# Potentia est Scientia: Security and Privacy Implications of Energy-Proportional Computing

*Shane S. Clark, Benjamin Ransford, Kevin Fu*
Dept. of Computer Science, Univ. of Massachusetts, Amherst
{*ssclark,ransford,kevinfu*}*@cs.umass.edu*

## Abstract

The trend toward energy-proportional computing, in which power consumption scales closely with workload, is making computers increasingly vulnerable to information leakage via whole-system power analysis. Saving energy is an unqualified boon for computer operators, but this trend has produced an unintentional side effect: it is becoming easier to identify computing activities in power traces because idle-power reduction has lowered the effective noise floor. This paper offers preliminary evidence that the analysis of AC power traces can be both harmful to privacy and beneficial for malware detection, the latter of which may benefit embedded (e.g., medical) devices.

## 1 Introduction

The parallel trends of greater energy efficiency and more aggressive power management are yielding computers that inch closer to energy-proportional computing [3], but these improvements also escalate the risk of information leakage via fluctuating power consumption.

In a recent paper on energy-efficiency trends, Koomey et al. point out that energy efficiency (computations per kilowatt-hour) doubled every 1.57 years from 1946 to 2009, with much room for improvement remaining—seven orders of magnitude, extrapolating from an estimate by Feynman—until designs hit theoretical limits [15].[1] Modern processors increasingly employ techniques such as dynamic voltage and frequency scaling (DVFS), clock gating, and dark silicon [6] to decrease their static and dynamic power consumption. Long-term energy-efficiency campaigns such as ENERGY STAR [28] and growing consumer dependence on battery-constrained mobile devices have helped drive this trend.

[1] Transistor counts have doubled at the somewhat slower pace of once every $\sim 1.8$ years [15].

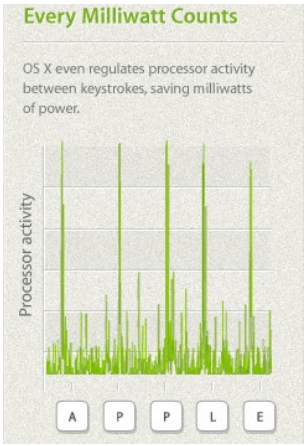


**Figure 1:** The Mac Mini's product page [2] touts energy-efficiency gains that also happen to reveal keystrokes in power traces.

A major side effect of the trend toward greater energy efficiency is that modern computers and operating systems cooperate to reduce idle power consumption. The webpage for the Apple Mac Mini [2] claims that its operating system

> *. . . never misses a power-saving opportunity, no matter how small. It even regulates the processor between keystrokes, reducing power between the letters you type.*

Figure 1 shows the accompanying graphic. Microsoft's "Building Windows 8" blog also addresses OS power efficiency [25], referring to power management as

> *. . . a core OS capability that is critical on any chip architecture and any PC form factor.*

In light of these advances and manufacturers' keen interest in furthering them, this paper takes the position that *improvements in energy efficiency are making whole-system power analysis more effective over time*.



Published in Proceedings of USENIX Workshop on Hot Topics in Security, 2012

Intuitively, reducing idle power consumption lowers the effective noise floor, making activity patterns "stick out" more by raising the signal-to-noise ratio. The increasing leakage of information has both positive and negative repercussions, including an elevated risk of accidental information disclosure (§3) and an enhanced ability to identify certain indicators of malware infection (§4).

## 2 Background on Power Analysis

Power analysis is the process of making inferences from changes in power consumption over time. Given a power trace, one would like to know what information it contains, and how to extract that information. Both laptop and desktop computers use switched-mode power supplies (SMPSes) to convert AC from the power grid to the DC that their components use. Past research has revealed three primary sources of information leakage from SMPSes that manifest themselves in a computing device's power consumption: current fluctuations, reactive power, and changes in switching speed.

**Current fluctuations.** If any particular component conveys information via its power consumption, then that information *may* appear in a power trace, where its power consumption is reflected as part of a single sum—or it may be lost to destructive interference from other signals. Past work on DC power analysis [11, 14] thoroughly explores the phenomenon of information leakage via power consumption, particularly for its value in discovering key material. These studies have employed a variety of measurement techniques that involve varying degrees of contact with the device under test. In the simplest cases, an observer can measure current by placing an inexpensive sense resistor in series with the power supply or a less intrusive Hall effect sensor near the power line; both techniques are equally applicable to AC power measurement.

**Reactive power.** Unlike simple resistive loads (e.g., incandescent lights), the capacitive and inductive components inside SMPSes distort the shape of the sinusoidal AC waveform, drawing the current and voltage waveforms out of phase and returning power to the source in the form of *reactive power*. The amount of reactive power varies with the load, leaking information about whole-system power consumption onto the power line.

**Changes in switching speed.** Like any electrical device that switches on and off, an SMPS emits electromagnetic interference (EMI) that other devices can detect. The switching speed varies with the components' aggregate demand for power. To meet emissions standards (e.g., FCC Title 47 CFR Part 15 [7]), SMPSes contain inductors that filter out noise at frequencies above the voltage regulator's switching frequency. This EMI filtering does not prevent activity information from appearing on the AC power line, as Enev et al. demonstrate [5].

| Year | Processor | Idle (W) | Active (W) |
|---|---|---|---|
| 1997 | Pentium 120 MHz | 39 | 50 |
| 2002 | Pentium 2.8 GHz | 99 | 150 |
| 2011 | Core i5 3.1 GHz | 42 | 65/78/94/110 |

**Table 1:** The minimum and maximum power consumption observed for three desktops, as measured by a P3 Kill-a-Watt meter. For the Core i5, there are four load values because we measured the power consumption while loading each additional core. The load in this experiment was one infinite loop per core.

### 2.1 Information in Power Traces

A power trace is a sequence of $\langle time, power \rangle$ pairs, where a related quantity such as current or voltage may be substituted for power (if appropriate variables are held constant). Many signal-processing approaches may be applicable to power traces; this paper can cover only a subset of them.

Several previous papers have used predictive models to find information in leaked signals. Extending the approach of van Eck [29], Kuhn used precise video-timing information to tune the parameters of an EMI decoder [16, 17]. Vuagnoux and Pasini developed models of keyboards' electromagnetic emanations and matched recorded transmissions against them [30]. Such an approach requires detailed knowledge of the device under test and may be sensitive to peculiarities of make and model, but it provides more insight than untuned approaches when it is applicable.

Under the assumption that modeling a general-purpose computer's state space is infeasible, this paper adopts a black-box approach to identify multiple occurrences of the same signal. Our method performs basic feature extraction and trains a classifier based on supervised-learning algorithms. This approach enables recognition of signals that are consistent but difficult to model.

### 2.2 Capturing Power Traces

A sense resistor, a simple circuit element that could be confined discreetly in a standard AC power outlet or power strip, suffices to make power signals available to conventional measurement equipment. This section describes one of many possible setups to measure whole-system power on the AC line. *Note: We consulted with qualified electrical technicians. Do not try this at home.*

Modern AC outlets have three terminals: *hot*, *neutral*, and *ground*. To measure a power supply's instantaneous current on the hot–neutral circuit, we placed a $0.1\,\Omega$ sense resistor in series with the hot terminal of a standard outlet (Figure 2). For ease of experimentation, we extended the outlet from the wall by stripping one



Published in Proceedings of USENIX Workshop on Hot Topics in Security, 2012

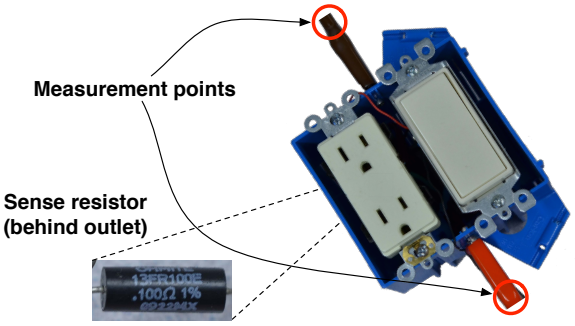


**Figure 2:** An instrumented AC outlet for capturing power traces. A data-acquisition unit connects to measurement points on either side of a 1 cm sense resistor.

end of an extension cord and plugging the other end into an uninstrumented lab outlet. We attached an Agilent U2356A data acquisition unit (DAQ) to the terminals of the sense resistor. The DAQ samples the voltage across its probes and sends the data via USB to another PC (not the computer being measured). For the experiments described later in this paper, we recorded 16-bit samples at a rate of 250 kHz (i.e., 4 Mb/s) to capture workload artifacts occurring at up to 125 kHz.

## 2.3 Previous Approaches

Prior work has demonstrated attacks against user privacy via power analysis by modeling the device under attack or the parasitic modulation of a well-defined signal [11, 14, 16–18, 29, 30]. Some recent work, for example, reveals keystrokes via EMI when provided with detailed knowledge of keyboard signaling protocols, with different attack strategies for PS/2 and USB keyboards [30]. Enev et al. have demonstrated privacy attacks against HDTVs using a model-free classification approach with recurrent neural networks [5]; they apply frequency filtering and a feature-extraction algorithm that automatically identifies changes in SMPS switching speed that correlate with shifts in video brightness. Past work has identified LCD or plasma-based displays, which dominate the power consumption of HDTVs, as power-proportional components [12, 23].

In contrast to this previous work, Section 3.1 evaluates the effectiveness of whole-system power analysis using the tracing setup described above and a straightforward conversion to the frequency domain without filtering.

# 3 Malicious Uses of Power Monitoring

Because of the dynamic range of computers' power consumption is increasing, power signatures of computing activities increasingly leak via the power supply. Using webpage visits as an example of a private activity, this section describes how an eavesdropping adversary with access to a victim's power line (or outlet) can use whole-system power analysis to compromise privacy.

## 3.1 Case Study: Webpage Identification

Using the instrumented outlet described in Section 2.2, we gathered traces from a 2008 MacBook while it loaded 50 popular webpages drawn from Alexa's list of top sites [1]. For each page, we gathered ∼90 samples. We then converted all of these samples into 500-feature vectors using the Fourier transform, with each feature representing a 250 Hz swath of the frequency spectrum. We adopted this simple frequency-domain approach to feature extraction because it allows us to compare traces of different lengths and ignore alignment issues that could arise in the time domain. We used half of the samples from each page to train a set of 50 binary support vector machine (SVM) classifiers (as implemented by the open-source libsvm package [4]) and the other half for testing against these classifiers.

Because our classifier uses standard machine-learning techniques, we use standard metrics from machine learning to evaluate its performance. In the following definitions, *tp* and *tn* refer to true positives and true negatives (correct labelings), and *fp* and *fn* refer to false positives and false negatives (incorrect labelings). *Precision*, $tp/(tp+fp)$, is the fraction of positively labeled examples whose labels are correct. It measures the classifier's ability to exclude negative examples. *Recall*, $tp/(tp+fn)$, is the fraction of all the examples that *should* have been positively labeled that *are* correctly positively labeled. It measures the classifier's ability to identify positive examples (i.e., to correctly match an unlabeled webpage sample to a webpage in the training set).

With ∼45 training examples per webpage label, the SVM classifier achieves on average 87% precision and 74% recall over all webpages, though the classifier's performance varied among the tested webpages. The classifier maintained 100% precision and recall for 5 of the 50 webpages, and ≥ 90% precision and recall for 18 of the 50 webpages. In contrast, the classifier's precision and recall for the worst page, `godaddy.com`, were only 30% and 17% respectively. Figure 3 summarizes the results.

We also gathered single traces of 441 popular webpages not appearing in the training set to ensure that the SVMs were able to reject pages outside of the training set with high probability. This separate test was neces-



Published in Proceedings of USENIX Workshop on Hot Topics in Security, 2012

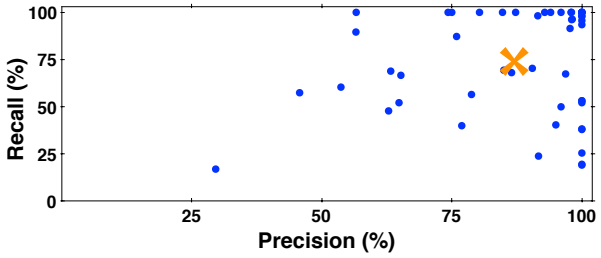


**Figure 3:** Precision and recall for the 50 webpage labels. The arithmetic mean over all webpages is plotted as a large X. The outlying webpage label with the lowest precision and recall is `godaddy.com`, a pathology we believe is due to its heavy use of dynamic content and large download size.

sary because the negative samples appearing in the testing set for the previous experiment were drawn from the same set of pages as those appearing in the training set. It was possible that the SVMs had actually learned to reject only these specific pages rather than all negative samples. We tested all 50 trained SVMs against this set of unknown pages. The 1.6% false-positive rate over all classifiers shows that the trained SVMs were indeed able to reject unknown pages with high probability.

These experiments demonstrate that whole-system power analysis can determine detailed information about a modern computer's workload, including which webpage the browser is loading. Whether it is feasible to make similar inferences about generic computation using similar techniques is an open question.

## 4 Constructive Uses of Power Monitoring

"Forever-day" vulnerabilities—those that the vendor does not plan to fix—in civil and industrial control systems have recently put unpatchable software in the spotlight. A key question for system and network administrators is whether they can detect intrusions without compromising these machines' operations or violating terms of use. This section describes how whole-system power analysis can help identify anomalous behavior. Such analysis will be useful in concert with conventional network intrusion-detection systems.

Most of the published work on constructive uses for power monitoring leverages nonintrusive load monitoring (NILM) techniques or smart meters to accurately report load changes for the sake of personal energy-use monitoring [9, 22]. Our power-analysis approach instead focuses on understanding the internal state of a complex, standalone, integrated device. There is also published work on power analysis for smartphone malware detection [13, 20]. Smartphones are convenient platforms for power analysis because they tend to comprise simpler hardware and software stacks than do traditional computers; many expose APIs for battery monitoring.

One particular class of devices stands to derive outsize benefits from constructive power analysis. Many medical devices used in hospitals and other sensitive domains are based on commodity hardware and software and are networked, exposing them to the same malware threats as consumer devices; however, manufacturers advise against routine patching, ostensibly for compliance with regulations that require re-certification for every significant change.

According to testimony by Roger Baker, Assistant Secretary for Information and Technology at the U.S. Department of Veterans Affairs (VA), more than 122 medical devices at VA facilities were found to be infected with malware in the 14 months preceding May 2010 [19]. Baker said:

> *The major challenge with securing medical devices is that, because their operation must be certified, the application of operating system patches and malware protection updates is tightly restricted.*

While few patches require FDA clearance [27], anecdotal evidence suggests that such updates are a gray area for customers and manufacturers—resulting in customers adopting a hands-off attitude toward devices ostensibly under their control.

A tempting solution to the malware-propagation problem is to mandate that medical equipment never come in contact with the Internet, but some devices require network connectivity. USB drives used by employees or field technicians are also a threat. The VA report from which Baker drew malware statistics includes data on infection sources: while 65% of the incidents had an unknown infection vector, 10% were traceable to local networks via repair technicians, USB drives, or LAN sharing [24]. With respect to malware infections more generally, Symantec reports that 72% of the malware they analyzed in 2009 could propagate via USB drives or other "sneakernet" mechanisms [8] against which network intrusion-detection systems are ineffective.

### 4.1 Case Study: Detecting Embedded Malware

Concerning embedded devices such as medical equipment, oscilloscopes, and industrial control systems, we make a key observation about the suitability of power analysis for malware detection: while these devices use commodity hardware and software, they are not intended to be used as general-purpose devices. This fact simplifies the task of identifying aberrant behavior because there is little variety in expected workloads. Just as an



Published in Proceedings of USENIX Workshop on Hot Topics in Security, 2012

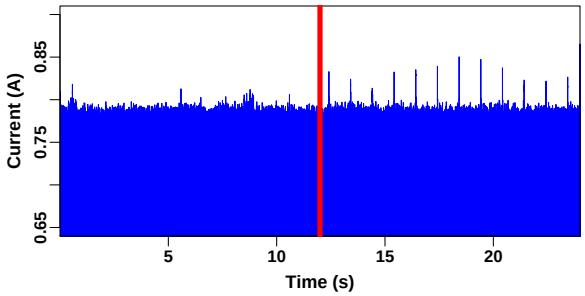


**Figure 4:** Side-by-side traces of current consumed by an oscilloscope running Windows XP. The left half of the plot shows idle consumption before infection. The right half shows idle consumption after infection—with distinctive spikes every second caused by resident malware.

oscilloscope should be used only to measure signals, rather than to browse the Internet, medical-imaging systems or industrial control systems have a small number of intended uses. Each intended use should correspond to one, or few, power consumption profiles, suggesting that whitelisting approaches have promise in this domain.

To test whether power analysis can be applied to malware detection for embedded devices, we performed a simple experiment using an embedded device that runs standard software and the measurement technique described in Section 3. We took power traces of an Agilent Infiniium 54832D MSO oscilloscope, which runs Windows XP on an Intel Pentium 3 platform with additional custom hardware and software. After establishing baseline power consumption by recording traces while the scope was both idle and taking measurements, we intentionally infected it with malware by visiting the top two URLs from a honeypot site [21] and allowing each site to run an executable. Finally, we gathered a second set of idle traces after infection.

Inspection of the before and after power traces revealed obvious differences in power consumption patterns. The malware running on the scope, later identified as Kolab and PWS-Zbot by McAfee Enterprise Security, created a transient power spike once per second while the scope was idle. The power spikes were brief and infrequent enough that they did not have a significant effect on the mean or variance of the power consumption, but they stood out clearly in a time-series plot (Figure 4).

Other types of malware may cause similar power anomalies. First, it is typical for worms designed to create a botnet to regularly search for or communicate with command and control servers. Some types of malware also poll for the presence of antivirus software or for certain conditions, such as password entry, that will trigger a response. Finally, some malware has the explicit goal of hijacking an infected computer's computational resources. As an example, researchers last year identified a trojan that mined Bitcoins on infected hosts—even including GPU support for faster mining [26].

## 5 Open Problems

Research challenges remain in the area of whole-system power analysis. The effectiveness of the simple black-box power-analysis method described in Sections 3 and 4 suggests that more sophisticated methods might yield more information. Candidates include:

- Statistical analysis to automatically identify salient features for classification;
- An adaptation of behavior-based malware detectors such as Panorama [31] to accept power traces as input signals;
- Correlating AC power traces gathered at the plug with DC power traces gathered between the power supply and components, to identify which components' power consumptions scale most closely with workload and thus leak the most information;
- Modeling-based approaches in which models mapping activity to power consumption would inform trace classification.

**Threat modeling.** The threat model implicit in this paper posits an eavesdropper who can place a resistor in series with a victim's power supply. However, even covertly modifying an electrical outlet may be infeasible in many scenarios. A more realistic adversary might instead have access to the same AC power circuit (the NILM scenario [10]) or be stationed nearby. Future work could investigate the degree to which these alternative eavesdropping modalities are useful for activity recognition, and whether the effective classifier described in Section 3 remains effective.

Future work may need to consider adversaries who are aware that their activities may be detected via power analysis. Such adversaries may attempt to conceal their activities by, for example, waiting for power-hungry activities to begin or activating at random times. Other than defense-in-depth approaches, how to respond to this risk is an open question.

## 6 Conclusion

As an inevitable consequence of energy-proportional computing, modern computer components leak increasing amounts of information about their internal states onto the power line. Paradoxically, this information can harm privacy but may help at detecting malicious software. Like a thermometer for a human patient, constructive power analysis can reveal indicators of malicious activity or "sneakernet" infections of embedded devices.



Published in Proceedings of USENIX Workshop on Hot Topics in Security, 2012

## Acknowledgments

We thank Dan Holcomb for feedback on early drafts and Quinn Stewart for feedback on late drafts. This material is based upon work supported by the National Science Foundation under grants CNS-0923313, CNS-0845874, and two Graduate Research Fellowships. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This publication was made possible by Cooperative Agreement no. 90TR0003/01 from the Department of Health and Human Services. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the HHS. This work was funded by a Sloan Research Fellowship.

## References

[1] Alexa Internet, Inc. Alexa top 500 global sites. `http://alexa.com/topsites/global`. Loaded Sept. 2011.

[2] Apple Inc. Mac mini – The most energy-efficient desktop computer. `http://www.apple.com/macmini/environment.html`. Loaded May 1, 2012.

[3] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *Computer*, 40:33–37, Dec. 2007.

[4] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3), Apr. 2011.

[5] M. Enev, S. Gupta, T. Kohno, and S. Patel. Televisions, Video Privacy, and Powerline Electromagnetic Interference. In *Proc. ACM Conference on Computer and Communications Security*, CCS '11, Oct. 2011.

[6] H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger. Dark silicon and the end of multicore scaling. In *Proc. ACM International Symposium on Computer Architecture*, ISCA '11, June 2011.

[7] Federal Communications Commission. Code of Federal Regulations, Title 47, Part 15, Sections 101–103, Oct. 2010.

[8] K. Haley. Sneakernet revisited. `http://www.symantec.com/connect/blogs/sneakernet-revisited`, Aug. 2010. Loaded June 2012.

[9] G. W. Hart. Residential energy monitoring and computerized surveillance via utility power flows. *IEEE Technology and Society Magazine*, June 1989.

[10] G. W. Hart. Nonintrusive appliance load monitoring. *Proc. IEEE*, 80(12):1870–1891, Dec. 1992.

[11] T. Kasper, D. Oswald, and C. Paar. EM Side-Channel Attacks on Commercial Contactless Smartcards Using Low-Cost Equipment. In *Workshop on Information Security Applications*, WISA '09. Aug. 2009.

[12] M. Kazandjieva, B. Heller, P. Levis, and C. Kozyrakis. Energy dumpster diving. In *Workshop on Power Aware Computing and Systems*, HotPower '09, Oct. 2009.

[13] H. Kim, J. Smith, and K. G. Shin. Detecting energy-greedy anomalies and mobile malware variants. In *Proc. International Conference on Mobile Systems, Applications, and Services*, MobiSys '08, June 2008.

[14] P. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Advances in Cryptology*, CRYPTO '99, Aug. 1999.

[15] J. G. Koomey, S. Berard, M. Sanchez, and H. Wong. Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing*, 33:46–54, Mar. 2011.

[16] M. G. Kuhn. Electromagnetic Eavesdropping Risks of Flat-Panel Displays. In *Workshop on Privacy Enhancing Technologies*, PET '04, May 2004.

[17] M. G. Kuhn and R. J. Anderson. Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations. In *Information Hiding*, Apr. 1998.

[18] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong. Power Signature Analysis. *IEEE Power and Energy Magazine*, 1(2):56–63, Mar. 2003.

[19] N. Lewis. VA security compromised by medical devices. `https://www.informationweek.com/news/healthcare/security-privacy/225200097`, May 2010. Loaded June 2012.

[20] L. Liu, G. Yan, X. Zhang, and S. Chen. Virusmeter: Preventing your cellphone from spies. In *Recent Advances in Intrusion Detection*, volume 5758 of *Lecture Notes in Computer Science*, Sept. 2009.

[21] NovCon Solutions LLC. Novcon minotaur analysis system. `http://minotauranalysis.com/exetweet/`, Loaded Apr. 2011.

[22] S. N. Patel, T. Robertson, J. A. Kientz, M. S. Reynolds, and G. D. Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In *Ubicomp 2007: Ubiquitous Computing*, volume 4717 of *Lecture Notes in Computer Science*, Sept. 2007.

[23] PCSTATS. Blackle vs. Google Monitor Power Consumption Tested. `http://www.pcstats.com/articleview.cfm?articleID=2649`, Loaded Apr. 20, 2012.

[24] L. Sherrill. Medical device infection data. Personal communication, Jan. 2012.

[25] P. Stemen. Building a power-smart general-purpose Windows. `http://blogs.msdn.com/b/b8/archive/2011/11/08/building-a-power-smart-general-purpose-windows.aspx`, Nov. 2011. Loaded June 2012.

[26] Symantec Corporation. Trojan.Badminer technical details. `http://www.symantec.com/security_response/writeup.jsp?docid=2011-081115-5847-99&tabid=2`, Aug. 2011. Loaded June 2012.

[27] A. Taylor. Virtual patient safety: Worms, viruses, and other threats to computer-based medical technology (presentation). ECRI Audio Conference, Nov. 2003.

[28] United States Environmental Protection Agency. ENERGY STAR program requirements for computers. `http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/Version5.0_Computer_Spec.pdf`, July 2009. Loaded June 2012.

[29] W. van Eck. Electromagnetic radiation from video display units: An eavesdropping risk? *Computers & Security*, 4:269–286, Dec. 1985.

[30] M. Vuagnoux and S. Pasini. Compromising Electromagnetic Emanations of Wired and Wireless Keyboards. In *Proc. USENIX Security Symposium*, Security '09, Aug. 2009.

[31] H. Yin, D. Song, E. Manuel, C. Kruegel, and E. Kirda. Panorama: Capturing system-wide information flow for malware detection and analysis. In *Proc. ACM Conference on Computer and Communications Security*, CCS '07, Oct. 2007.