MIT Open Access Articles

EMI Spy: Harnessing electromagnetic interference for low-cost, rapid prototyping of proxemic interaction

EMI Spy: Harnessing Electromagnetic Interference for Low-Cost, Rapid Prototyping of Proxemic Interaction

Nan Zhao, Gershon Dublon, Nicholas Gillian, Artem Dementyev, and Joseph A. Paradiso
Responsive Environments Group, Massachusetts Institute of Technology Media Lab
75 Amherst St., Cambridge, MA 02139
{nanzhao,gershon,ngillian,artemd,joep}@media.mit.edu

Abstract— We present a wearable system that uses ambient electromagnetic interference (EMI) as a signature to identify electronic devices and support proxemic interaction. We designed a low cost tool, called EMI Spy, and a software environment for rapid deployment and evaluation of ambient EMI-based interactive infrastructure. EMI Spy captures electromagnetic interference and delivers the signal to a user's mobile device or PC through either the device's wired audio input or wirelessly using Bluetooth. The wireless version can be worn on the wrist, communicating with the user's mobile device in their pocket. Users are able to train the system in less than 1 second to uniquely identify displays in a 2-m radius around them, as well as to detect pointing at a distance and touching gestures on the displays in real-time. The combination of a low cost EMI logger and an open source machine learning tool kit allows developers to quickly prototype proxemic, touch-to-connect, and gestural interaction. We demonstrate the feasibility of mobile, EMI-based device and gesture recognition with preliminary user studies in 3 scenarios, achieving 96% classification accuracy at close range for 6 digital signage displays distributed throughout a building, and 90% accuracy in classifying pointing gestures at neighboring desktop LCD displays. We were able to distinguish 1- and 2-finger touching with perfect accuracy and show indications of a way to determine power consumption of a device via touch. Our system is particularly well-suited to temporary use in a public space, where the sensors could be distributed to support a pop-up interactive environment anywhere with electronic devices. By designing for low cost, mobile, flexible, and infrastructure-free deployment, we aim to enable a host of new proxemic interfaces to existing appliances and displays.

I. INTRODUCTION

More than a decade has passed since Bill Buxton challenged the nascent 'ubiquitous computing' research community to incorporate proximal sensing and context into our computing interfaces [2], but we have not yet achieved many of the seamless user experiences he envisioned [10]. Buxton famously compared the interfaces of a simple motion-sensing light and a computer screen saver, pointing out the ways in which the light might appear to a user to be more intelligent, recognizing them at a distance and adjusting the environment accordingly. Buxton's example succinctly highlighted the shortcomings of our smart devices, especially with respect to sensing. Today, despite near-ubiquitous network connectivity, many devices remain, in the words of Greenberg et al., "blind to the presence of other devices" [10] and often, their users. To paraphrase Buxton, how hard would it be to create a reactive environment?

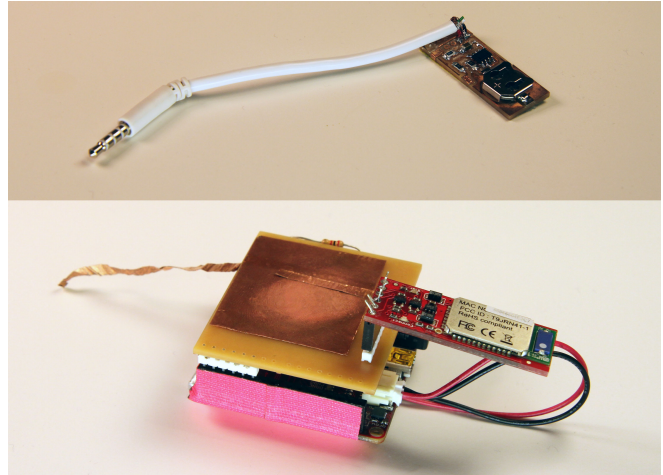
Fig. 1. Top: the $10 audio input configuration of EMI Spy. Bottom: the Bluetooth wearable EMI sensor.

The emergence of the Internet of Things (IoT) has brought this question to the fore. Researchers and developers are increasingly interested in identifying connected devices proximate to the user. This approach to local discovery allows the number of distributed devices to scale and saves power by allowing devices to remain in a sleep mode until a user comes within range. There are many research and commercial approaches to this problem—Bluetooth beacons or RFID tags, for example [12]. But these approaches typically require instrumenting the devices in the environment and equipping users' mobile devices with readers. They also do not usually have the sensing resolution to detect user touch and gestures, or to measure fine-grained range, limiting applications to detecting presence. Body sensor networks (BSN) offer the possibility of measuring and classifying electronic signatures of devices in the vicinity of a user over a range comparable to RFID, without requiring instrumentation in the environment. These signals can also be used to enable gesture and touch recognition on existing devices.

In this paper we present the EMI Spy, a low-cost mobile peripheral for rapid prototyping of ubiquitous proxemic interaction with displays and other electronic devices with high-frequency switching, such as power supplies. Leveraging the wearer's body in the signal path, EMI Spy captures electromagnetic interference (EMI) and delivers the signal to

a user's mobile device or PC. Spectral features are extracted from the signal and fed through a fast-learning pattern recognition pipeline [9]. By touching different displays in a training process, users can quickly train the system to later recognize device-specific electromagnetic fingerprints in a 2-m radius around them. Preliminary results show that other devices containing high-frequency switching components also have similar spectral signatures that could be fed into a machine learning pipeline to recognize the devices, as well as their internal states in some cases. For example, we were able to identify a varying load on a power supply by detecting changes in the EMI spectrum.

Besides the mobile peripheral and a computer, our system requires no infrastructure in the environment. We developed two versions of the hardware. In the first version, to keep hardware costs below $10, EMI Spy connects to the existing audio input on its host device, a laptop. In this version, the human body acts as an antenna, and we measure a signal between the body and the local ground of the laptop, which is directly coupled to the room ground when plugged in and ambiently coupled to the room ground when running on battery (as in our experiments). We also developed a Bluetooth wearable version; in this version, because we had no direct AC ground reference, we changed the topology to use the human body as an ambiently-coupled path to the building ground [15].

We considered several concept applications that motivate the capabilities of EMI Spy. Because of its low cost and zero infrastructure requirements, our system is particularly well suited for temporary use in a public space, where the sensors could be distributed to support a pop-up interactive environment anywhere with displays. We also considered a pointing and touch-to-connect scenario, for example in which a user could approach their fridge to start a recipe application, or move media from display to display by pointing. EMI Spy could also support device-centric, RFID-like location tracking. In this scenario, it would add a rich, device-centric dimension to an activity recognition system.

We demonstrate the feasibility of these concepts in several preliminary user studies. In a test of the audio input configuration with 5 users and 6 identical model LCD displays distributed throughout a building, we achieved 96% recognition rates at close range, approaching the performance of a fixed infrastructure RFID system currently in place. For the same test using the wearable configuration, the detection radius dropped by a factor of 10, to 20-cm. The difference is a result of the different topology of the wearable design; we have identified strategies for improving the sensitivity of the wearable that are detailed in the discussion below. In a test of the audio input configuration with 3 typical (non-touch) LCD displays on a desk, we achieved 90% accuracy classifying pointing gestures, as well as 1-finger and 2-finger touching. The signal is not dependent on the content displayed. By designing EMI Spy for low cost, mobility, and infrastructure-free deployment, we aim to enable rapid prototyping of proxemic interfaces to displays, and in the future, generic appliances.

II. BACKGROUND

Most appliances emit electromagnetic noise, and because devices have different internal switching mechanisms, each produces a unique EMI signature. For identical devices, there is sufficient variability in production to make individual recognition possible. Gupta, et al. investigated transient and continuous electromagnetic phenomenon in the power line, and were not only able to reliably detect the powered devices, but also distinguish between different states [11]. We point interested readers to that work for a more in-depth examination of the sources of EMI. With a similar approach, Chen et al. were able to recognize hover and touching gestures in front of a LCD monitor [3]. The difference between their system and ours is that in order to observe signals in the power line the user needs to tap into the infrastructure. EMI Spy, on the other hand, provides a mobile and wireless solution.

Artists and designers have explored EMI as a medium, mapping the invisible electromagnetic terrain to user experience. In [16], Vaucelle et al. introduce a bracelet that captures and visualizes EMI to the wearer for immediate feedback as well as review later on.

Researchers have used the body as an antenna to capture EMI signals. Cohn et al. have developed a number of applications based on capturing EMI through the body, focusing on classification of indoor location, movements and gesture recognition [6], [5], [4]. In [5], the authors mention that they are able to recognize kitchen appliances with 100% accuracy. However, in these experiments, the authors make an assumption that the electromagnetic noise profile is static. In general, this is not the case; the EM background in a typical home or office changes when other devices go into standby mode, turn off, or are moved. In contrast, our system uses features that are not affected by changes in the background. By classifying the EMI signature of each target device independently, our approach is insensitive to day-to-day changes in the environment.

Others have demonstrated active transmission of EM signals through the human body. [7] and [18] implemented interaction between user and devices using the body as both a physical and metaphorical carrier of information from one device to another.

III. PROXIMAL SENSING AND PROXEMIC INTERACTION

Borrowing from anthropology, a theory of proxemic UI has been formalized in the absence of a ubiquitous supporting sensor technology. Greenberg et al. define five dimensions of proxemics for ubicomp: distance, movement, identity, location, and orientation [10]. EMI Spy can supply the first three and determine location wherever there are displays. While our system does not provide orientation directly, it can be used to detect pointing and other gestures, from which user orientation cues might be inferred.

Researchers have taken on Buxton's proximal sensing challenge in myriad ways. Kortuem et al. developed an ultrasound sensing USB dongle for peer-to-peer localization that allows devices to build models of their spatial

relationships [13]. Using a combination of cloud-based data sharing and touch sensing, Mistry et al. created SPARSH, software for transferring media from device to device by touching one and then the other [14]. Others have developed proxemic interfaces enabled by specialized sensors or radios, like RFID [1]. Increasingly, the near field communication (NFC) standard is enabling very short range proximal sensing and communication on equipped mobile devices. In contrast, our system can perform peer-to-peer ranging at low cost on "dumb" appliances without built-in touch, radios, or special sensor technologies. Like SPARSH, EMI-aware applications could then use ubiquitous networking to push media from device to device in the background, with very little specialized hardware.

IV. System Design

The EMI Spy consists of hardware for sensing and signal conditioning, and software for feature extraction and machine learning.

A. Capturing EMI

The EMI is a differential signal between the local ground and a local signal plane, measured through the body. The switching frequencies of interest in typical home appliances range from approximately 100Hz to as high as 500kHz [11], or even higher for other types of switching circuits. Ideally, this signal would be digitized at 1Msps or higher; in practice, a much lower sampling rate is sufficient to identify many devices. For signal components higher than our sampling frequency, we intentionally took advantage of the aliasing effect to shift those out-of-band frequencies to the lower spectrum, and found the technique effective for our machine learning system, despite the loss in fidelity [17].

When a user touches a device, it dominates the energy of the spectrum, but we can still observe signals from other nearby devices. The spectral energy of the device decreases nonlinearly with distance from the source; the drop-off depends on the geometry and composition of the emitting surface. For example, the signal amplitude increases with the size of the emitting surface. Similarly, if the source is touching a metal surface, the signal will be stronger.

B. Hardware and Signal Processing

We developed two hardware configurations to deliver the EMI signal to a host device or server for processing. Figure 1 shows the two hardware configurations. The simplest makes use of the host's existing audio input; in this arrangement, the sensing hardware biases, buffers, and high-pass filters the signal with a 70Hz cutoff. In this configuration, the user contacts an electrode connected to the buffer input. Some built-in audio cards are capable of digitizing the signal at a 96kHz or 192kHz sampling rate, which is sufficient to capture much of the relevant spectrum. Most modern computers expose these high sampling rate options, enabling extremely low-cost sensing. Other devices, including many smartphones, top out at 48kHz sampling rates with aggressive anti-aliasing filters.
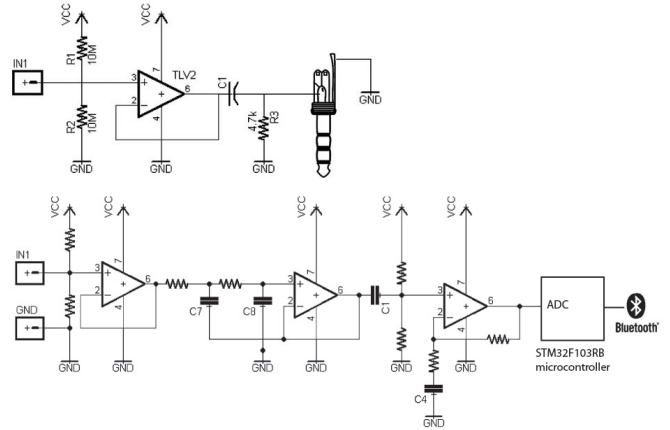

Fig. 2. Top: the schematic for the low-cost version of EMI Spy that requires only an audio input on the mobile device. Bottom: the schematic for the wearable version, which uses Bluetooth to communicate with a mobile device.

We also developed a wearable version of the hardware that would use an external microcontroller to digitize the signal and Bluetooth to interface to the host device, allowing us to tailor the conversion to our needs. In this configuration, the EMI signal is captured though a local signal plane, piped through a high impedance input, buffered, and then filtered with a second order Sallen-Key low-pass filter with a 200kHz cutoff. The signal is then fed to a built-in 12-bit ADC on the STM32F103RB microcontroller at a 100ksps sampling rate. The signal is buffered on the microcontroller and then sent to the host device over Bluetooth. In this configuration, the user contacts the ground plane of the wearable device.

Once the signal is digitized from either hardware module and fed to the server, we apply the Discrete Fourier Transform (DFT). Each device has unique switching characteristics, which manifest as different peaks in the frequency spectrum. Some devices have multiple significant peaks. For example, a typical laptop shows a peak for its power supply, display driver, and battery switching. We experimented with the FFT window size, which determines the resolution of the spectral analysis. With higher resolution, it is easier to reliably separate two devices with similar EMI signatures (e.g. identical models), but a larger window slows the system update rate. For the signal from the audio input, a 2048-sample window size with 1024 bins was empirically sufficient. For the noisier signal from the microcontroller, we used a 4096-sample window size with 2048 bins.

C. Machine Learning

To train EMI Spy to recognize an appliance, a user simply touches the appliance and records the signal for several seconds (corresponding to 100 2048-sample or 4096-sample windows). For each sample window, we compute the average energy of the spectrum. This value is used to reject bad samples caused by lost connections or other discontinuities, which have abnormally high spectral energies. We then find the maximum spectral peak in each of the remaining samples, and calculate the median frequency of the maximum peaks
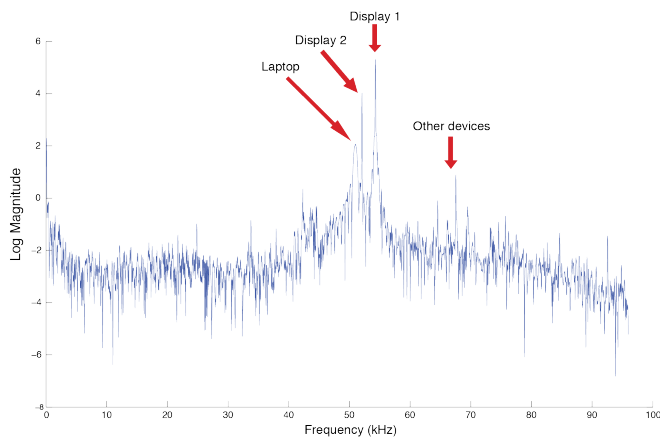
Fig. 3. 4096-bin spectrum of EMI signal captured at 192kHz, showing two LCD displays and a laptop.

Fig. 4. (a) This image shows the software GUI for the case where the user is touching the EMI Spy but not touching a display. On the left is the enlarged EM spectrum from both displays. The blue peak corresponds to the display on the left and the yellow peak to the display on the right. (b) In this image, the user touches the EMI Spy and the display on the left. The blue peak corresponding to the touched display rises significantly whereas the yellow peak does not change in amplitude. Spectral signatures of other nearby devices are visible in the GUI on the laptop screen.

across the training set. As a testament to the stability of the feature, we have observed very little variation in the median peak feature over days of measurement with different devices. However, a baseline level of noise persists when the user is not directly contacting the device. To account for this effect, we add random noise of +/- 1 DFT bin to the median frequency, retaining 10 artificially noisy samples as the training set for the device class. When using the audio input configuration, it is necessary to calibrate out the signal contributed by the host device, which is nearest the sensor and can dominate the training set. The calibration process also enables EMI Spy to identify its host later on.
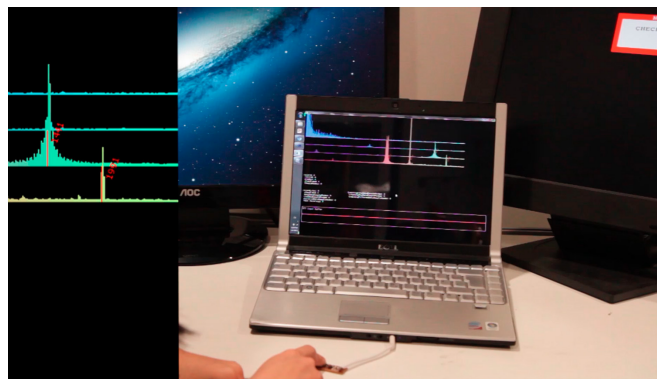
A user touches the sensing device and approaches an appliance. The incoming signal is digitized and the DFT is computed. The resultant DFT spectrum is then filtered to dampen high-frequency temporal changes. Next, we extract the spectral peaks exceeding an empirically determined threshold. The prediction is performed on each peak using a k-NN classifier with k = 10. The software displays the labels for all the detected devices and a measure of the distance to each one as a relative amplitude to the initial training data, assumed to be a full hand touch (the maximum of the amplitude range).

V. EVALUATION

We evaluated the performance of EMI Spy in 2 application scenarios: detecting as users approach and touch different public display kiosks, and detecting pointing and touch gestures at an array of 3 LCD displays on a desk.

A. Public display kiosk

In the public display kiosk scenario, we invited 5 users to move between 6 large LCD displays of the same model across several floors in a building; the displays are normally used as public digital signage utility. We chose identical model displays to maximally challenge the classification, as we found different model displays to be more spectrally distinct and therefore easier to identify. It is not uncommon to see identical model appliances distributed throughout a

building. The study was conducted during different hours of a workday to ensure a realistic level of background noise. We performed this study twice, once for each version of the hardware. In the audio input configuration, an unplugged laptop was placed on a rolling table and moved along with the study subject in fixed increments towards the display. The user touched the signal input of the EMI spy with one hand. In the wearable case, the user simply held the device in their hand, and a nearby computer processed the signals transmitted over Bluetooth.

100 samples of data per class were collected by one person and used to train the system; tests were conducted with 2500 samples collected from the others. We found that a single model from one person generalized across all the study participants and over days. Using the audio input configuration of EMI Spy and a laptop, we were able to uniquely identify the displays from 2-m away with an accuracy of 72% and a precision of 0.92. At 1-m and closer, the accuracy increased to 83% with a precision of 1. With the mobile device (user holding the locally-grounded battery and probing with the devices exposed input electrode as in Figure 5), the detection range dropped significantly; at 2-

Fig. 5. An example usage scenario with a display kiosk. The user holds the EMI Spy at its grounded battery, hence captures the electromagnetic interference from proximate electronic devices via its sensing electrode. These noise signatures are then used to uniquely identify the devices and detect gestures such as pointing at a distance and touch in real-time.

m, no devices were detected. At 1-m, we achieved 24% accuracy and a precision of 0.48; at 30-cm and closer, the classification performance increased dramatically, with an accuracy of 96.5% and a precision of 0.99. A possible solution to improve the performance of the wearable version will be discussed in the next section. Nevertheless, under many of the conditions we tested, we were able to approach the performance of the fixed infrastructure RFID system currently in place on the kiosks.

B. Pointing and touching gestures

Next, we explored the suitability of our system for detecting pointing at a distance and touching gestures at an array of LCD displays. Using the audio input configuration of the EMI Spy, we asked 5 participants to point at different devices including monitors and laptops to evaluate the classification accuracy in this scenario. The user was standing at an approximately 1-m distance to two ordinary LCD displays and one laptop computer, while touching the EMI Spy signal input with one hand. The computer and displays were standing next to each other with about 10-cm distance from each other. For touch classification, the user was allowed to touch any position on the screen, without further instruction except for the numbers of fingers to use.

The system was trained with data collected from a single user, with the data collected from the remaining 5 participants being used for testing. The features used are the frequency signature and corresponding amplitude for each of the three displays. We achieved an average classification rate of 90% for pointing across the 3 neighboring devices. We were also able to classify 1- and 2-finger direct touching of the devices with perfect accuracy.

VI. DISCUSSION AND FUTURE WORK

Our method learns a device-specific spectral signature that is independent from background noise. This signature was sufficiently unique for different displays in the environments we tested. We experimented with 6 displays of the same model and discovered that there was sufficient variation across the set to uniquely identify them. We also tested with a number of different models and found the variance to be even greater. We conducted these experiments in a busy, real-world environment across days, and found our method to be robust to a changing EMI background. These results were also independent of the content displayed on the LCD monitors. We cannot conclude from this that all displays would have unique signatures, but are confident that the results demonstrate feasibility. A more extensive study of the differences between LCD drivers would provide a clearer picture of the reliability of our approach in the general case.

Our results show EMI sensing to be a promising technology for supporting proxemic user interaction with displays, and there are many avenues to pursue. On the sensing side, we will add variable gain to make the sensor more sensitive and flexible. We have also identified ground path modeling strategies that would improve system performance; specifically, placing the sensing device between the user and building ground, for instance as in a shoe [19], would likely improve its performance and enable the hands and body to directly probe for signals as sensing antennas.

We would like to replace the microcontroller and Bluetooth with a frequency shifter so that we can use the much lower-cost audio input configuration with any mobile device. Frequency shifting would offer the same tradeoff as intentional aliasing; we would be able to detect higher frequency switching with a lower sampling rate, but would risk cluttering the smaller bandwidth. Moving the device from the duct tape prototype to a more robust form factor would enable further testing and development. We are particularly interested in using EMI Spy to build pop-up interactive environments in public spaces, where electrical appliances could be easily transformed into anchors for user interaction.

Our current peak frequency feature does not account for time-varying or other complex spectral features. We are planning to experiment with temporal features and more sophisticated machine learning techniques that would improve classification performance for a broad range of devices beyond displays. In informal testing, we were able to visually identify many different appliances present in the EMI signal that would not be well-separated by our single feature, but that might be classifiable with others. A typical laptop emits unique signals from its various internal components, like the fan, battery, touchpad, and screen; with better features, it might be possible to capture a fine-grained EMI profile of a given device. This could enable free-air gestural control around the device that would be invariant to the background EMI.

Additional information, such as energy consumption, is also embedded in the EMI signal. Dutta et al. were able to identify energy consumption by counting the cycles in a switching regulator [8]. These effects will also influence the EMI noise coming from the supply. We investigated this phenomenon for different power supplies and were able to
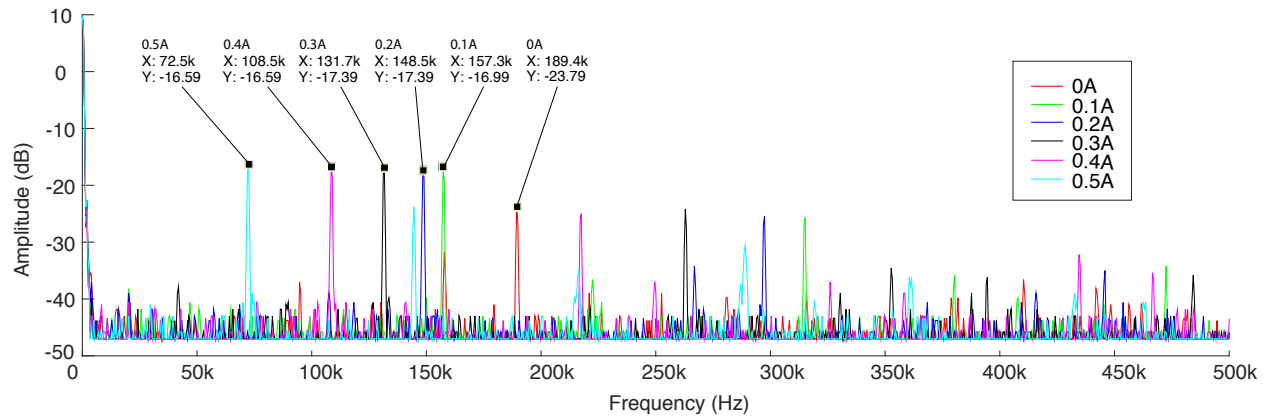
Fig. 6. This graph illustrates EMI spectra from one ordinary 5V/1A power supply under different load. The different colors represent 0A, 0.1A, 0.2A, 0.3A, 0.4A, and 0.5A load current respectively.

detect changes corresponding to the load current. To observe broader frequency range, we used an oscilloscope to capture these data. Figure 6 illustrates the load-dependent variation in peak frequencies for one example supply. We also observed linear changes of amplitude. These changes could potentially be recognized by a machine learning model to predict power consumption of nearby devices.

VII. CONCLUSION

In this paper, we demonstrated the feasibility of a system that senses EMI noise generated from LCD displays to detect users' proximity and gestures to the displays. We successfully identified different displays of the same model from a 2-m distance in an office building. We also demonstrated classification of 1- and 2-finger touch as well as non-contact pointing on ordinary (non-touch) displays. We built a wearable EMI sensing device based on this principle. Furthermore, we showed potential for other applications using EMI sniffing, such as remote load monitoring, and a number of directions for improvement of our low-cost prototyping tool. We see a wearable proxemic sensing approach as comparable to instrumented approaches such as RFID, without the requirement for instrumentation and with the added benefit of gesture recognition. The EMI Spy could become a device that enables a user to determine not only the ID of the device, but also infer aspects of its power consumption, internal operation, or overall functional 'health'.

REFERENCES

[1] Michail Bletsas. The MIT Media Lab's Glass Infrastructure: An Interactive Information System. *IEEE Pervasive Computing*, 11(2):46–49, 2012.
[2] W Buxton. Living in augmented reality: Ubiquitous media and reactive environments. *Video Mediated Communication*, pages 363–384, 1997.
[3] Ke-Yu Chen, Gabe A Cohn, Sidhant Gupta, and Shwetak N Patel. utouch: sensing touch gestures on unmodified lcds. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2581–2584. ACM, 2013.
[4] G Cohn, S Gupta, T J Lee, D Morris, J R Smith, M S Reynolds, D S Tan, and S N Patel. An Ultra-Low-Power Human Body Motion Sensor Using Static Electric Field Sensing. In *12th ACM international conference on Ubiquitous computing (Ubicomp)*, 2012.
[5] G Cohn, D Morris, S N Patel, and D S Tan. Your noise is my command: Sensing gestures using the body as an antenna. *ACM annual conference on Human Factors in Computing Systems (CHI)*, pages 791–800, 2011.
[6] G Cohn, D Morris, S N Patel, and D S Tan. Humantenna: Using the Body as an Antenna for Real-Time Whole-Body Interaction. In *ACM annual conference on Human Factors in Computing Systems (CHI)*, pages 1901–1910, 2012.
[7] David Cranor. Prototouch - a system for prototyping ubiquitous computing environments mediated by touch. Master's thesis, MIT Media Lab.
[8] Prabal Dutta, Mark Feldmeier, Joseph Paradiso, and David Culler. Energy metering for free: Augmenting switching regulators for real-time monitoring. In *Proceedings - 2008 International Conference on Information Processing in Sensor Networks, IPSN 2008*, pages 283–294, 2008.
[9] Nicholas Gillian. Gesture Recognition Toolkit (GRT). http://www.nickgillian.com/software/grt, 2012.
[10] Saul Greenberg, Nicola Marquardt, Til Ballendat, Ro Diaz-Marino, and Miaose Wang. Proxemic interactions. *interactions*, 18(1):9, January 2011.
[11] S Gupta, M S Reynolds, and S N Patel. ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. In *12th ACM international conference on Ubiquitous computing (Ubicomp)*, pages 139–148. ACM, 2010.
[12] Apple Inc. ibeacon for developers. https://developer.apple.com/ibeacon/, March 2015.
[13] G Kortuem, C Kray, and H Gellersen. Sensing and visualizing spatial relations of mobile devices. In *18th annual ACM symposium on User interface software and technology (UIST)*, pages 93–102. ACM, 2005.
[14] P Mistry, S Nanayakkara, and P Maes. SPARSH: touch the cloud. In *ACM Conference on Computer Supported Collaborative Work (CSCW)*, pages 585–586. ACM, 2011.
[15] Joseph A. Paradiso and Neil Gershenfeld. Musical applications of electric field sensing. *Computer Music Journal*, 21:69–89, 1997.
[16] Cati Vaucelle, Hiroshi Ishii, and Joseph Paradiso. Cost-effective wearable sensor to detect EMF. *CHI '09: Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, 2009.
[17] R.G. Vaughan, N.L. Scott, and D.R. White. The theory of bandpass sampling. *Signal Processing, IEEE Transactions on*, 39(9):1973–1984, Sep 1991.
[18] Jamie Zigelbaum and Marcelo Coelho. Six-forty by four-eighty: An interactive lighting installation. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '11, pages 485–485, New York, NY, USA, 2011. ACM.
[19] T. G. Zimmerman. Personal Area Networks: Near-field intrabody communication. *IBM Systems Journal*, 35:609–617, 1996.