11 participants (3 female) on 6 LCD monitors (M1-M6) and 2 laptops (L1-L2) to validate how well the state of LCD could be inferred from the EMI. Table 4.11 shows the variety of panel technologies used in the study.

Five different touch gestures were used to induce operating states as shown in Figure 4.12: full-hand touch (FH), five-finger touch (FF), hover, push and pull. In a push gesture, the user was asked to first perform a five-finger touch, followed by a full-hand touch (i.e., the hand is pushed toward the panel), and the pull was exactly the opposite. Each participant performed 6 repetitions of each touch gesture on each of 6 monitors (M1-M6), and 5 of the 11 participants also collected data on 2 laptops (L1-L2). The order in which the gesture were performed was randomized to mitigate any temporal bias. Furthermore, for consistency, all

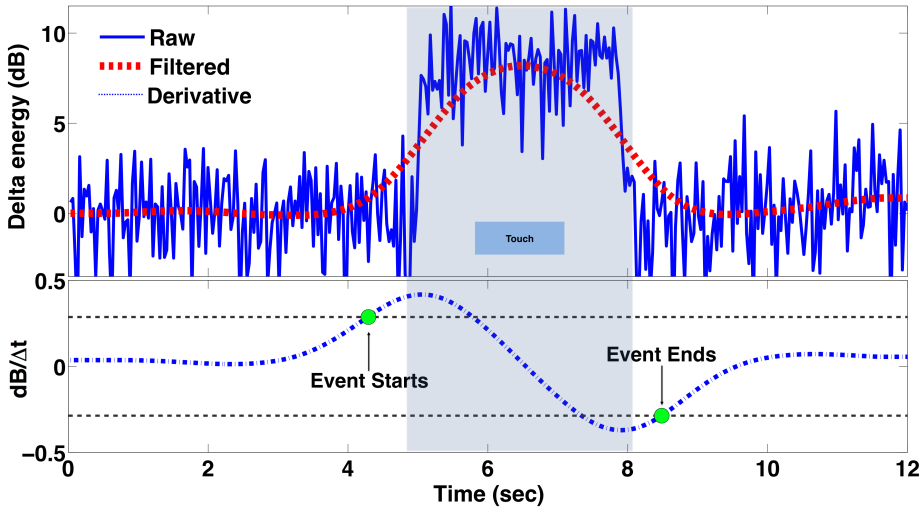

Figure 4.15: The summed amplitude (top, blue) of the EMI during a touch event, indicated by the gray highlight. The filtered amplitude (red) and derivative (bottom) are shown along with the threshold used to determine the start and end of the event.

monitors used the same background image.

The hardware used to capture the EMI and analyze used was similar to that described in section 3.1.2.

Touch Gesture Detection & Classification:

The touch EMI is produced at many harmonics of the row rate; however, for processing, a single harmonic with the highest power for each monitor was manually selected. After selecting the EMI peak, the energy of the magnitude of the FFT in the selected frequency bin with the 2 adjacent bins was summed, and the result filtered with 3 passes of a Savitzky-Golay filter with a degree of 1 and a frame length of 39 (Figure 4.15, top). To identify the end points of the performed touch gesture, a 1st-order derivative (i.e., sample-to-sample difference) of the filtered summed energy curve was taken and smoothened again with 2 more passes of the Savitzky-Golay filter (Figure 4.15, bottom). The end points are chosen to be the positions when the derivative curve exceeds a globally defined threshold selected to be the equal error rate of 4.94% using all of the data from the user study.

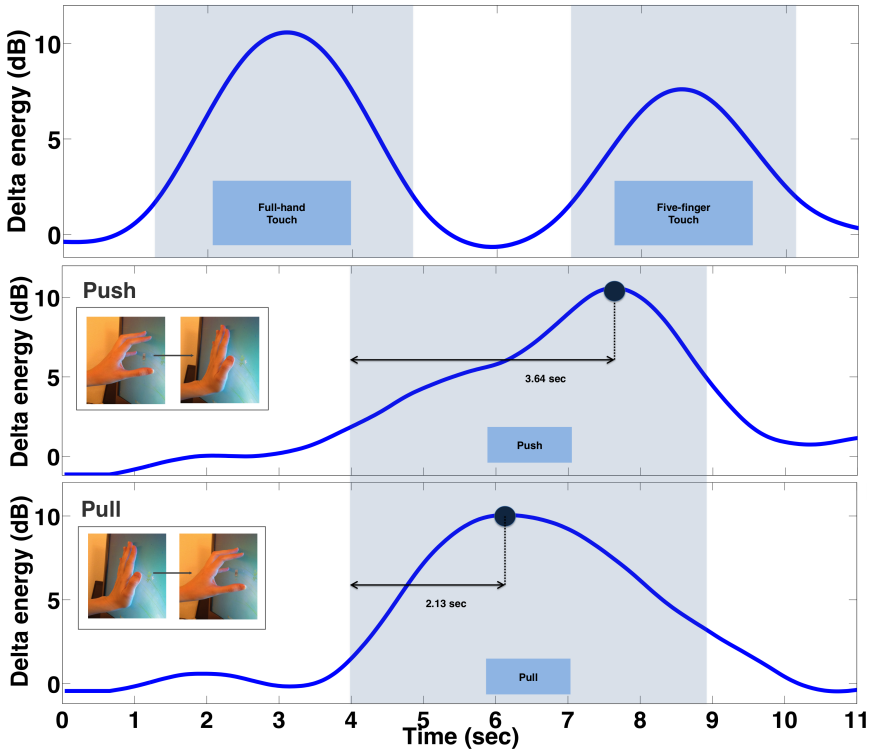

Figure 4.16: Summed energy curves representing full-hand and five-finger touches (top) as well as the push and pull gestures (bottom). The gray shaded area indicates the touch duration.

After an event is detected, features were extracted from the filtered energy curve for gesture classification. From Figure 4.16 (top), it can be seen that the amplitude change is much greater for a full-hand touch compared to a five-finger touch. To capture these amplitude differences the following three features were computed: (1) maximum amplitude: maximum value during the touch duration, (2) average amplitude, and (3) change in amplitude: difference of the average amplitude during the touch duration and the average energy in the 3 seconds prior to the touch. The push and pull gestures can be distinguished using the asymmetry in the capacitive coupling (see Figure 4.16, bottom) using the following features: (4) peak amplitude position: position in time of the point of maximum amplitude relative to the segmented touch duration, and (5) amplitude asymmetry: difference between

the average amplitude in the first and second half of the segmented touch duration.

4.2.3.3 Results and Discussion

Average detection rate of 96.4% (=9.5) was obtained, with the rate being above 97% for all touch gestures, except for hover and push. The lower detection rate for hover is due to the hand being farther from the panel, resulting in less capacitive coupling, and thus less change in EMI. The lower detection rate for the push gesture is attributed to the way most users performed the gesture. We observed that users tended to approach the screen more slowly when performing the push gesture compared to the other gestures. Since our detection algorithm is based on the derivative of the EMI, this slow approach results in fewer detections.

Using a 5-class (i.e., representing the 5 induced operating states as a result of 5 touch gestures) support vector machine (SVM) and only the first two examples of each touch gesture for each monitor and user as a trainign sample, an average classification accuracy of 68.3% (=22.5%, chance=20%) was obtained. It should also be noted that the high standard deviation in classification accuracy is almost entirely due to differences in the monitors.

These preliminary results support the hypothesis that the variations in EMI could be used to infer an LCD's internal operating states. More importantly, the EMI in this case originates from not the SMPS as with other appliances studied previously in this chapter, but from the LCD panel driver circuitry. This is an important distinction as it suggests that the approach to using variations in EMI does not have to be limited to appliances with switched mode power supplies.

4.3 Extending to Other Appliances

The fundamental concept explored in the three case studies presented above is that of inferring device's internal operation via an indirect method: by tracking the variations in the EMI they produce. The appliances analyzed in these case studies represent a small subset of both what appliances exhibit variable EMI and in what operating states could be inferred from it. Given the promising results with CFLs, TVs and LCDs, I briefly explored other appliances that may exhibit EMI that varies as a function of their internal operating

states. In this section I discuss EMI signal characteristics that I have observed for a few appliances that show promise. The preliminary signal analysis for the following appliances strongly suggest that EMI could be used to infer internal operating states; however further in-depth analysis is required.

4.3.1 Laptop Power Supply

I observed that the EMI from a laptop's charger (Macbook Pro Retina 2013) varies in strong correlation to the laptop's CPU load. In particular, the EMI in this instance demonstrated a change in the energy density. Recall that this type of EMI exhibits a change in the EMI's amplitude over time.

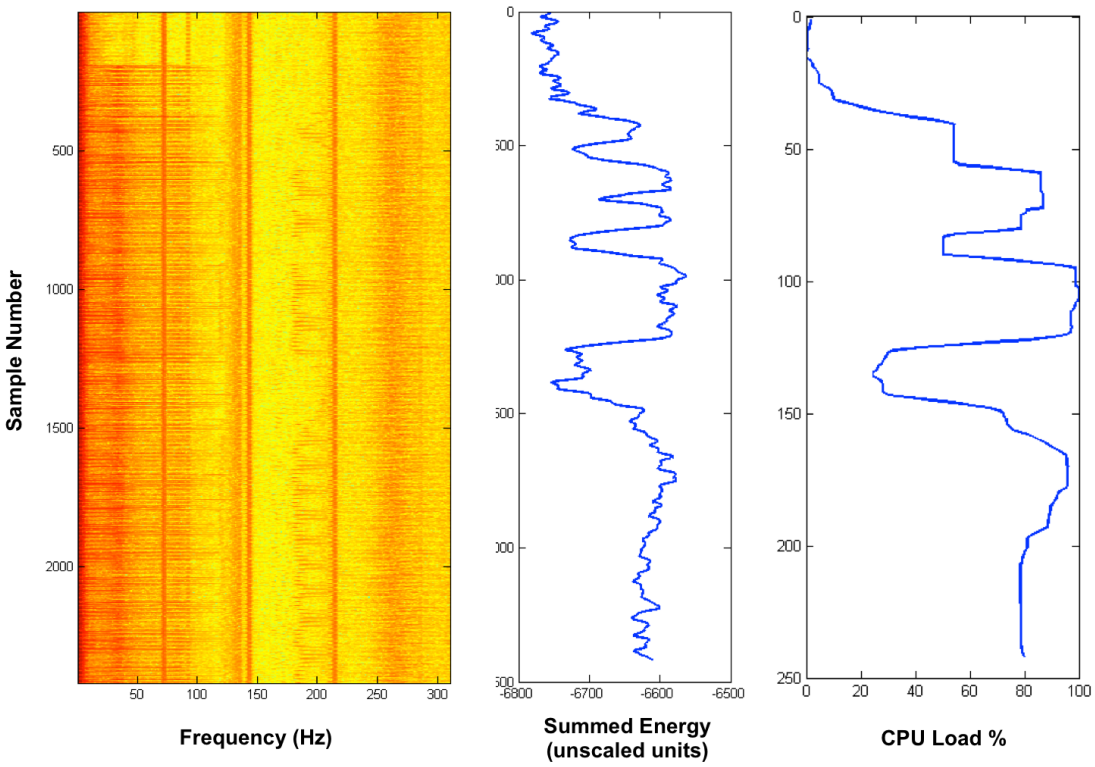

Figure 4.17: (*left*) EMI from a Macbook Pro laptop charger. (*center*) Change in energy density extracted from the EMI. *(right)* Laptop's actual CPU load. Notice how the extracted signal from EMI closely correlates to the CPU load.

Figure 4.17 shows the EMI from the laptop's charger observed when it is charging, the extracted EMI from the varying energy density using the algorithm presented in section 4.4 and the actual CPU load of the laptop. The strong similarily between when the energy density peaks to that of the CPU load peaks strongly suggests that the EMI from a laptop could be used to infer the laptop's internal operating state — in this instance the CPU behavior. I hypothesize that this varying EMI could be further used to not only profile a laptop's CPU, but also for possibly profiling the software running on the computer. In addition, from an activity recognition perspective it will be possible to infer whether a user is actively using her computer or not by keeping track of the EMI variability. Further study is required to explore these scenarios in detail.

4.3.2 Desktop PC

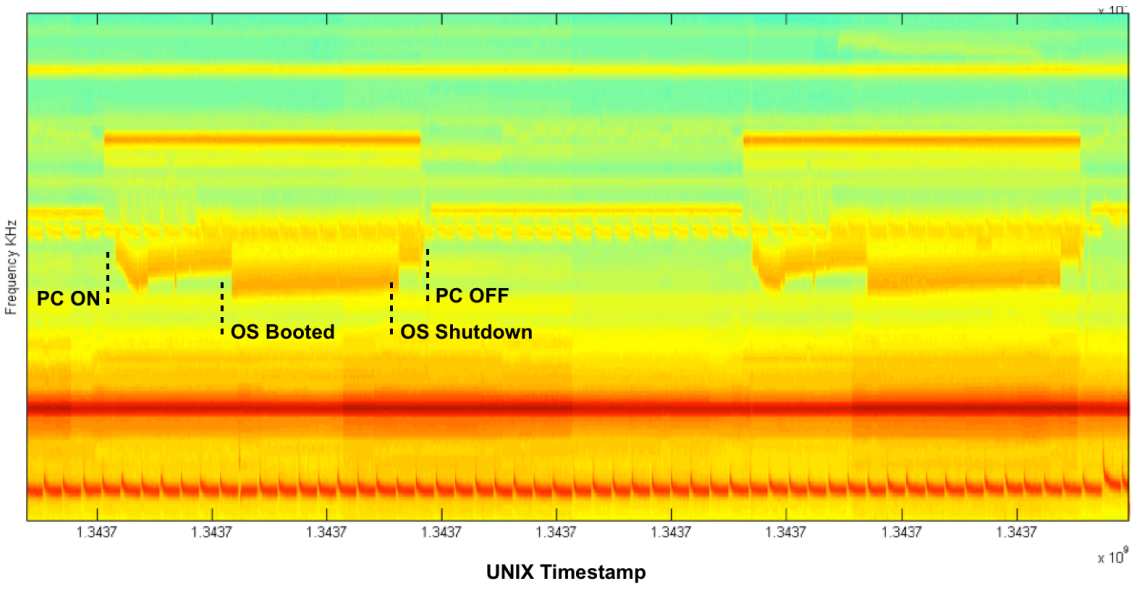

Figure 4.18: EMI of frequency-shift type observed from a desktop PC. Dotted lines indicate manual labels of when the computer was turned ON/OFF and when the OS was deemed to be fully booted. OS shutdown label indicates when the user clicked the Windows OS shutdown button.

Similar to the laptop, I observed that a desktop PC exhibits varying EMI as it is being operated. Of particular interest was the observation that when the OS was fully booted, it would exhibit EMI different from that when it is either starting up or shutting down. Figure 4.18 shows the change in EMI (frequency shift type) along with the boot state. Though the EMI variation could be attributed to the CPU load, however its correlation to the OS state suggests otherwise. My hypothesis is that the OS initializes a particular hardware subsystem, most likely the graphics card, which is responsible for the observed EMI pattern.

The implications for this observation are quite interesting: for instance, it should be possible to detect when particular hardware subsystems of a desktop PC are engaged, which can further lead to inferring user behavior. From an energy disaggregation perspective, a system taking advantage of such information would be able to give precise feedback to the user on not just the desktop PC's power consumption, but even the various subsystems. For example, imagine suggesting to user to switch to on-board graphics card instead of the PCI one when she is engaged in a non graphic intensive activity.

4.3.3 Vacuum Cleaners

Vacuum cleaners make use of high RPM motors ($\geq$ 10,000 RPM) to create an effective suction and a high rate of air flow. Usually vacuum cleaners have one motor that drives the main suction fan and through a reduction gear the agitator brush. There are variations that include a separate motor to directly drive the brush and keep the main motor cool; however compared to the main motor these run at much lower RPM.

A vacuum cleaner's main motor is often a universal motor - a type of electric motor that can run on both AC and DC power, has a high starting torque, runs at high RPM and is lightweight; making it an apt choice for portable electrical appliances like vacuum cleaners. Unlike other motor designs, universal motors are able to run at rotational frequencies that are well above the AC line frequency. In fact, they are often run at high speeds of 4000-20000 RPM. At such high RPMs, as the commutator brushes make and break electrical contact, it results in large amount of EMI at frequencies proportional to the motor RPM

(typically in the 5-20 KHz range).

Since the dominant frequency at which the EMI is produces is proportional to the motor RPM, as the vacuum cleaner's RPM changes, it results in a change in the frequency at which EMI is observed.

More specifically, as long as the load on the motor is static, the motor will run at a constant RPM. Thus, the speed setting selected by the user, if any, can be inferred. More interestingly, as a person drags the vacuum cleaner across a surface back and forth, the load on the motor changes resulting in a change in the EMI. I observed in my preliminary experiments that depending on the surface being vacuumed (wooden/hardwood floor vs. carpet) both the amount of change in EMI and the center frequency of EMI varies.

These initial observations strongly suggests that it should be possible to infer the type of surface being vacuumed after an initial machine learning training sequence. Likely such training will likely need to be per-home, however it presents an interesting opportunity to further understand and model human activity in the home. On a macro scale, collecting energy usage statistics on vacuuming wooden floors vs carpet may lead to previously undiscovered efficiency insights.

4.4 Generalized Temporal Tracking Algorithm

Unlike the one-shot difference algorithm as described in Chapter 3 used to detect appearance or disappearance of EMI from the powerline, inferring operating states requires detecting and tracking EMI variations over a period of time. On investigating the various classes of appliances, I have found that temporal variations in EMI are primary of two types: (1) center frequency of the EMI changes, and (2) energy density, i.e., total energy in the band limited region around the center frequency changes. Figure 4.19 shows examples of these two types of EMI variations. In the appliance on the left changes in the operating state cause a shift in the center frequency of the EMI. The white strip on left shows the signal obtained by tracking center frequency of the EMI on left. For device on the right, changes in operating state manifest as change in the energy density of the EMI. Tracking total energy in the band limited region over time yields the signal shown in white strip on the right.

There are two critical steps that must be performed to extract the time-varying EMI

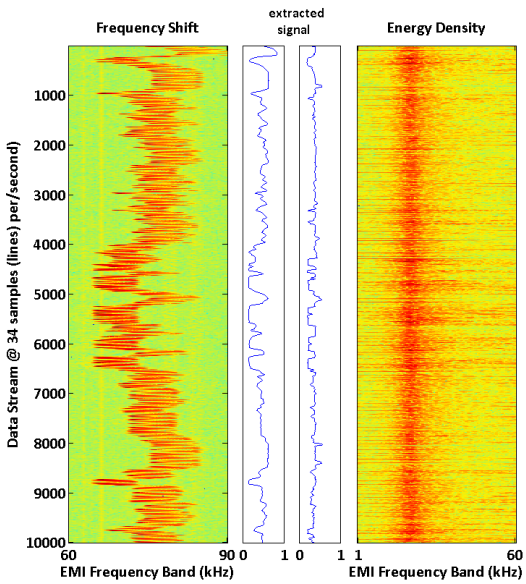

Figure 4.19: Two types of EMI variations and the extracted signals are shown. Appliance operating state changes manifest as (left) change in the center frequency of the EMI causing it to shift; and (right) as change in the energy density in a band-limited frequency region around the center frequency.

signal's pattern: (1) detect in what frequency range does the EMI signal resides, (2) in what way does the EMI signal vary with changes in appliance's state and (3) extract the EMI variation.

Although it is possible to resolve both of these steps by visual inspection of a spectrum analyzer during device operation, I developed a search procedure where given a manually or synthetically generated "target variation signal", the algorithm performs a range sweep to find EMI changes that are strongly correlated in time with the target signal. There are several ways to generate this target variation signal. For example, in case of LCDs a target signal was generated by collecting the power profile of the LCD in isolation. That is, as the LCD cycles through different brightness levels, it consumes varying amounts of power (as well as causes variation in EMI) and thus we can find the type of EMI variation by searching for a similar temporal pattern to the power profile.

As a specific implementation example of this search procedure; for my experiments with TVs, this search algorithm required FFT samples extracted from EMI over the 0 to 250 KHz frequency range collected during controlled video signal playback (N total FFT samples). Given this input the procedure analyzes subsets (windows) of the frequency spectrum and correlates their EMI trace (computed as both energy density and frequency shift) to the trace of brightness changes in the software-generated video. At the end of the sweep, the optimal 1-dimensional EMI trace is returned which produced maximal correlations to the control signal (the search is performed across all frequency windows and signal types). The functions applied are shown below in pseudo-code using the array syntax *initial-value:final-value* common to the MATLAB programming language. Smoothing is done using a 4th order Butterworth filter with a cutoff frequency of 10 KHz; decimation is done using eighth-order lowpass Chebyshev Type I filter with a cutoff frequency of $0.8*(Fs/2)/r$ where $r = 20$ is the decimation factor; standard normalization is performed to shift the domain to 0:1.

```
1  EMI:= fft(range:=0:250KHz, time:=1:N)
2  target := brightnessGradient(controlVideo)
3  for r:=1:2:250
4          for len:=10:10:60
5                  if r+len > 250
6                          continue;
7                  end
8                  rEnd := r+len;
9                  [freqEMI, densEMI] := extractEMI(r, rEnd);
10                 fCorr(r,len) := xcorr(freqEMI, target);
11                 dCorr(r,len) := xcorr(densEMI, target);
12         end
13 end
14   [optRange, sigType] := max_{r, len}(fCorr, dCorr);
15         function extractEMI(start, end){
16                 type1 := sum(EMI(start:end, 1:N);
17                 type2 := maxIndex(EMI(start:end, 1:N);
18                 densEMI := normalize(decimate(smooth(type1)))
19                 freqEMI := normalize(decimate(smooth(type2)))
20         }
```

4.5 Overview of Contributions & Implications for Energy Disaggregation Algorithms

In this chapter I presented an advancement of the EMI sensing technique presented in the previous chapter and showed how it could be extended to allow inference of an appliance's internal operating state, that is, the mode an appliance is running in. More specifically, I showed that the answer to the following question is *yes*: "Could EMI from electrical appliances be used to infer their internal operating states?". This is a significant development that can have a profound impact on complementing past energy disaggregation techniques making them more robust and accurate. To the best of my knowledge, this is the first technique that allows for fine-grained inference of operating states using a single-point sensor.

The contributions I made in this chapter are as follows:

- I presented a new sensing approach that turns ordinary compact fluorescent light bulbs into human proximity sensors. In addition, I presented evidence that CFLs can also be used for more general purpose sensing applications, such as monitoring the changes in ambient temperature, human motion detection (if a person is in range of a CFL), or as a sensor for incipient failure of CFL failure. All of these new uses are enabled by monitoring the EMI over the power line from a single point. This laid the foundation that EMI over the powerline can be used to infer the changes that the CFL's or more generally any appliance's internal oscillator is exhibiting which can then be used to infer operating states that affect the oscillator and thus indirectly the EMI produced.

- I extended the approach of studying variations in EMI to modern televisions and found that the EMI could not only be used to infer whether a TV is on or off, but also infer the screen flux matching the video content that is played. I demonstrated through in-home and laboratory experiments that given a time varying EMI pattern, it is possible to infer what video content is playing on a TV's screen.

- In addition to EMI from an appliance's power supply, I presented evidence that other oscillator in an appliance also produce EMI - such as the display driver in LCDs that also exhibit varying EMI as a result of state changes within the appliance.

- Several other appliances such as vacuum cleaners and desktop computers were also briefly studied to show that variations in EMI correlate to their operating states. These investigations strongly support my argument of using variations in the EMI to extract fine-grained operating states of electrical appliances.

- Studying a variety of appliances and their time-varying EMI, it was found that there are primarily two types of EMI variations: frequency-shifting type and energy-density type. An algorithm to extract both types of EMI was also presented.

Implication for energy: The finding that EMI sensing can be extended to estimating an appliance's operating mode or state has direct and profound implications for energy disaggregation. It redefines the state-of-the-art in single point sensing capabilities. Here I list a few implications that this additional sensing capability can have on energy disaggregation algorithms and techniques.

- First and foremost, being able to infer operating states of appliances would allow distinguishing between appliances that otherwise have similar EMI signature. For example, two appliances could have oscillators that produce overlapping EMI, however due to their unique usage models and the operating states they exhibit it would be possible to differentiate between them. This overcomes one of the key limitations of EMI sensing mentioned in Chapter 3.

- Beyond disambiguating appliances with similar EMI, the additional feature extracted from time-varying EMI could be used to make machine learning classifiers more robust. As discussed in Chapter 3, the features are currently limited to the center frequency, bandwidth and the amplitude of the EMI, however leveraging the findings presented in this chapter, new features could be extracted. Consider an additional feature of how the EMI varies over a time quantum could be added. One approach is to extract the variance in the three features over a period of few minutes to capture the unique operating characteristic of an appliance. This may not seem particularly useful at first, but compare for instance the variance for a CFL to that of a washing machine.

The variance for the former would be small while for the latter would be much higher as it goes through different wash cycles. In sum, operating state estimation can lead to more robust machine learning models for modern appliances.

- Currently research in eco-feedback has been limited to studying appliance level energy feedback. I contend that with estimation of finer grained energy usage — at the level of operating states of an appliance, eco-feedback researchers would find new insights and more engaging feedback mechanisms. For example, building on the findings from my TV case study, it is possible to detect not only how much energy a TV is consuming, but also its approximate brightness level. In such a case, an eco-feedback application could ask the user to reduce the brightness by a few points that saves energy without compromising the viewing experience. Similarly, a future eco-feedback system could recommend alternate energy saving modes for an appliance that the user may not be aware of.

- Sensing energy use down to different operating states of an appliance could allow mode-specific energy feedback. For example, users could be made aware how the majority of energy in a dishwasher is consumed by the dry/heat cycle. On a hot day, an eco-feedback system could recommend the user to rack-dry rather than run the power hungry dry cycle. Of course, this information is also highly relevant for policy makers and appliance manufactures to better design future appliances taking energy savings into consideration.

Beyond implications for energy disaggregation and eco-feedback, tracking appliance behavior over different operating states and over time can allow for early failure detection. A CFL's startup EMI pattern for instance can drastically change when the lamp is not operating optimally, often indicating deteriorating operation. I observed this in my experiments over the past years with various appliances. In particular, I observed a distinct start-up EMI pattern from a washer that failed a few weeks later. Past literature in non-invasive early motor failure detection that makes use of a magnetic pickup sensor placed closed to the motor housing (in effect sensing radiated EMI) shows similar EMI patterns that I observed

over the powerline using my EMI sensing approach [4, 5, 9, 20, 27, 29, 48, 49, 53].

Finally, as illustrated in the CFL case study, the time varying EMI can also be used for more general purpose sensing applications, such as monitoring the changes in ambient temperature, human motion detection (if a person is in range of a CFL), or as a sensor for incipient failure of CFL failure. All of these new uses are enabled by monitoring the EMI over the power line from a single point.

Chapter 5

SYSTEM FOR COLLECTING LONGITUDNAL & NATURALISTIC APPLIANCE LEVEL POWER

This chapter summarizes the data acquisition hardware and software that I built for collecting real world energy usage data. To advance the development of energy disaggregation algorithms, especially those leveraging high frequency EMI, there is a need in the research community for a comprehensive dataset. In particular, a dataset with ground-truth labeled, *real-world* energy use events. Here I describe the hardware and software platform I built to gather such a dataset. It should be noted that though I designed the hardware system described here for building energy use dataset, it is sufficiently general purpose to scale to other dataset collection efforts that researchers may come across.

This chapter is organized as follows. Section 5.1 described my motivation and reasoning behind creating the necessary tools to collect a new energy use dataset. In section 5.2.1 I give an overview of the capabilities of the new wireless data acquisition system called FireFly that I built. Section 5.2.2 onwards detail the hardware design, firmware logic and scalable cloud infrastructure that supports a FireFly node.

5.1 The Need for Labeled Appliance Level & Naturalistic EMI Dataset

A key challenge in evaluating any new sensing technique is acquiring real-world, groundtruth data. In the experiments presented in Chapter 3 and Chapter 4, groundtruth datasets were collected by visiting various homes, installing the system and *manually* labelling the electrical events that happened. There are two issues with this approach: (a) the datasets capture electrical activity that span only a period of a few hours and (b) the labelling needs to performed manually which is susceptible to human error. Exception to the length of dataset was the 6-month, one home dataset I collected in Chapter 3. Though such manually labeled and small datasets are sufficient to discover new signals and prove a new sensing technique's viability, they are limiting when developing real-world deployable systems.

In previous chapters, I have argued and scientifically shown through experimentation that EMI from appliances could be used to detect their use and then be applied to energy monitoring applications. I also presented how it could be extended to infer the operating states of appliances. However, to further study if my proposed sensing approach is deployable in real-world scenarios and how well it works, *longitudinal analysis that spans several months of naturalistic electrical activity is needed.* To this end, I built a data collection system that would allow further work to be performed on my proposed sensing technique and better understand where it excels and where it fails over long periods of time and under naturalistic scenarios.

An ideal dataset to support this research effort would include automatically generated usage labels for every appliance in the home, their power consumption profiles, harmonic content at 60Hz, impedance and capture the EMI they generate. Unfortunately such a dataset does not exist at the time of writing this dissertation. This is primarily because of lack of interest in EMI for disaggregation purposes so far, however that is quickly changing as several research groups across the world have begun experimenting in this domain. The closest dataset in spirit to what I envision as an ideal dataset is the Belkin Energy Dataset [6] that was published as part of my commercialization efforts of ElectriSense. Though this dataset includes EMI and labels for several appliances, the labels were manually collected and under staged appliance usage scenarios. In other words, it does not capture a typical home's naturalistic nuances of energy use - which I strongly believe is necessary for deeply exploring practical energy disaggregation algorithms.

To capture the detailed information for the aforementioned 'ideal dataset' to support future EMI based, non-intrusive single-point sensing research, ironically, I needed to build intrusive sensors that would attach to each and every appliance. Since these sensors capture electrical events at the level of the appliance, I could generate a dataset where usage of every appliance was labeled. Once installed, these sensors would be left running for several months to capture human activity and nuances of naturalistic energy usage.

5.2 FireFly: Ground Truth Datacollection System

In this section I present the hardware and software data acquisition system I built called the FireFly for the purposes of collecting naturalistic, groundtruth energy use data from homes. The following section describe in detail the hardware subsystem, design challenges, optimization for in-field provisioning, backend data aggregator, data cleansing and organization, choice of database, optimizing for scalable architecture and system health monitoring dashboard.

5.2.1 Overview of the FireFly System

On a high level, the FireFly system is a high-performance wireless data acquisition (DAQ) system. In simple terms DAQ hardware acts as an interface between a computer and analog signals from the outside world. Its primary function is to digitize analog signals observed with help of sensors so that a computer system can interpret them as necessary. In our case, the primary analog signals we are interested in are raw voltage and current waveforms of each appliance in a home. The key component responsible in a DAQ for this digitization of analog signals is an analog-to-digital converter (ADC). Once digitized, the samples are transferred to a computer via a communication bus. Since our measurement sensors need to be connected directly on the appliance to be monitored, a wireless communication link becomes necessary to transfer the digitized samples to a computer.

DAQs and even wireless DAQs are commonplace and can easily be bought off-the-shelf for most data collection needs, however I built a custom DAQ system because of the following features that are necessary for a successful in-field deployment and groundtruth data collection. This feature list was realized after consulting with several fellow researchers in the energy disaggregation space and my own personal experience with real-world data collection. Unfortunately, commercial DAQ systems meeting this feature set (partially in many instances) were priced higher than $1200 for an 8-channel system. Since we are interested in capturing both the current and voltage waveforms per appliance (i.e. 2-channels per appliance), each DAQ unit would allow monitoring up to 4 appliances. A typical home has about 30-45 appliances making this solution unsustainable for larger deployments and

multi-home studies.

Desired DAQ Features for Ground Truth Residential Energy Use Data Collection

- Capture 2 channels per appliance @ 1 KSPS, 12-bit resolution; namely voltage (V) and current (I) waveform.
- Zero offset between V and I samples. That is, simultaneous sampling of V/I channels for each appliance.
- Global time synchronization for very small time drift across different DAQ nodes ensuring various sensor streams could be aligned well for analysis.
- Easy in-field WiFi provisioning during deployments.
- In-field configurable ADC gain, sampling rate, conversion mode (single ended or differential) and polarity (signed or unsigned).
- Deterministic delay in global timestamping sampled data. That is, the data flow should be such that the expected delay is static and predictable to allow accurate timestamp management.
- Reliable and proven wireless communication channel for long term deployment.
- Redundant local data backup on the DAQ node in case of communication link failure.
- Scalable backend architecture for larger homes and more appliances. In particular, be able to support at least 100 DAQ nodes streaming data simultaneously.
- Backend supports centralized management of the DAQ nodes.
- Backened handles generation of dynamic configuration parameters, detects anomalies and corrects them as necessary and provides early warning for timing issues and imminent failures.

FireFly DAQ was designed to *meet all of the above features* and ensure that in-field deployments could be performed with as little overhead and manual configuration by the researcher as possible. In addition, anomaly detection and basic warning system ensures that impending failures are communicated to the researcher immediately for resolution before loss of data occurs. FireFly DAQ is a complex system and consists of several vital subsystems that each need a detailed explanation. Figure 5.1 shows a block diagram of the key subsystems. These are described in detail in the rest of this chapter in the order of how

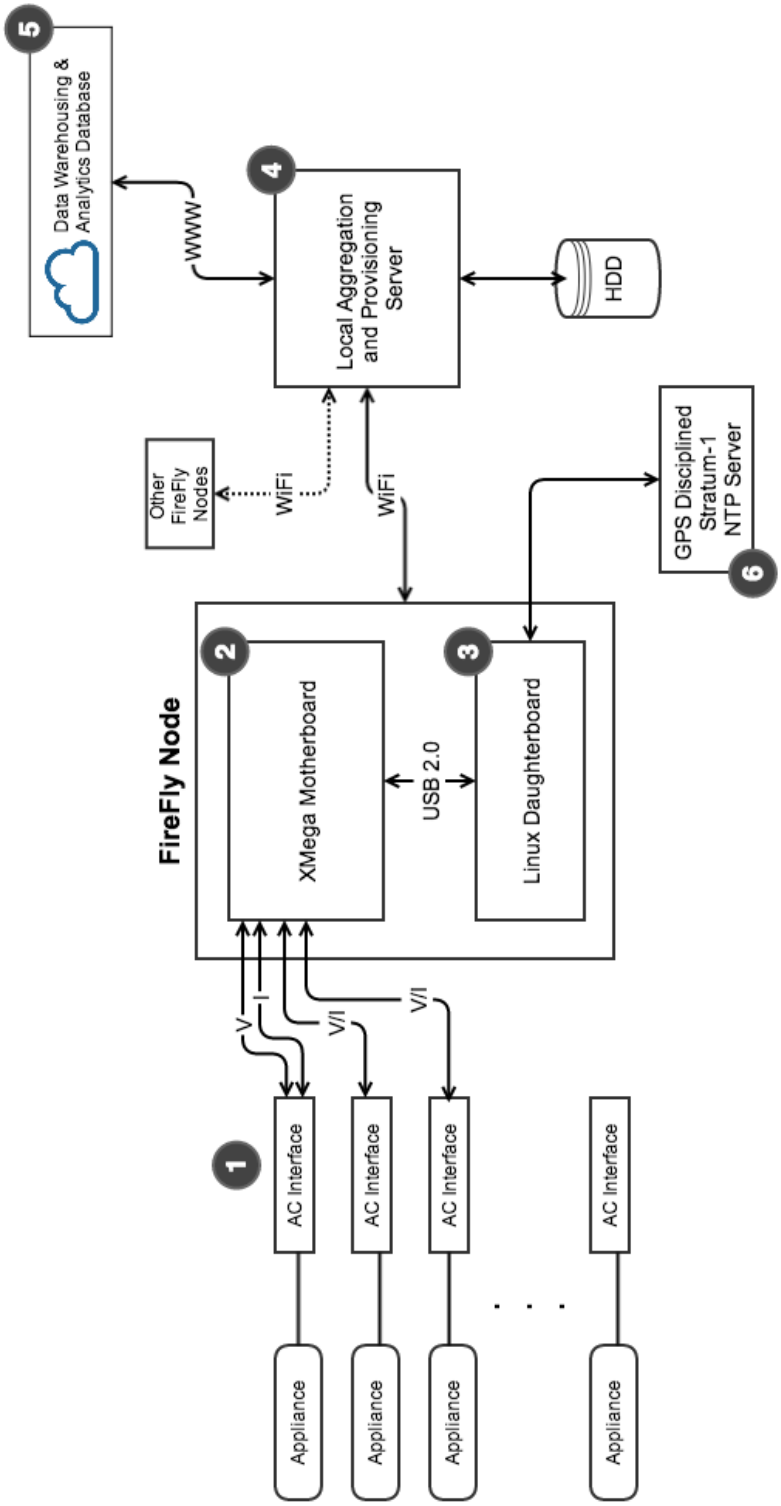

Figure 5.1: Block diagram showing the various components of the FireFLy data acquisition system.

the data flows through the system starting with analog signals that need to be measured, all the way up to their digitization, storage and analysis in the cloud. It should be noted that for brevity, only the key subsystems and modules customized for application to energy monitoring are described. General purpose embedded systems and software modules have been omitted, however are available for analysis by the reader as an open source project repository (please contact me to gain access).

5.2.2 Block 1: AC Interface

Fundamentally, we are interested in recording the amount of power consumed by each appliance in a home for the purposes of generating groundtruth appliance usage labels. However, as discussed in Chapter 2 many disaggregation algorithms in published literature utilize current harmonics as features. Though these harmonics are derived from a home's aggregate current draw, it could be beneficial for future NIALM algorithm development to study isolated harmonic content and EMI of each appliance and how that manifests at the single sensing point location.

Since collecting groundtruth labels necessitates putting a sensor on every appliance, with marginal extra effort, we can instead capture the raw current and voltage waveforms; an option that was the case for the system described here. Of course, these waveforms could then be used to compute real power as a software processing step and consequently generate appliance usage labels. It should be noted that for applications where only the power consumption is desired, a power measurement sensor could be connected to the FireFly DAQ. This also has the advantage of reducing the number of channels (and hence bandwidth) by half.

AC signals from the appliances are too dangerous and can be noisy to measure directly. Thus signal conditioning circuitry is needed to manipulate it into a form that is suitable for input into the FireFly's ADC. In addition to converting the two signals of interest: voltage (V) and current (I) to voltage levels appropriate for the ADC, they are also galvanically isolated for safety and to prevent accidental electric shock to human users. Both of the voltage conversion and safety functions are performed by an AC interface circuitry as shown

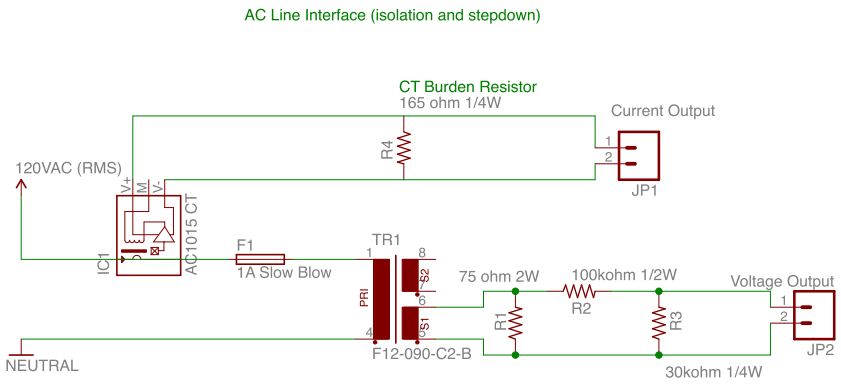

Figure 5.2: Schematic of the plug-in AC line interface for galvanically isolated and stepped down current and voltage waveforms suitable for input into FireFly DAQ ADC. An AC line interface acts as a pass-through device that is connected to each appliance being monitored.

in Figure 5.2.

First and foremost, the AC interface circuitry consists of a 2 amp fuse for safety purposes. The fuse is of "slow blow" type to ensure that inrush current or a sudden surge of power does not trip the fuse. Next, to isolate the voltage waveform and reduce it to an optimal voltage level that matches the FireFly's full scale (FS) 2.5V ADC, a step down transformer is used. A burden resistor is placed on the secondary of the transformer such that it operates in its rated loaded state and results in a nominal 115V:6.3V RMS step down. This is further reduced with help of a voltage divider such that the peak voltage does not exceed 2.5V.

To isolate and measure AC current, an industry standard current transformer (CT) is employed. The CT is placed around the 'live' wire connecting the AC mains to the appliance being monitored. Similar to any transformer, a CT has a primary winding, a

Burden Resistor (Ohms)	**Max. Current at 2.5V FS**	**Current Resolution (mA)**
100	77.50	18.92
120	64.58	15.77
145	53.45	13.05
165	46.97	11.47
180	43.06	10.51
191	40.58	9.91
200	38.75	9.46
250	31.00	7.57
330	23.48	5.73

Table 5.1: Choosing optimal burden resistor for the CT. Maximum current at 2.5V FS and resolution at 12-bit are shown for each burden resistor.

magnetic core and a secondary winding. The primary winding is comprised of the wire around which it is placed and through which the current flowing is to be measured. Thus the primary winding is of 'single turn'. The secondary winding is galvanically isolated from the current carrying conductor and produces a current proportional to the current carrying wire, which can then be measured. Since our ADC measures voltage, a burden resistor is placed across the output of the CT. This resistor was selected to optimize for the maximum current to senses at the 2.5V FS and the desired resolution at 12-bits. Table 5.1 lists maximum current and resolution for various burden resistors modeled using the CT's specific transformer ratio. Taking into account typical home appliances' energy consumption and algorithm requirements in literature, it was decided to use a burden resistor of 165 Ω. As described later, a programmable gain amplifier (PGA) precedes the input of ADC. This allows software selectable gains of 0.5X, 2X, 4X, 8X, 16X, 32X and 64X, thus allowing the system to adapt to appliances with much lower current consumption without altering the burden resistor or the circuit.

The CT and voltage step-down transformer, together allow us to measure the voltage and current waveforms of an appliance in a safe and suitable manner. During installation, the target appliance to be monitored is plugged into the AC interface with the latter plugged into

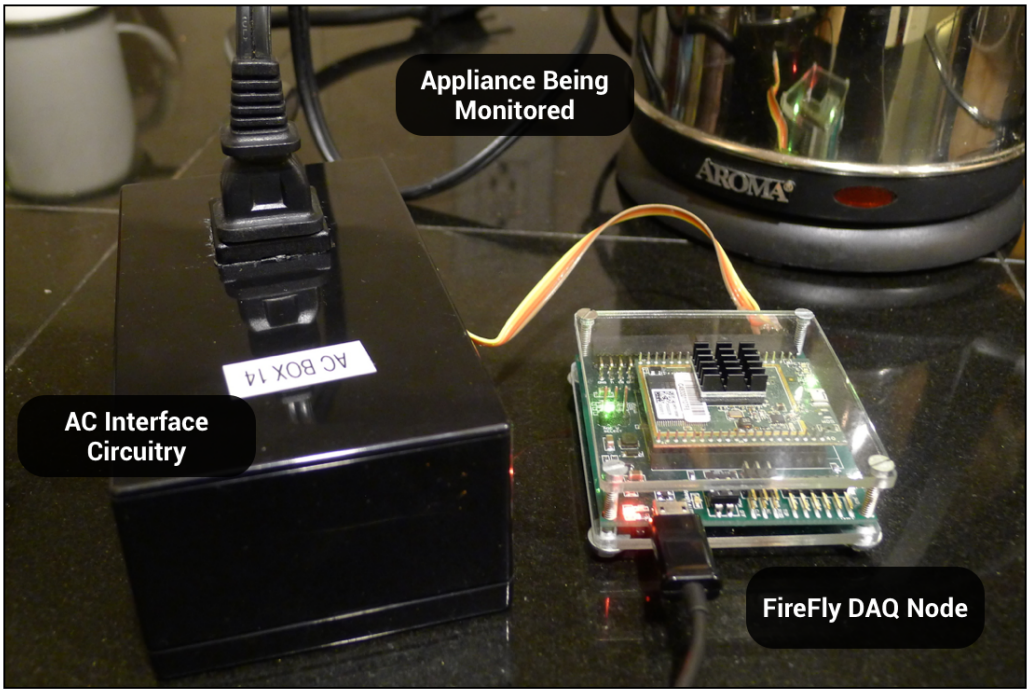

Figure 5.3: The black box on the left consists of the the AC interface circuitry that acts as a pass-through circuit for the electric Kettle plugged into it. The voltage and current waveforms for the kettle are captured by the FireFly unit connected to the AC interface via a 4-conductor ribbon cable.

the AC mains outlet. The interface circuit acts as a transparent pass-through that powers the appliance without effecting it in anyway while allowing simultaneous measurement of the V and I waveforms. Figure 5.3 shows an example installation with an electric Kettle being monitored.

5.2.3 Block 2: Firefly Node Motherboard

To allow for flexibility in choice of the communication channel and CPU resources available for on-node signal conditioning and processing, the FireFly node was designed with a modular motherboard and daughterboard architecture. The motherboard provide the core DAQ functionality with a high-rate ADC, microcontroller, switching power supply and precision

voltage reference. It also has a built-in USB 2.0 bus to allow direct connection to a PC for data logging or to support the daughterboard. In the current instantiation of FireFly DAQ node, any application specific processing of digitized samples is handled by a comparatively more powerful CPU on the daughterboard while real-time sampling and bit-level hardware control is managed my the motherboard microcontroller.

Choice of Microcontroller:

Figure 5.4 shows the schematic for the motherboard and in particular highlights the circuitry designed to support a microcontroller. This microcontroller is the brains of the DAQ node, and is crucial for timing control, low-level hardware management, sampling data and managing the communication link to transfer sampled data to a logging computer. The choice of the microcontroller used on the DAQ node was crucial. To meet the specification laid out in Section 5.2.1, I investigated a large number of microcontrollers ranging from 8-bit, low power ones to larger and more powerful 32-bit microcontrollers. In particular, I considered the Atmel 32-bit UC3 series, Microchip Propeller series and TI's C2000 series and performed a detailed comparison of parameters while optimizing for overall functionality and least part count.

Atmel ATXMega192A3U was particularly chosen because of its high performance *dual ADCs*, which ensured that two channels can be sampled simultaneously with guaranteed zero time offset between the samples. The architecture of the microcontroller allows simultaneous hardware triggering of both ADCs with the silicon signal path being identical. This was a key requirement for accurate V/I waveform capture ensuring that the capture hardware does not introduce an offset which would later result in incorrect power factor and real power computations. Combined, the two ADCs offer 16 single-ended or 8 differential channels, 12-bit conversion resolution at a maximum of 2000 ksps. 16KB of SRAM ensures that sufficient data samples could be buffered before being transferred out using a full-speed on-board USB controller.

In addition, this microcontroller features hardware direct memory access (DMA) controllers further adding to the list of favorable features to realize a high performance DAQ system. DMA is a feature that allows certain hardware subsystems of the microcontroller

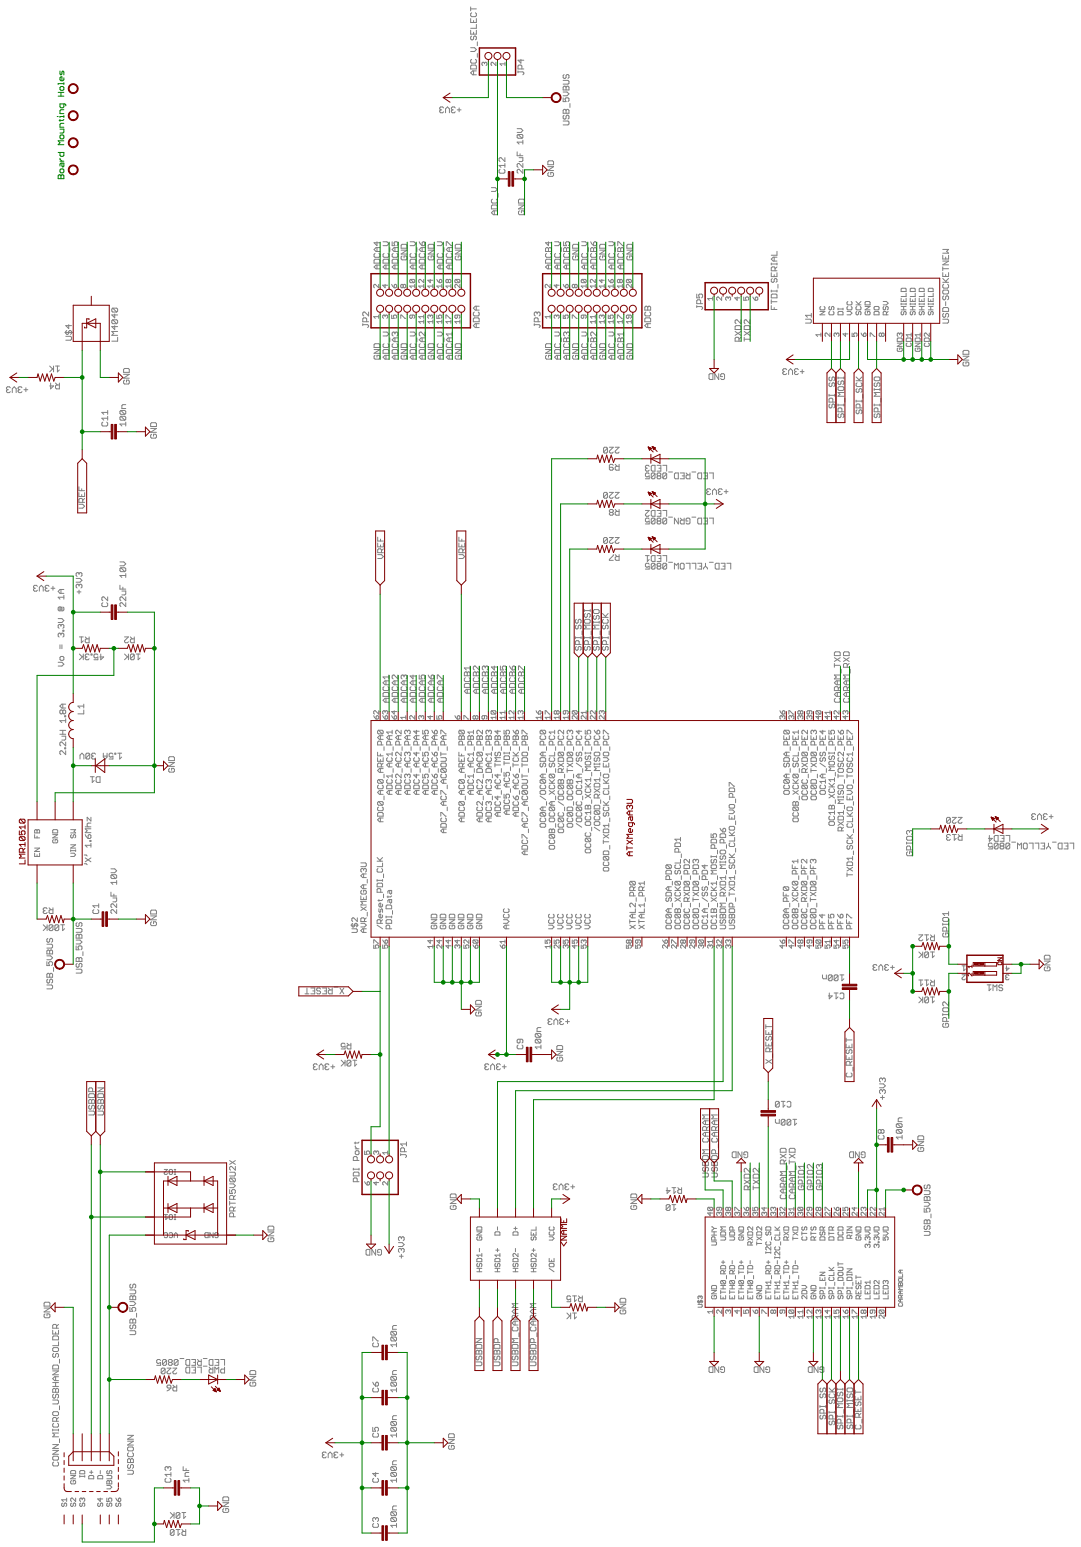

Figure 5.4: Schematic for the FireFly Node motherboard.

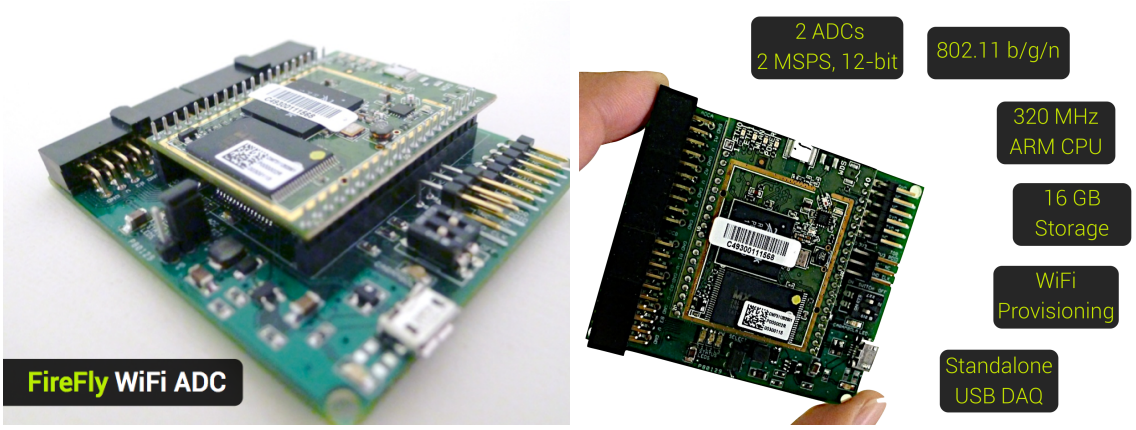

Figure 5.5: Firefly DAQ node with the daughterboard plugged into the base motherboard.

to directly access the memory without intervention from the CPU. Specifically, in our case this means that samples from the ADC's output registers could be directly transfered to the RAM without spending any CPU cycles. Since the primary function of the microcontroller in this setup is to sample analog signals and transfer digitized data to USB, a DMA allows this to happen automatically (and in background) leaving the CPU free for housekeeping tasks such as timing control, error checking and orchestrating the DMA controllers. It should be noted that the USB on this particular microcontroller is also a hardware implementation, further freeing up CPU cycles from performing repetitive IO tasks. In sum, the choice of the microcontroller was driven by availability of dual ADCs, DMA controllers, large enough RAM for double buffering and a high speed communications bus such as USB. Figure 5.5 shows the completed FireFly node with the WiFi daughterboard plugged into the base motherboard.

ADC Configuration and Precision Caveats:

Configuring the dual ADCs on an Atmel ATxMega192A3U (referred to as XMega hereon) is straightforward for most applications however it poses several non-trivial challenges when guarantee on high precision is required. This section described how I obtained 12-bit resolution on the XMega series using the signature row values and apt reference

selection.

Our selected XMega microcontroller consists of two 12-bit successive approximation register (SAR) type ADCs. These type of ADCs are often the choice for medium to high resolution applications with sampling rates under about 5 megasamples per second, provide a good balance between power consumption and form factor and with a resolution ranging between 8 to 16 bits. Thus, microcontrollers often make use of SAR ADCs and with XMega it provides more than sufficient performance for the intended application. XMega's ADCs can be configured in several ways to match their intended use including:

- Input mode (single-ended vs. differential)
- Resolution (8-bit or 12-bit)
- Front end programmable gain stage before ADC conversion
- Variety of reference voltages
- Conversion Mode (signed or unsigned)

Before the ADC can be used all of the above listed parameters needs to be initialized. This is accomplished by downloading a configuration structure from the local provisioning and aggregation server (see Section 5.2.5). For each node a specific configuration is specified that can either be kept static or be dynamically updated during operation to adapt to changing installation environment and appliances. Since our AC interface circuit outputs differential signals for each of V and I, the input mode is set to differential and conversion mode to signed (voltage can swing from -2.5V to +2.5V). The resolution is set to 12-bit at the cost of increased bandwidth and data storage requirements, but ensures NIALM algorithm development without compromise.

In my preliminary deployments, I found that none of the parameters except the ADC gain need to be updated in-field. This is because, ideally we want the gain to be set in such a way that the incoming analog signal's peak voltage is as close to the ADC's FS 2.5V to allow utilization of the full 12-bit quantization. For voltage waveforms this is not an issue as it does not vary significantly, however the current is both dependent on the appliance and the mode it operates in. The researcher could hand-pick the gain for each appliance during installation and also anticipate the peak current it may draw, however this is tedious and error prone. To counter this and automate the process, an algorithm on the local aggregation

and provisioning server was implemented which automatically adjusts the gains for optimal performance. When the server detects that the signal is about to clip or has clipped, it pushes a new ADC configuration to the DAQ node with updated ADC channel gains. The microcontroller on receiving the new configuration structure immediately applies it. It should be noted that updating the ADC configuration parameter while it is running requires pausing and flushing the pending samples pipeline which results in dropping 11-samples. For most appliances the gain never changes once set by the server, however appliances that are operated in different modes from time to time may require configuration updates.

Once initialized, the XMega ADC can be triggered to perform a conversion as a result of many events, including analog comparator level change, timer counter overflows, DMA transfer requests and explicit software triggers. It can also be configured to run in a 'freerunning' mode where a new conversion begins as soon as the previous is complete. Once triggered, a conversion begins and within a few clock cycles (due to propagation delay) the result of the conversion is placed in a 16-bit register. The end of conversion also fires an event through the XMega's event system that allows us to have the DMA controllers take over the task of copying the result registers rather than the CPU.

To get the highest precision possible (11-bit in differential mode and 12-bit otherwise), certain caveats need to be kept in mind steps taken to counter them. First, the ADC calibration from the NVM production signature row should is loaded into the ADC module to null out any hardware inaccuracies. Second, to remove any positive offset, the ADC channel is first internally connected to ground and an average of 1000 samples used as an offset to consequently apply to any samples taken. This is done once after the configuration step is performed. Next, the ADC clock frequency (and hence the maximum sampling rate) effects the precision. For XMega, the manufacturer recommends no greater than a 200 KHz clock rate for maximum 12-bit precision. Since our desired sampling rate is 1 KHz per channel, I chose to keep the ADC module clock at 50 KHz.

It is important to note that for keeping the noise low and frequency resolution high, it is crucial that the ADC samples at a consistent rate. To ensure this, I trigger the ADC conversion using a timer counter overflow event with a rate set to a multiple of the ADC clock. Any other mode or frequency will have a 1/2 ADC clock jitter and will result in

frequency modulation of the signal. Finally, I make use of a precision 2.5V voltage reference (see schematic in Figure 5.4, component LM4040). The TI LM4040 used on the FireFly node provides very stable performance over large temperature range and is stable to typically 100 ppm/C. These steps ensure high accuracy and precision, which I confirmed with a calibrated test signal source.

DMA Chaining:

One of the most critical requirements of a data acquisition system such as the FireFly system is meeting *hard real-time deadlines*. In other words, it must maintain strict timing while buffering and processing data to ensure no buffer overflows occur or samples dropped as a result of CPU resource unavailability. Missing a timing deadline constitutes as a system failure as we are designing an infrastructure to collect "groundtruth" data. Ensuring strict timing therefore has been a high priority design goal when I began development on FireFly.

To meet this stringent timing constraint, I have implemented a ping-pong buffer managed by chained DMA controllers (see Figure 5.6). To appreciate this design choice and form an intuition, consider a scenario where a CPU collects N number of samples, buffers it and sends it off-chip using a standard interface. In this scenario the CPU needs to transfer the data off-chip before the next sample is ready; if it fails then previous samples will be overwritten causing data loss. To counter this, a double-buffer or ping-pong buffer approach is used where while the samples are collect in the first buffer, the CPU uses the time between each sample to handle the transfer of data in second buffer and vice versa. This optimizations results in the CPU having twice the amount of time to process data transfer than a single buffer approach.

On a resource constrained microcontroller, handling the data movement from and to ping-pong buffers could consume a large amount of CPU time leaving insufficient resources for other tasks. Instead, data movement on-chip can be done efficiently using special data transfer resources such as DMA controllers. The CPU is only involved in setting up the DMA controllers and once done, the data movement can happen in the background leaving the CPU free. In case of FireFly, I have implemented a DMA chaining mechanism which ensures that each buffer of the two ping-pong buffers of each ADC (total 4 buffers and DMA

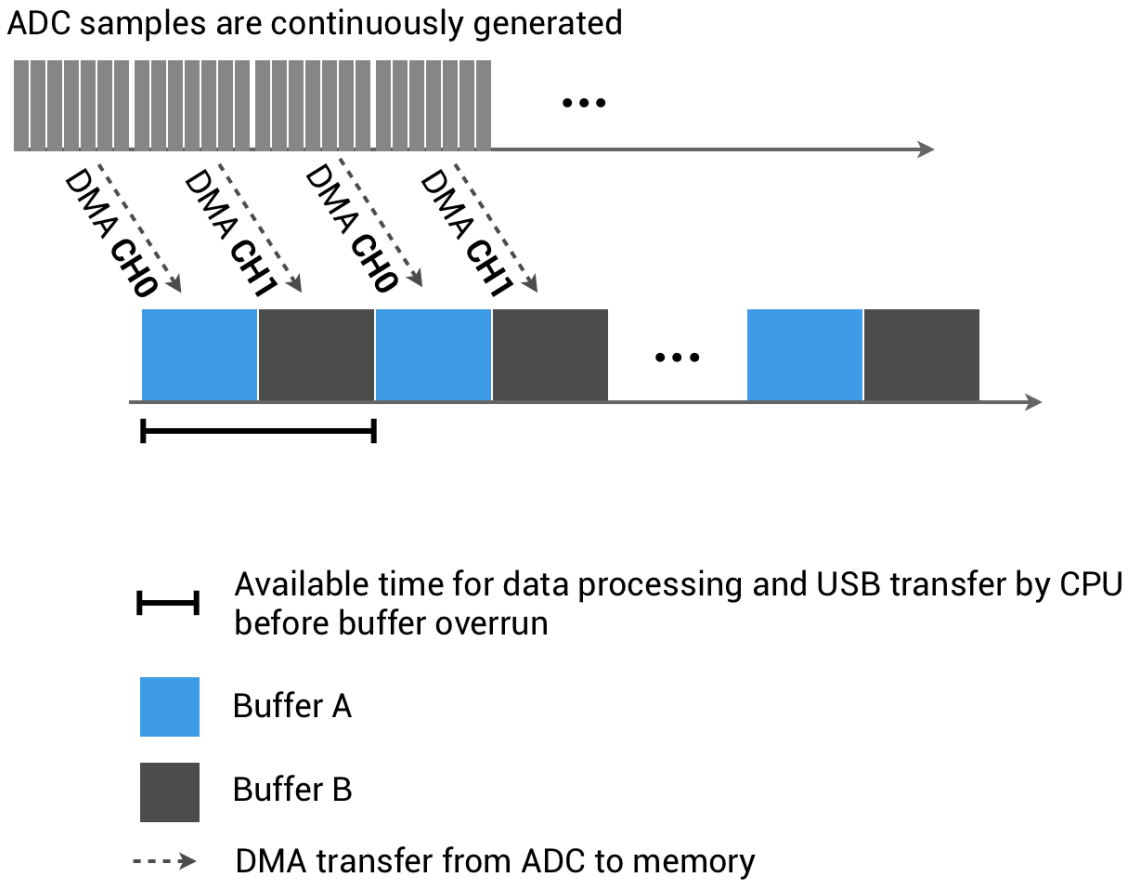

Figure 5.6: Strict timing and hard real-time deadlines are maintained on a FireFly DAQ by doubling the amount of time the CPU has to process each buffer using a ping-pong buffer approach. To further minimize CPU load, DMA chaining is employed with a dedicated DMA channel for each buffer bank that hands over to the alternate DMA channel when done requiring no reconfiguration of DMA controllers while operational.

channels) has a dedicated DMA controller managing data movement. This has proved to be quite efficient resulting in a very low CPU usage and high data throughput.

Optimization of USB Throughput:

Once the data is buffered into one of the ping-pong buffers, it needs to be efficiently transferred off-chip using the on-board USB. USB supports many different transfer mechanisms of which the 'Bulk Transfer' is implemented as it is most suitable for large bursty data transfers. Bulk transfer mechanism provides error detection via CRC with guaranteed delivery, however no guarantee of bandwidth or minimum latency. The latter two are not an issue since in case of FireFly, the USB bus is not shared with other transport mechanisms. That is, it is enumerated as *only* a bulk transfer device and thus does not contend for prioritization.

Bulk transfer supports several different bulk packet sizes, namely 8, 16, 32 and 64 bytes long, with the XMega implementing a 64 byte packet. For ideal utilization of the underlying protocol, the XMega USB driver waits until at least 64 bytes have been buffered for off-chip transfer before invoking the bulk transfer. On the surface it may seem that the matching the ping-pong buffer size to 64 bytes would work well. It indeed does, however the overhead of invoking USB transfer is sufficiently large that a larger data buffer could perform better or as well. To evaluate what the apt buffer size is that results in the maximum throughput, low CPU usage and has the *least jitter*, I performed a detailed experiment measuring the timing for USB calls and DMA completion calls.

To ensure that my timing experiment has minimal impact of the true timing of the system, I chose to introduce *single-instruction* debug commands in the firmware that would toggle a pin on DMA completion and USB transfer initiation calls. I then used a high-speed logic analyzer to capture the timing. This was done for several different buffer sizes and each of the Low-speed and high-speed USB modes supported by the XMega. Table 5.2.3 details the measurements for readers interested in studying the trade-off between various options. The key takeaway from the experiment was that a buffer size of 322 bytes and USB mode of high speed yielded optimal results. The jitter between each USB transfer was least compared to other modes and zero buffer overruns occurred. A larger buffer size also means

that the CPU gets a large amount of time between each buffer (4 seconds with 322 bytes buffer at 1KSPS) to complete other housekeeping tasks such as updating configurations received from upstream server, compute CRC and manage the user interface event loop.

Astute readers will notice that a buffer size of 322 bytes performed slightly better than 320 bytes *on average*. This can be explained by the fact that though 320 bytes (multiple of 64 byte bulk packet size) should have performed best, it lags slightly behind 322 bytes packet because the extra 2 bytes in the latter force the USB driver to immediately initiate a bulk transfer. Thus, on a single instance basis 320 byte packet is still the correct choice, however taking into account the fact that we buffer for a loonger time than it takes to perform the USB transfer, a forced USB transfer (at the cost of under-utilizing the 64 byte packet) performs better on average to reduce the jitter.

Firmware Event Loop:

Once ADCs and DMAs are initialized the firmware enters an event loop where it only responds to either events from the DMA controllers or interrupts pertaining to configuration updates. The latter are command and control messages that are received on a high-priority backchannel (TTL232) allowing control of the firmware even in an event of an error or for forcing a processor reset. Figure 5.7 depicts the various functions the firmware performs, some of which have been discussed in detail in this section.

5.2.4 Block 3: WiFi Linux Daughterboard

As the FireFly node motherboard digitizes the various signals and offloads them via USB, they are received and processed on the daughterboard. The daughterboard is a plug-in module that features a much more powerful CPU compared to the microcontroller on the host board and performs higher level functions such as application-specific data conditioning, compression, local backup, timestamping and remote backup.

For the hardware instance shown in Figure 5.5, the daughterboard is a pre-fabricated RaLink System on a Chip (SoC) based Linux computer called the Carambola [1]. The Carambola module features a 320MHz processor, 8 MB flash, 32 MB RAM and 802.11N

Configuration	USB Write A					DMA A Toggle					DMA B Toggle					Violations
	Min	**Max**	**Avg**	**Std**	**Spread**	**Min**	**Max**	**Avg**	**Std**	**Spread**	**Min**	**Max**	**Avg**	**Std**	**Spread**	
HS_66	0.0228	0.0604	0.0249	0.0058	0.0376	0.7644	0.8361	0.8005	0.0054	0.0718	0.7644	0.8369	0.8006	0.0055	0.0724	0
HS_64	0.0223	0.0599	0.0255	0.0075	0.0376	0.7641	0.8365	0.8004	0.0063	0.0724	0.7639	0.8360	0.8002	0.0064	0.0721	0
LS_64	0.0223	1.0020	0.0371	0.0388	0.9798	0.7664	0.8348	0.8003	0.0068	0.0684	0.7679	0.8318	0.8003	0.0064	0.0639	2
HS_306	0.0794	0.1204	0.0859	0.0092	0.0409	3.7939	3.8159	3.8026	0.0018	0.0220	3.7934	3.8141	3.8026	0.0018	0.0206	0
HS_320	0.0826	0.1154	0.0895	0.0090	0.0328	3.9969	4.0134	4.0010	0.0019	0.0164	3.9966	4.0133	4.0011	0.0019	0.0166	0
HS_322	0.0886	0.1206	0.0944	0.0062	**0.0320**	3.9966	4.0133	4.0013	0.0017	0.0167	3.9962	4.0143	4.0013	0.0017	0.0181	**0**
LS_66	0.0283	0.5699	0.1625	0.1605	0.5416	0.7595	0.8409	0.8002	0.0097	0.0814	0.7616	1.6004	0.8006	0.0206	0.8388	662
LS_320	0.3414	2.2579	0.8609	0.3307	1.9165	3.9978	4.0069	4.0014	0.0013	0.0091	3.9972	4.0284	4.0015	0.0020	0.0312	1
LS_322	0.4059	1.4329	0.9866	0.2962	1.0269	3.9987	4.0149	4.0030	0.0017	0.0163	3.9988	4.0129	4.0030	0.0016	0.0141	0

Table 5.2: Timing delay between successive USB transfer function calls and each DMA controller completion. All timings are presented in seconds. Configuration is formatted as LS_X for low-speed USB and HS_X for high-speed USB, where X is the size of buffer in bytes. Notice that HS_322 performed optimally with low average USB write time and least jitter (see spread and standard deviation).

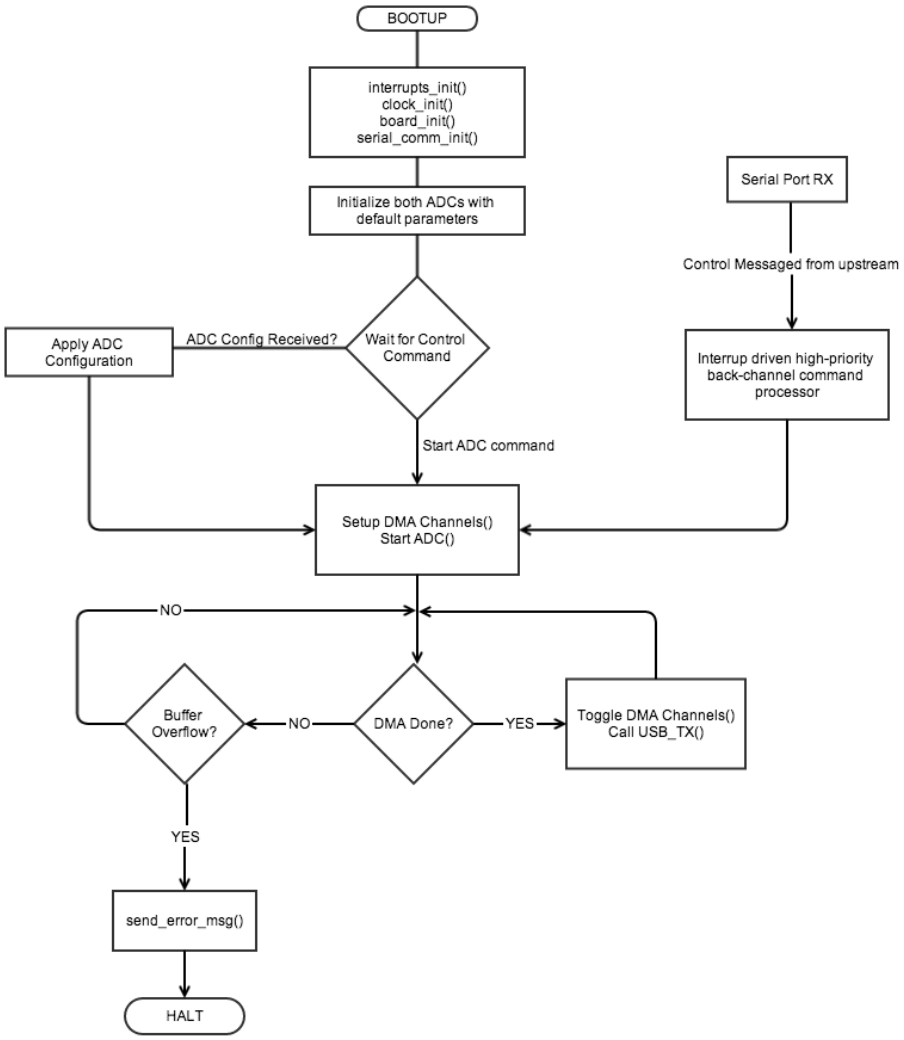

Figure 5.7: High level flowchart depicting the various functions the microcontroller firmware performs.

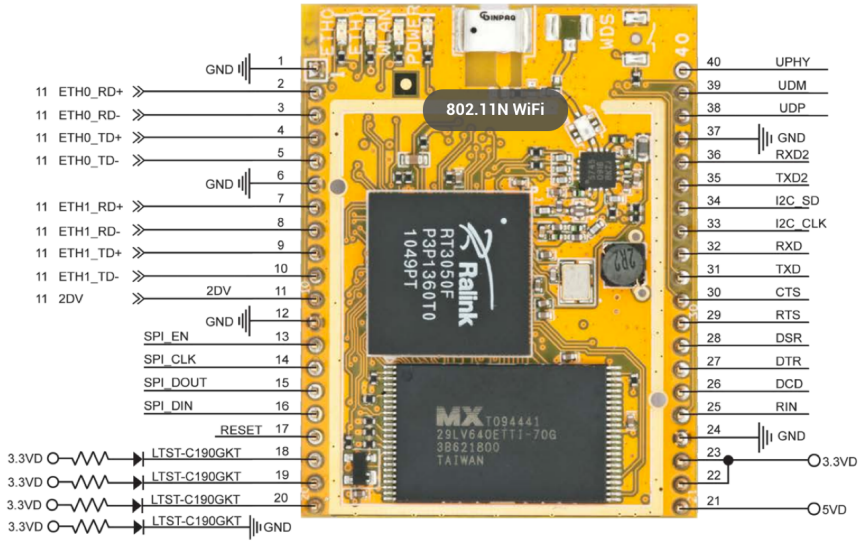

Figure 5.8: The daughterboard on a FireFly DAQ node is a micro Linux computer running the popular OpenWRT distribution. Note the various interfaces such as Serial Peripheral Interface (SPI), Ethernet, USB, TTL232, GPIOs and WiFi that make it a versatile platform. This board communicates with the XMega motherboard using USB and TTL232. *Image courtesy of 8devices.com.*

2.4 GHz radio with on-board chip antenna. It supports the popular OpenWRT Linux distribution often used on routers. The choice to use Carambola was driven by its low cost, high-speed WiFi capability, wide range of communication buses and support for a mature Linux distribution allowing flexible development and long-term support (see Figure 5.8 for list and pinout of various interfaces). A high-priority process runs on the daughterboard that apart from moving data over the wireless link to an aggregation server for storage also performs vital functions as described below.

Persistent Serial Back-channel with the Microcontroller:

First and foremost the Linux process maintains a persistent TTL232 serial connection with the the motherboard microcontroller. On the microcontroller side as well, it is made sure that this connection is live and active before starting the main sampling loop. This

back-channel serves several purposes.

It is responsible for bootstrapping the USB communication link between itself and the microcontroller that would later be used to transfer all the sampled data. Since the Linux process has connectivity via its WiFi link to the local aggregation and provisioning, this channel is also used to push the ADC configuration packets downstream. On the XMega this serial backchannel is interrupt driven and has a high priority so that control messages like dynamic updates to the ADC (like updating amplifier gains) can be performed irrespective of what state the microcontroller is in. Another use of this serial link is to update the firmware using a serial bootloader on the microcontroller. This allows the researchers to push firmware upgrades remotely and update nodes deployed in the field with a small downtime (60 seconds).

I have also utilized this serial link to implement a debugging interface during development. One of the may features of this interface allows error notification, such as buffer overrun to cascade up from the microcontroller all the way up to the Internet connected local aggregation server.

Data Conditioning, Global Timestamping and Uploading:

The Linux process is setup in a collector-relay architecture. Both the collector and the relay are lightweight processes (or threads) that continuously run until explicitly shutdown by the user. A watchdog process makes sure that if either the collector or relay threads die, they are immediately re-spawned. The two threads leverage shared memory to communicate information back and forth and to transfer the large amounts of streaming data samples from the collector to the relay thread.

The collector thread is the FireFly microcontroller facing process, i.e., its job is to communicate with the low-level processor over USB, retrieve data as soon as it is ready, condition and packetize the data, apply a global timestamp and prepare it to be transferred to the relay thread. Similar to the microcontroller, the collector threads manages a ping-pong buffer into which data over USB is buffered. Once the buffer is full, the samples are put in a packet with header that includes a CRC of the data and the global timestamp at which the samples were taken.

The global timestamp is derived from the system time, which is set as soon as the Linux boots up from a local high-precision Network Time Protocol (NTP) server (detailed later in this chapter). Having accurate global timestamps guarantees that data from different FireFly nodes could be correlated in time. This is important from energy disaggregation research purposes as a researchers we often would want to correlate appliance's power usage to the whole home's power consumption as well as other nearby appliances. For example, stove usage in kitchen would often be preceded by the kitchen lights being operated. It should be noted, that avoiding jitter due to non-deterministic delays in the Linux USB data path (downside of a non-realtime OS) is of paramount importance. To this end, I designed the system such that the timestamp for each packet is computed in the collector thread with a separate process syncing a software clock between the collector and the microcontroller using the low-latency TTL232 backchannel. This is to make sure that the clock drift never exceeds 1-ADC sample time. Once the packet has been created and timestamped, it is transferred over to the relay thread using a shared memory.

The relay thread is the Internet facing or local aggregation server facing process which primarily deals with higher-level protocols such as HTTP to upload the data packets to the server. As soon as a packet is ready for relay, i.e., consumed by the relay from the collector thread, it is queued for delivery to the aggregation server. The relay continuously polls the queue and makes HTTP POST requests to the server with the data packet as the payload. If successful, the packet is de-queued and deleted. However, if the upload fails even after several retries, the data is immediately written to a persistent local storage — a SD memory card mounted as an EXT3 partition. From hereon, the relay thread enumerates through both the data packets on the SD card and in-memory and constantly tries to upload it.

Local storage in form of a SD card is an important feature of the FireFly DAQ system that sets it apart from existing commercial systems available. This model guarantees that network failures would not result in loss of data. In my preliminary field deployments, I found that a 16 GB SD card would hold up to 15 days of uncompressed data (6 differential channels sampled at 1ksps, 12-bit). This is more than sufficient if the researcher has access to the deployment site and can remedy the network problem(s), if not, a larger SD card could be used - up to 64GB SD cards are currently supported by the system.

WiFi Provisioning Script:

One feature worth noting that makes deploying FireFly DAQ nodes easier than many commercial competitor products is the ease of provision them. In particular, connecting them to the local WiFi network for connectivity at the target deployment site. Of course the researcher could program the WiFi credentials into each node manually using a debug interface, but I found through field experience that this is both tedious and time consuming.

I have developed a simple, yet effective scheme where an un-provisioned node, on bootup scans for nearby wireless access points (AP) and the SSID they braodcast (the name of the Wifi network). Traditionally one would select the desires SSID , supply a password if needed and connect. Instead, I wrote a tool that looks for an SSID that is formatted in a particular format. In particular, it is looking for a SSID that has the form "ffly_X_Y", where X is the name of the target AP that the node should connect to and Y is the password it will need. For instance, if an un-provisioned node sees an SSID of "ffly_SG_uw507", then it will attempt to connect to network "SG" with password "uw507".

Of course, this scheme is not secure as the SSID and password are broadcasted in plain text. TO remedy this, I make use of symmetric cryptography where the network password is encrypted before broadcasting, which can then be decrypted on the FireFly nodes pre-programmed with the symmetric key. This results in a system where the researchers have to program their nodes once and from there on can deploy them in different environments and connect to new networks without the need to have physical access to the debug interface for each node. It should be noted that a provisioned node can be made to 'forget' the network by setting DIP switches and power cycling - similar to how a network router is reset to factory defaults.

5.2.5 Block 4: Local Aggregation Server

Each deployment site runs a local aggregation server to which each DAQ node uploads the data packets to and downloads bootup configuration data from. This server is also responsible for managing all the nodes, monitoring their status, report anomalies and in

general the central command and control system for the site deployment.

Each DAQ node is configure to find the local aggregation server through ZeroConf service discovery and have their relay thread upload data continuously as it becomes available. As evident, this means that the server needs to have a high uptime and be able to scale efficiently to a large number of nodes. I designed the server keeping in mind that it would often run of a laptop deployed by the end researcher and should support at least 100 DAQ nodes constantly streaming data to it. This local server is complex and performs a large number of tasks, however for brevity only a few of the key features are discussed in this dissertation.

The first role this server plays is in acting as a central point for each node to download their ADC configuration from. For each node, the server maintains a configuration file which can requested using a RESTful API call. In other words, on bootup the nodes could issue a HTTP call and retrieve their specific configuration. In addition to the low-level configuration, the server also maintains a mapping between the DAQ node's ID, its ADC and their respective channels to high-level application specific labels. For example, when a user installs Node-1 and connects a Kettle to ADC channel 0, the server constructs a XML file from user's input that maps Node-1.ADC_A.0 to Kettle_Current and Node-1.ADC_B.0 to Kettle_Voltage. The format of naming scheme is flexible and can be changed as needed.

The core functionality of the server lies in acting as an 'aggregation' server, i.e. take incoming data from all the nodes, process it and organize it in a suitable data structure for consumption by the disaggregation algorithms. There are several steps in attaining this functionality. When architecting the server, I explicitly chose to build on top of a non-blocking asynchronous HTTP server. A non-blocking server means that a request made to it would not block the entire server for the duration of the IO operation. Instead, it would return or relinquish control immediately whether the request was successful or not. Asynchronous property ensures that a 'background' process is started to handle the request and the API would result in a callback when the background process has finished. Combined, they ensure that large, time consuming and blocking IO operations happen in the background and the server made available to handle other tasks. In our case, since the primary request being handled by the server is data upload (or largely IO), this is particularly useful as other requests could be handled while the previous IO is in progress.

The asynchronous server of choice was the Tornado web framework written in Python. In my performance experiments with the server performing all the tasks including processing and archiving data as described later, I attained an average HTTP request rate of 649 requests/second with the server running with a single CPU thread. On a multi-core machine, multiple thread could be spawned further increasing the throughput. In terms of the number of nodes this server could support with each node sampling 6 channels @ 1ksps, a 4-core server (typical laptop processor) would scale to 500 nodes. This is sufficiently large for even the most dense DAQ node deployments.

The web server on receiving a packet from a node performs several tasks including those specific to energy disaggregation (see Figure 5.9 for an overview of the data flow and various tasks). First, as soon as the packet is received, a CRC check is performed. If this check fails, the server returns with an error code that forces the node to resend the data packet. If no errors are detected, the packet is saved to persistent storage (often a large external hard disk drive) immediately. This is necessary, since the steps performed afterwards cause irreversible or destructive changes to the data packet. Once backed up, the data packet is sent to a serialization routine. The purpose of this routine is to de-packetize the data and reconstruct a continuous stream of samples. Additionally, timestamp for each sample is computed while ensuring that no packets were dropped and that the clock jitter is within specification. Any deviations are logged and reported via a user facing graphical dashboard. If the deviation is too large, corrective measure is taken by re-initializing the stream and marking the gap as missing samples.

The stream of samples is then used alongside the ADC channel-to-appliance mapping to extract its voltage and current waveforms. These are then utilized to compute real-power, power factor, RMS current and voltage and AC line frequency at a predetermined update rate. In my test deployments this rate was set at 10Hz, which I deemed sufficient for my energy use groundtruth data. The current waveform is constantly monitored to detect any clipping as a result of ADC gain setting being too high, in which case, a new gain value is computed and pushed downstream to the XMega as an ADC configuration update. Additionally, the computed power metrics are further decimated and displayed on a dashboard allowing the researcher to manually monitor the integrity of the DAQ system in

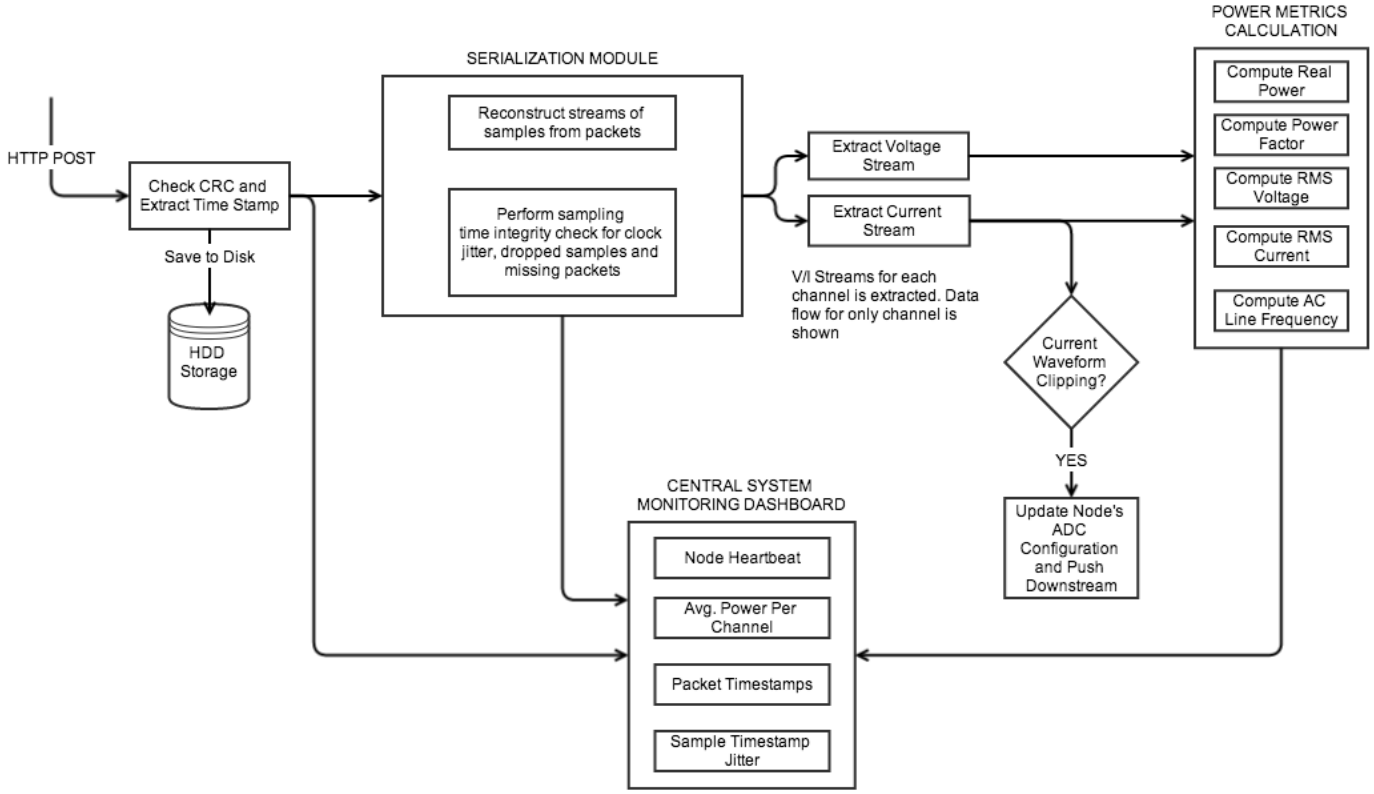

Figure 5.9: High level overview of the various tasks the local aggregation server performs. All DAQ nodes deployed at a site communicate with one such server that archives the data, computes application specific metrics and also acts as a central point for the researcher to command and control the various aspects of the node hardware. Not shown is the configuration functionality that allows DAQ nodes in field to update their configuration by requesting from the same server.

real-time and catch anomalies early. Lastly, the server outputs the power metrics (or other application specific metrics) in a format suitable for storage and analysis in the cloud. In this particular instance, the data is saved in a SQL format which can directly be imported into a data warehousing database called InfoBright.

5.2.6 Block 5: Remote Aggregation Server

The data output from the local aggregation server is ready for consumption for energy disaggregation research directly, however it can quickly become cumbersome to deal with large number of flat files. I implemented a database server, routable over the Internet that the local aggregation server across deployment sites can upload the data to. This allowed me to conveniently access and analyze all of the data from a central point. Another reason for this server was the need to query the data. For instance, with the data in a database I could query it for all time ranges where an electric kettle was turned on in a particular home. I could then use these time ranges to extract the various power metrics data from the database and use for algorithm development and research. Another feature the database affords is to allow generating appliance ON/OFF labels using thresholds specified via SQL. This allowed me to experiment with various thresholds and generate ground truth labels for use in machine learning algorithm.

The choice of database for use with the data from the FireFly DAQs is non-trivial, primarily due to the large number of data points (which translate to rows) generated per minute. In choosing an appropriate database technology, I considered SQL and NOSQL databases and studied them in-depth to finally settle on the InfoBright database. InfoBright is a data warehousing database, specifically designed to manage massive amounts of machine-generated data – precisely the case with data from FireFly DAQs installed in homes. In addition to easily scaling to even terabytes of data, one of the key features of InfoBright is its ability to run analytical queries in a matter of seconds as opposed to many hours with traditional databases. This particular capability is the reason I could run queries to select the right threshold for groundtruth label generation directly on the database.

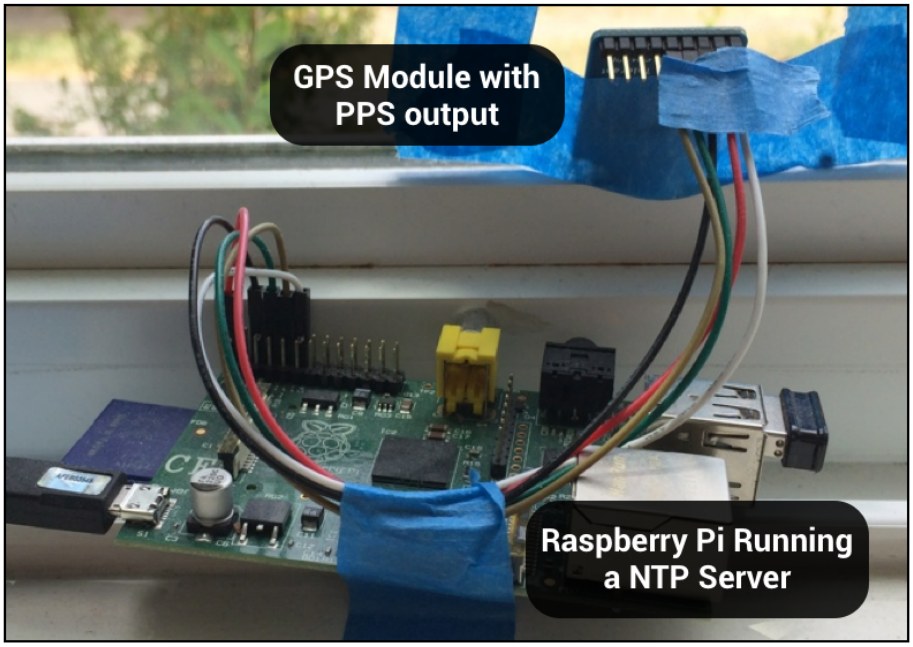

Figure 5.10: A custom GPS disciplined Stratum-1 NTP server built using a GPS module and a Raspberry Pi running Linux. On a local area network, timing coupling or offset of an average of 300 nanoseconds with a jitter less than 3 microseconds can be achieved.

5.2.7 Block 6: GPS Disciplined Clock and NTP Server

As I mentioned previously in this chapter, maintaining accurate global time on each of the FireFly DAQ nodes is quite important so that the data from these disparate nodes can later be correlated in time. The most direct way of doing this is to let the Linux OS on each of the nodes sync to a NTP server. This is exactly what is done on the nodes however, not just any NTP server could be used.

NTP is a networking protocol that allows computers on a network to synchronize their time. NTP servers can be categorized into "stratums" depending on how close they are to a high-precision reference clock. The highest level, a stratum-0 NTP server is directly connected to a high precision time keeping device such as an atomic clock, GPS clock or other radio clock. In other words, startum-1 server would sync to a stratum-0 server and would be the next best source of time.

Typically, when synchronizing to a NTP server on the internet, as most internet connected devices do, the time can be synced to a few milliseconds of each other. Unfortunately, for our data collection needs, a few milliseconds is not acceptable – it translates to an offset of a few samples at 1ksps sampling rate. The next best approach is to run a local stratum-2 server, i.e., it syncs to a NTP server on the internet and the FireFly nodes would sync to the NTP on the LAN. This works quite well and time synchronization across nodes is on the order of 35 microseconds.

An offset of about 35 microseconds is quite acceptable for most applications, however to further tightly couple the timing, I designed and implemented a custom stratum-1 NTP server (see Figure 5.10) based around a small Linux computer (Raspberry Pi, Model A) and a GPS module. This use of a GPS is a common design and often referred to as a GPS disciplined clock. The GPS module provides an accurate 'pulse per second' as derived from satellite signals that serve as stratum-0 clock. This pulse is used to keep the internal Linux clock in check – more accurately adjust for the drift every second using the pulse. With the FireFly DAQs syncing to this custom NTP server over wireless LAN, I was able to achieve an average jitter of ≤ 3 microseconds and an average offset of *300 nanoseconds*. Such tight coupling ensures that even at high sampling rates, data from different nodes can be reliably synced and correlated for analysis.

Chapter 6

CONCLUSION AND FUTURE DIRECTIONS

In this dissertation I presented a new sensing technique called ElectriSense that requires minimal instrumentation for sensing electrical appliance usage in the home. The core contribution of my work is in leveraging a signal that is otherwise considered "noise" by mainstream approaches. This noise, Electromagnetic Interference is produced by most consumer electronics and in this dissertation I demonstrated how it can be re-purposed as a signature or fingerprint to detect and classify the use of electrical appliances in real-time.

In addition I demonstrated that the new EMI sensing approach provides a signal that is rich in information to enable new applications that go beyond just estimating whether an appliance is on or off. Specifically, I demonstrate that EMI can be used to infer the internal operating states or modes of an appliance, which can be used to enable new applications, better energy disaggregation models and build better human activity models. I discussed two such applications in this dissertation - application to novel user interaction that allows inferring how a human is interacting with a CFL or non-touch LCD; and inferring TV viewing habits of a user. Since both these applications improve human activity modeling and how they interact with electrical appliances, it has profound implications for energy disaggregation.

6.1 Implications of EMI Sensing for Different Communities

The applications for this sensing technique are myriad. Here I discuss how I envision this new sensing technique impacting various fields of research. I encourage researchers in these fields to consider this new, easy to install and low cost sensing technique as an alternative to conventional distributed sensing approach.

Energy Disaggregation and Sustainability Community:

In my view EMI sensing has the most concrete and immediate implication for the energy disaggregation community. This community is largely made up of policy makers, regulatory authorities, utilities, renewable resources government agencies and behavioral change researchers focusing on eco-feedback. As discussed in Chapter 1 and 2, these communities often need to monitor fine-grained energy usage, often at the individual appliance level so that the energy conservation measures they take and propose can have the most impact.

Additionally, in my communications with government agencies interested in compiling energy models across the nation, I found out that they can easily instrument new buildings with a plethora of sensors while they are under construction to study and evaluate them. However, they run into the 'retrofit' problem, i.e. 99% of the buildings that already exist do not have such sensors and alternate sensing methods are required. NIALM techniques promise a solution that would allow sensing of fine-grained energy use without the expensive overhead of installing large number of sensors, but unfortunately these techniques have not yet matured for widespread adoption, largely due to lack of new features and the resulting stagnation in performance improvements.

To this end, I believe that EMI sensing significantly advances the state-of-the-art in features available to energy disaggregation community for designing improved algorithms. EMI based features from ElectriSense not only add an additional set of features but presents a much richer and higher resolution feature space that has not been previously explored. For instance, multiple appliances could be operating simultaneously, yet their EMI could be different - a scenario that poses a hard problem for prior approaches. Additionally the intensity of EMI observed is largely independent of the power an appliance consumes. This observation overcomes an important practical limitation of current and power based approaches that small loads such as a 10W phone charger's use cannot be detected; its EMI however can be observed with a high SNR. Of note is also my findings that EMI based sensing could be used to infer an appliance's internal operating state. This is a very valuable and novel finding that will allow energy disaggregation researchers to tap into a new feature space and go further than simply detecting if an appliance if being used or not. I strongly believe that leveraging this operating state information can profoundly impact many facets of disaggregation research positively.

Human Activity Sensing and Ubicomp Community:

Beyond application to energy monitoring, EMI sensing has a direct application in sensing human activity in the home and supporting home sensing goals for the ubiquitous computing community. For instance, location and activity of people in the space could be inferred by observing the electrical activity and actuation of certain appliances in the home. This information could then be used to enable applications that rely on such contextual information. For example, sensing kitchen lights being tuned on, followed by a fridge door opening and closing and stove activity would strongly suggest that a meal is being prepared. In contrast, sensing freezer door being opened followed by microwave usage may indicate reheating or cooking frozen food. Such fine-grained activity sensing without the need to install intrusive sensors throughout the home could very well support research goals for non-technical communities as well. For example, researchers in the medical domain studying in-situ rehabilitation would be interested in tracking how much a person moves around her home and what activities they perform. An activity sensing system based on ElectriSense could potentially be a lucrative solution.

Sensing appliance level electrical activity also has implications for home automation and ubicomp applications. For example, detecting a room's lights being tuned on could indicate that someone entered that room, which could then automatically adjust the room heating or cooling for a comfortable environment. Lastly, since ElectriSense can distinguish between on and off events, these events could be used as triggers to control other appliances not originally intended for control by a particular light switch or appliance. One can imagine an intelligent home automation system that maps the on event of a stove turning on to smoke hood lights and exhaust system without having to install new wiring between the two.

Security and Privacy Community:

Discovering new signals that can be used to infer appliance usage, an appliance's operating state and human activity in the home has direct implications for breach of privacy and 'information leakage' over the powerlines. At the simplest level, just monitoring a home's

EMI could provide a strong indication of the resident's coarse activity patterns, which could later be used against them, such as infiltrating the home when everyone's asleep. A more sophisticated attack based on inferring operating states from EMI could potentially extract information about what computer applications a user is using on her computer. It may also be possible to mount a side-channel timing attack on cryptographic processes to extract a user's private keys. I encourage security researchers to consider the information leakage over the powerline in the form of EMI as a viable attack vector.

6.2 Future Directions

Though this dissertation summarizes a large amount of research I conducted into EMI based sensing over the past few years, I believe I have only scratched the surface and demonstrated few of the many things that are possible with these newly discovered signals. Like with any research, the answering of questions has led to new questions. There are several future directions that this research work can take, several of which can have a profound impact on multiple application domains and research communities.

Analyzing Time Varying EMI:

Though time varying EMI and its potential application to operating state estimation is detailed in Chapter 4, I believe further in-depth experimentation is required to fully understand its scope. In particular, it would be an interesting research direction to evaluate whether the variations in EMI could be used to predict an impending appliance failure. This could be quite useful in a commercial setting, such as a factory where machinery downtime is expensive.

Reducing Training Effort:

This is a research direction that can be pursued by machine learning researchers as well as those interested in circuit and systems. In the current setup, ElectriSense requires supervised training in order to predict which appliance is being turned on or off. As more sophisticated machine learning models are built for better performance, the supervised training approach is unfortunately unsustainable due to its tedious nature. Homeowners

could not be expected to provide a training label for each and every appliance in the home. To alleviate this, several approaches can be explored.

One approach is to *synthetically generate the EMI* that an appliance is expected to produce given its circuit model. One of the key strengths of my EMI based work is that the EMI is a byproduct of an engineered circuit. Thus, it should be possible that given the circuit model, one could predict what EMI will be produced. Though at a high level, Figure 3.1 demonstrates this concept briefly, however further experimentation is required. Another similar approach to learning EMI signatures without explicit training sample is to *analyze Federal Communications Commission's (FCC) emissions reports* made by appliance manufacturers. In the US, the FCC mandates that manufacturers file reports indicating how much and at what frequency certain classes of appliances produce EMI before they could be commercially sold to consumers. These reports are publicly accessible and it may allow harvesting details of expected EMI from an appliance.

A complementary approach to reducing training effort is to bring the human in the loop and perform *active learning*. In this case, the system starts with performing unsupervised clustering and when a particularly interesting (either statistically interesting or explicitly specified the user) electrical event happens the user is prompted with a question asking if they operated an appliance on their smart phone. Their responses could then be used to train the system over time and assign labels to clusters created using unsupervised machine learning techniques.

Yet another approach is to build models using *crowdsourcing*. Since the EMI signatures remain largely consistent across similar models of an appliance, homeowners could be provided an incentive to label as many appliances as possible, which are then shared globally. When a different user installs the ElectriSense system, they could seed their machine learning models with the global database reducing the amount of training required. Overtime a robust database of a large number of appliances would exist reducing training effort for everyone using the system.

Beyond EMI on Powerlines:

Though I primarily focused on the EMI I observed on the powerlines, I found that

my technique of EMI sensing can potentially be extended to encapsulate other sources and mediums of EMI and spur innovation in other application domains. In a home, this approach could be extended to capturing *radiated EMI* that may allow us to analyze EMI from appliances that do not conduct it well over the powerline. A particular appliance for which this is most likely true is a microwave. I have observed radiated EMI in the 2.4GHz range when a microwave is operational, however none on the powerline. A broadband RF receiver could potentially be built into the plug-in sensor to capture both conducted and radiated EMI.

The EMI sensing approach does not have to be limited to applications inside the home. For example, I have observed *radiated EMI in a car* that strongly correlates to the engine's RPM. This is an important discovery as it could allow non-contact sensing of a vehicle's various parameters with implications for ubicomp research, eco-feedback and privacy. There is also potential for performing EMI sensing on a mobile phone, which could for example allow sensing when a person is in her office because of the characteristic radiated EMI from the PC.

I hope this work acts as an inspiration to anyone interested in leveraging EMI and in general 'noise' around us to further push forward the boundary of sensing technology. I strongly believe that identifying and capitalizing on new sensing opportunities are pertinent to enabling future computing platforms and driving innovation. Fortunately, such sensing opportunities are all around us, waiting to be discovered; waiting for someone to *make sense of the noise*.

BIBLIOGRAPHY

[1] 8devices. Carambola Embedded Module.

[2] DO Abdeslam, Jean Mercklé, and Raphaël Ngwanyi. Artificial neural networks for harmonic estimation in low-voltage power systems. *4th International ICSC ...*, 2004.

[3] DO Abdeslam and Patrice Wira. A unified artificial neural network architecture for active power filters. *Industrial*, 54(1):61–76, 2007.

[4] G G Acosta, C J Verucchi, and E R Gelso. A current monitoring system for diagnosing electrical failures in induction motors. *Mechanical Systems and Signal Processing*, 20(4):953–965, 2006.

[5] Peter R Armstrong, Chris R Laughman, Steven B Leeb, and Leslie K Norford. Fault detection based on motor start transients and shaft harmonics measured at the RTU electrical service. 2004.

[6] Inc. Belkin International. Belkin Energy Dataset, 2013.

[7] Bidgely. Bidgely.

[8] S Bin and H Dowlatabadi. Consumer lifestyle approach to US energy use and the related CO2 emissions. *Energy Policy*, 33(2):197–208, 2005.

[9] A J Marques Cardoso, S M A Cruz, and D S B Fonseca. Inter-turn stator winding fault diagnosis in three-phase induction motors, by Park's vector approach. *Energy Conversion, IEEE Transactions on*, 14(3):595–598, 1999.

[10] K. Carrie Armel, Abhay Gupta, Gireesh Shrimali, and Adrian Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52:213–234, October 2012.

[11] Gabe Cohn, Sidhant Gupta, Jon Froehlich, Eric Larson, and Shwetak N Patel. GasSense: Appliance-level, Single-point Sensing of Gas Activity in the Home. In *Proceedings of the 8th International Conference on Pervasive Computing*, Pervasive'10, pages 265–282, Berlin, Heidelberg, 2010. Springer-Verlag.

[12] A I Cole and A Albicki. Algorithm for nonintrusive identification of residential appliances. In *Circuits and Systems, 1998. ISCAS '98. Proceedings of the 1998 IEEE International Symposium on*, volume 3, pages 338 –341 vol.3, 1998.

[13] A I Cole and A Albicki. Data extraction for effective non-intrusive identification of residential power loads. In *Instrumentation and Measurement Technology Conference, 1998. IMTC/98. Conference Proceedings. IEEE*, volume 2, pages 812 –815 vol.2, May 1998.

[14] Small Commercial, Residential Buildings, K D Lee, L K Norford, and S B Leeb. Equipment Scheduling and Cycling Automated Diagnostics Based on Start-Up Signals Development of a Functioning Centrally Located Electrical- Load Monitor. *Signals*, 3, 2003.

[15] European Commission. Eurostat (Code: tsdpc320), 2010.

[16] Jetway Computers. Jetway Computers Compliance Report.

[17] PK Dash and DP Swain. Harmonic estimation in a power system using adaptive perceptrons. 1996.

[18] Willem Den Boer. *Active matrix liquid crystal displays: fundamentals and applications*. Elsevier, 2011.

[19] K Ehrhardt-Martinez, KA Donnelly, and S Laitner. Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities. 2010.

[20] Mohamed El Hachemi Benbouzid. A review of induction motors signature analysis as a medium for faults detection. *Industrial Electronics, IEEE Transactions on*, 47(5):984–993, 2000.

[21] Jon E Froehlich, Eric Larson, Tim Campbell, Conor Haggerty, James Fogarty, and Shwetak 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.

[22] Suman Giri, Mario Bergés, and Anthony Rowe. Towards automated appliance recognition using an {EMF} sensor in {NILM} platforms. *Advanced Engineering Informatics*, 27(4):477–485, 2013.

[23] Sidhant Gupta, MS Reynolds, and SN Patel. ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the home. *Ubicomp 2010*, 2010.

[24] GW Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 1992.

[25] P3 International. Kill A Watt.

[26] Lei Jiang, Jiaming Li, Suhuai Luo, Sam West, and Glenn Platt. Power Load Event Detection and Classification Based on Edge Symbol Analysis and Support Vector Machine. *Applied Computational Intelligence and Soft Computing*, 2012:1–10, 2012.

[27] Pengju Kang, David Birtwhistle, J Daley, and D McCulloch. Noninvasive on-line condition monitoring of on load tap changers. In *Power Engineering Society Winter Meeting, 2000. IEEE*, volume 3, pages 2223–2228. IEEE, 2000.

[28] Younghun Kim, Thomas Schmid, Zainul M Charbiwala, and Mani B Srivastava. ViridiScope : Design and Implementation of a Fine Grained Power Monitoring System for Homes. *Ubicomp 2009*, pages 245–254, 2009.

[29] G B Kliman, R A Koegl, J Stein, R D Endicott, and M W Madden. Noninvasive detection of broken rotor bars in operating induction motors. *Energy Conversion, IEEE Transactions on*, 3(4):873–879, 1988.

[30] J Zico Kolter and Tommi Jaakkola. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. XX, 2012.

[31] Christopher Laughman, Kwangduk Lee, and Robert Cox. Power signature analysis. *IEEE Power and Energy*, (april 2003), 2003.

[32] Steven B Leeb, Steven R Shaw, and James L Kirtley. Transient Event Detection in Spectral Envelope Estimates. *IEEE Transactions on Power Delivery*, 10(3):1200–1210, 1995.

[33] Jian Liang, Simon K K Ng, Gail Kendall, and John W M Cheng. Load Signature Study Part I : Basic Concept , Structure , and Methodology. *Power*, 25(2):551–560, 2010.

[34] Jian Liang, Simon K K Ng, Gail Kendall, and John W M Cheng. Load Signature Study Part II : Disaggregation Framework , Simulation , and Applications. *Power*, 25(2):561–569, 2010.

[35] B. Mettler-Meibom, B., and Wichmann. The influence of information and attitudes toward energy conservation on behavior. (Translated by M. Stommel, Michigan State University.). *H. Schaefer, ed., Einfluss des Verbraucherverhaltens au den Energiebedarf Privater Haushalte. Berlin: Springer-Verlag*, 1982.

[36] B Metz and O R Davidson. *Climate Change 2007: Mitigation: Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Intergovernmental Panel on Climate Change, 2007.

[37] Leslie K Norford and Steven B Leeb. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, 24(1):51–64, 1996.

[38] OPower. OPower, Inc.

[39] Shwetak N Patel, Sidhant Gupta, and Matthew S Reynolds. The Design and Evaluation of an End-user-deployable, Whole House, Contactless Power Consumption Sensor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 2471–2480, New York, NY, USA, 2010. ACM.

[40] Shwetak N Patel, Matthew S Reynolds, and Gregory D Abowd. Detecting Human Movement by Differential Air Pressure Sensing in HVAC System Ductwork : An Exploration in Infrastructure Mediated Sensing. pages 1–18.

[41] Shwetak N Patel, Thomas Robertson, Julie A Kientz, Matthew S Reynolds, and Gregory D Abowd. At the Flick of a Switch : Detecting and Classifying Unique Electrical Events on the Residential Power Line. *Ubicomp 2007*, pages 271–288, 2007.

[42] Shwetak N Patel, Khai N Truong, and Gregory D Abowd. PowerLine Positioning: A Practical Sub-room-level Indoor Location System for Domestic Use. In *Proceedings of the 8th International Conference on Ubiquitous Computing*, UbiComp'06, pages 441–458, Berlin, Heidelberg, 2006. Springer-Verlag.

[43] Dennis E Phillips. Supero : A Sensor System for Unsupervised Residential Power Usage Monitoring.

[44] PlugWise. Plugwise Home Efficiency.

[45] J T Powers, B Margossian, and B A Smith. Using a rule-based algorithm to disaggregate end-use load profiles from premise-level data. *Computer Applications in Power, IEEE*, 4(2):42–47, April 1991.

[46] A Prudenzi. A Neuron Nets Based Procedure for Identifying Domestic Appliances Pattern-of-Use from Energy Recordings at Meter Panel . *Electrical Engineering*, 00(c):1–6, 2002.

[47] Lazar Rozenblat. Functioning of a switched mode Power supply.

[48] S R Shaw, L K Norford, D Luo, and S B Leeb. Detection and Diagnosis of HVAC Faults via Electrical Load Monitoring. *HVAC&R Research*, 8(1):13–40, 2002.

[49] Parham Shenavar and Ebrahim Farjah. Novel embedded real-time NILM for electric loads disaggregating and diagnostic. *Computing*, pages 1555–1560, 2007.

[50] D Srinivasan, Senior Member, W S Ng, and A C Liew. Neural-Network-Based Signature Recognition for Harmonic Source Identification. *Power*, 21(1):398–405, 2006.

[51] Energy Star. Energy Star SMPS and linear power supplies compliance, 2010.

[52] Tendril. Tendril Volt II.

[53] William T Thomson and Mark Fenger. Current signature analysis to detect induction motor faults. *Industry Applications Magazine, IEEE*, 7(4):26–34, 2001.

[54] Thomas G Zimmerman, Joshua R Smith, Joseph A Paradiso, David Allport, and Neil Gershenfeld. Applying Electric Field Sensing to Human-Computer Interfaces. *Human Factors*, (May), 1995.