EMPower: Detecting Malicious Power Line Networks from EM Emissions

Richard Baker, Ivan Martinovic

EMPower : Detecting Malicious Power Line Networks from EM Emissions

Richard Baker[0000-0001-8215-1053] and Ivan Martinovic

Department of Computer Science, University of Oxford
first.last@cs.ox.ac.uk

Abstract. Power line communication (PLC) networks are commonplace today, particularly within consumer home environments. They permit simple plug-and-play networking by leveraging the existing electrical wiring in buildings to transmit data as well as power. However, the ubiquity of this networking opportunity is often overlooked and permits an attacker, with only one-time access to an environment, to establish free, unmonitored and high-bandwidth network connectivity to the victim. However, the unsuitability of power wiring for high-frequency signalling means that PLC leaks radiated emissions. We demonstrate the detectability of this phenomenon in a real-world setting and introduce EMPower; a system that identifies the presence of hidden power line networking from analysis of the characteristic EM emissions in the frequency and time domains. We demonstrate the effectiveness of EMPower using a COTS radio receiver — identifying the presence of a network near-perfectly within the same room, even when idle, and with 74.6% accuracy two rooms away and on a different floor. Thus realising the capability to monitor an environment for unwanted power line networks.

1 Introduction

Power line communications (PLC) technologies have been used for over 70 years. Whilst originally employed only for long-distance measurement and control over high-voltage distribution lines, advances in technology, increased demand from consumers and successful standardisation initiatives have permitted today's manufacturers to build interoperable, plug-and-play equipment that can communicate throughout most buildings at hundreds of megabits of data rate, using the building's existing power distribution infrastructure.

Today, power line adaptors are widely-available and inexpensive devices that are commonly deployed to overcome a lack of purpose-built networking infrastructure or to mitigate poor wireless connections. The HomePlug Powerline Alliance claimed in 2016 that 220 million devices were in use worldwide [6]. But just as these devices permit legitimate users to network devices, they also permit malicious users to construct networks at will that can easily go unnoticed in buildings that are increasingly populated by small, anonymous, electronic devices. While wired data networks are segregated and physically protected, and

while wireless networks are policed for rogue access points [16], little if any consideration is given to the openness of power networks. They present an easy target for a potential attacker and an attractive one for establishing access far into a secured area with much greater bandwidth than a homemade covert radio channel permits.

However, re-purposing power distribution wiring for high-frequency signalling is not without its own problems. In particular signal leakage, both by conduction and radiation, is a concern. We consider a beneficial use of these radiated emissions for a user seeking to protect themselves from the deployment of a malicious power line network on their premises. We present a system that permits the rapid detection of power line networks being deployed maliciously in the vicinity by detecting these radiated emissions.

In particular we make the following contributions:

- Highlight the security threats posed by unmonitored power networks — with an example attack implementation
- Introduce novel techniques that permit the detection of power line networks from EM measurements using time- and frequency-domain analysis
- Compare the relative performance of each technique in detecting the presence of power line communications in a real-world context

2 Security Risk

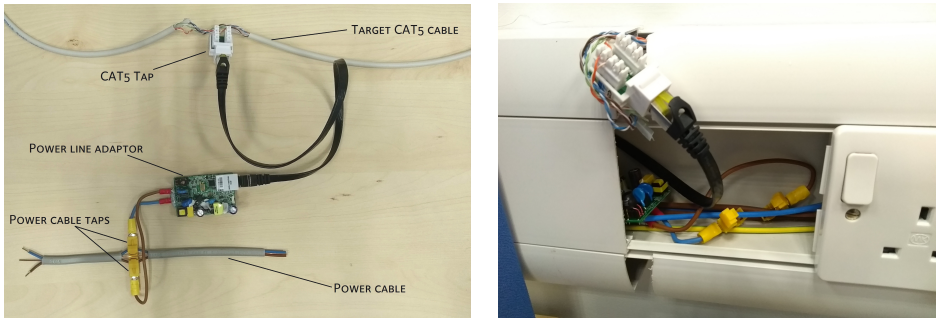

(a) A prototype covert attack device.

(b) The prototype attack device in situ inside service trunking on an office wall.

Fig. 1. A mock-up of a covert traffic capture and bulk data exfiltration attack using a maliciously-installed power line network.

Many studies of PLC LAN usage have been published. In general, they considered various legitimate deployments of power line networks and how susceptible they would be to eavesdropping, man-in-the-middle attacks and recruitment

of constituent devices into rogue networks (whether with malicious intent or by accident) and presented good arguments as to why these risks were well-controlled [10][9].
We are more concerned with a simple, secondary problem:

No one monitors their electrical network for the presence of a hidden data network

Electrical power is available everywhere in a modern building and usually segmented for supply management and safety reasons, rather than along security boundaries that exist for data networks. To demonstrate that this is a problem, we present the following attacks:

Minimum-effort attack An attacker connects a purchased adaptor and a length of CAT5 cable directly into a target computer or networking device. This can be achieved very quickly if the connection points are nearby and, while it is far from subtle and depends upon a convenient exposed network port, it immediately achieves the objective. In an out-of-sight area the adaptor can go unnoticed for a long time. Indeed, with the plethora of anonymous devices that populate modern buildings, even if a casual observer spots the adaptor, they may well not conclude that it was put there with malicious intent. It was trivial to demonstrate this at multiple locations within our building (a modern office), that crossed security boundaries within the same floor, although not between floors as each floor is served by a separate distribution board, the effects of which in tandem with the generally-high noise levels overwhelmed the signal.

Covert attack Figure 1a shows an example attack device constructed by the authors; using a single-board power line networking implementation taken from a Technomate TM-200 HP adaptor, with the two power lines connected via short leads to insulation-piercing crimp connectors. A tiny section of cable was connected from the adaptor's RJ45 port and into a CAT5 punchdown jack. Figure 1b then shows a mockup of an attack being performed. The attacker has unclipped a section of trunking to reveal the power and data cable runs. They tap the CAT5 cable in the normal way; by removing the outer sheath to reveal the data wires and then pushing the wires down into the jack. The power lines are tapped similarly (although in this case they are already separated). The device is hidden in the cable cavity and the trunking replaced. With proper installation there is no service interruption for either the data or power connections and the attack is completed in a couple of minutes. The unit is powered from the mains connection and can provide passive monitoring of traffic and forward it via the power line network in perpetuity.

3 Threat Model

The attacker attempts to establish connectivity to a target host or network, for bulk data exfiltration, real-time traffic monitoring or as a platform for further

attacks. The attacker has temporary access to the premises, for example as an insider or a brief visitor (e.g., a courier or cleaner), but intends to establish persistent access to devices or networks that are sited within a restricted area and they install a power line network adaptor to achieve this goal. They may connect the adaptor directly to a target host or to an exposed network port, or alternatively in a more complex fashion such as that described in Section 2 above. The attacker can then access the network from a location that is electrically-close to the target, despite being isolated from the conventional network, such as a reception area, an office-building café or a nearby house. They can do this without establishing a rogue wireless network for which commercial detection technologies are widespread and which may be too weak to communicate with successfully from their desired attack location.

We assume that no legitimate power line network exists in the target premises and that the attacker is restricted to commercially-available power line hardware. The standards for all broadband PLC are highly complex and in the case of HomePlug, the authors are aware of no implementation of any kind outside manufacturer members of the HomePlug Powerline Alliance.

4 Related Work

Power line security was studied along with the development of short-range, broadband systems in the late 1990s and early 2000s. Unintentional emissions were one of the main risks considered in security analyses, but these were almost invariably *conducted emissions* — the risk that the power sockets next door can also reach your network. Such analyses have dwelt primarily upon data confidentiality, the protocols for establishing networks [10] and the ability for users to administer their devices securely [9]. More recently, practical attacks have also been noted against weak implementations [13][4].

Work on the *radiated emissions* from power line communication has been largely absent from the security literature also, on the basis that the PLC channel presents such adverse signalling conditions that data recovery from elsewhere than the intended receiver would be infeasible (to say nothing of the payload encryption) [9]. However, there is considerable work on PLC radiated emissions for electromagnetic compatibility [17][11], and the security implications of unintentional emissions in general are a rich field of study.

Considerable attention has also been paid in the security community to the problem of rogue wireless access points and it is still receiving attention today both in academic circles [16] and for cyber security practitioners [2]. Wireless intrusion detection systems (WIDS) are de rigueur in modern wireless deployments, in an attempt to mitigate threats of rogue access points, banned devices, unauthorised ad-hoc networks or network bridging. These systems are powerful tools in securing wireless networks but are purpose-built for specific wireless technologies (usually 802.11 Wi-Fi) — both encouraging potential attackers away from those technologies and providing no protection if that occurs.

5 Background

Power line communication (PLC) systems have existed since 1838 in principle and the 1950s in practice [1]. Local-area communication variants appeared around the turn of the millennium. Some, such as X10, Universal Powerline Bus and latterly HomePlug GreenPHY permit robust, low-bandwidth communication for home automation, IoT and electric vehicle applications. However, the most well-known and commercially-successful application has been for broadband local-area networking; complementing or competing with common Ethernet-over-UTP or Wi-Fi deployments. The appeal of providing data networking over ubiquitous power-distribution wiring (the 'no new wires' benefit), while retaining some of the range and perceived security benefits of wired infrastructure has fuelled adoption. However, power networks were never designed for high-frequency signalling. They are unshielded (permitting radiated emissions and susceptibility thereto) and filled with impedance mismatches, impedance variation and noisy electrical devices. As such, they are a very challenging environment for communication; exhibiting frequency-selective fading, plentiful multipath interference and non-linear distortion — more akin to urban wireless communication than to purpose-built, wired data networks [8]. The dual effect of noise intrusion is that PLC signalling is also prone to leak out, by conduction and radiation. In general, higher-frequency signals radiate better and by signalling over a large bandwidth, broadband PLC adaptors will invariably produce at least some observable radiation from somewhere in the spectrum, where part of the local electrical wiring acts as a convenient, albeit unintentional, antenna. The potential problems caused by these emissions are widely acknowledged, and are the subject of academic work and regulatory intervention to ensure that unintended emissions are minimised [17][11].

The dominant, standardised, broadband LAN PLC technologies are the HomePlug and G.hn families (ratified in IEEE1901 and ITU G.9960 overarching standards respectively). Both standards make use of orthogonal frequency division multiplexing (OFDM) over bandwidths up to 100MHz, permit maximum theoretical data rates over 1Gbps and implement coexistence mechanisms for operating several virtual networks over the same physical media [12]. Contemporary devices advertise operating distances in domestic settings of up to 300m [15] and are often employed as Wi-Fi extenders to mitigate problems of poor coverage.

We concentrate in this work on devices implementing the HomePlug family of standards, in particular HomePlug AV. This selection is due to HomePlug AV introducing the vast majority of functionality that persists in later standards, making our findings generalisable to them. We discuss this in Section 9. HomePlug AV adaptors are available as host NICs [3], Ethernet bridges [15] and wireless access points [14], with or without a power pass-through capability.

5.1 HomePlug AV

HomePlug AV implements OFDM signalling over a frequency range of 1.8MHz — 30MHz. It distributes a total of 1,155 subcarriers over that range [5]. The

choice of OFDM in the standard's design was to mitigate the challenges of the medium discussed above. Individual stations exchange sounding packets to estimate the channel characteristics on each subcarrier and compute *Tone Maps*, which are used to adapt the number of bits sent per symbol on each. To limit electromagnetic compatibility issues, the use of spectral masks is mandated in the HomePlug AV specification [7]. The spectral mask is implemented by disabling a set of subcarriers from being used for signalling at all, creating gaps in the spectral usage akin to those created by bandstop or 'notch' filters. The notches correspond to ten amateur radio bands common across the world, as defined by the International Amateur Radio Union (IARU) and adhered to in the majority of spectrum enforcement jurisdictions. Usage of this spectral mask is hardcoded into power line adaptors and experimental results from emissions testing in [17] has shown that under lab conditions, emissions are consistent with these expectations. While some adaptors permit the addition of further spectral masks, removing notches is not possible without substantial modification to the hardware implementation of the device.

The HomePlug AV standard considers that individual logical networks may not be isolated, due to signal leakage. As such it implements a virtual network mechanism, with each virtual network electing one adaptor as the *Central Coordinator* to manage it. Virtual networks have a pre-shared *network membership key* (NMK) that is the basis for confidential communication. From this NMK is computed a *network encryption key* (NEK), that changes periodically and is used to encrypt data payloads with 128-bit AES in CBC mode. The standard also mandates a number of higher-level management systems as well; for quality-of-service provision, cohabitation with other virtual networks and the extension of the network via relays, the Central Coordinator manages these also. Communication to manage the virtual network, exchange Tone Maps between every pair of devices and operate inter-network cohabitation protocols, ensure a consistent minimum level of traffic is always present if a device is connected and powered.

The lowest-level transmission structure defined in the standard is the *PHY-layer protocol data unit* (PPDU). The PPDU is the concatenation of a preamble, frame control data and an encrypted payload consisting of a series of OFDM symbols encapsulating data from the rest of the network stack.

6 Designing EMPower

Figure 2 shows the EM emissions of a PLC network, as detected at short range in a normal office environment using a USRP N210 SDR with a short wire antenna. While this is not an ideal antenna, its deficiencies are minor compared to the more pronounced impact from the variability of effective radiating wiring in the building for each frequency across the observed band. The flat spectral occupancy observed in ideal conditions [17] has been corrupted substantially even at close range as in Figure 2a. Even a short distance from the source, as in Figure 2b (at 12.9m), the spectrum is barely distinguishable from the background and not recognisable to the eye. Indeed the wiring to which the power line adaptors are

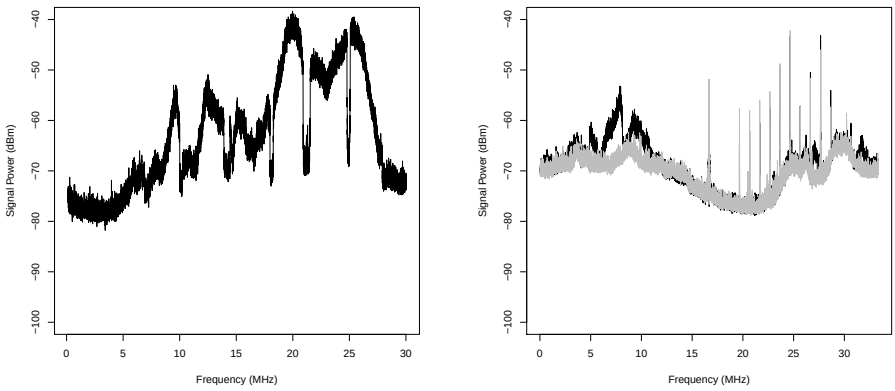

(a) Observed emissions at close range.

(b) Observed emissions three rooms away (from Location 3), with background level in gray.

Fig. 2. Observed radiated emissions from TP-Link TL-PA511 power line adaptor.

connected greatly affects the radiated emissions. The presence of certain types of common devices, such as halogen bulbs, switching power supplies, compact fluorescent lamps and dimmer switches, all introduce noise and poor impedance matches that set up various sections as antennas at different frequencies[17].

While the emissions from an adaptor vary by environment, they are broadly consistent between adaptors. We compared the emissions of a TP-Link TL-PA511, TP-Link TL-WPA281 and Technomate TM-200 HP (all HomePlug AV adaptors) and a Sumvision SVW1000 (HomePlug AV2). In each case the pattern was very similar, with only minor amplitude changes between them.

7 The EMPower Detector

EMPower performs analyses of the received signal in the frequency and time domains. Examining spectral content can provide useful information even with much of the signal attenuated, while time domain analysis can permit insight into the protocol taking place. Figure 3 shows the structure of the system. A received signal is first normalised by an automatic gain control implementation, before the values are passed to each processing chain for analysis.

7.1 Frequency Domain

The frequency domain method detects the presence of a spectral mask. With a received signal filtered to the HomePlug AV band, short-term Fourier transform (STFT) is computed at regular intervals and the signal power calculated. For

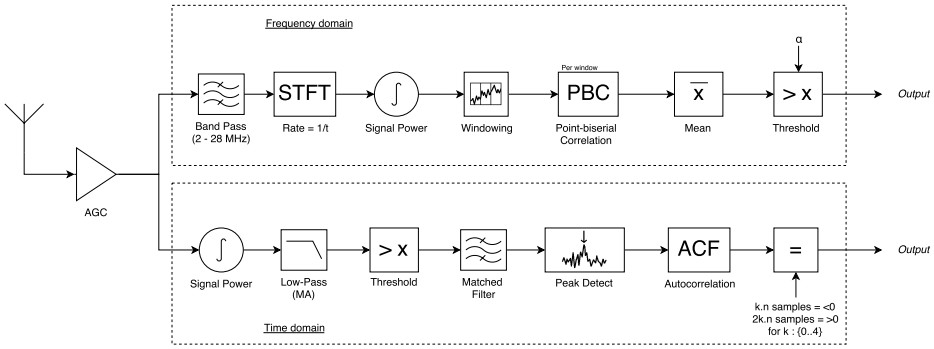

Fig. 3. Block diagram of the system, showing the full processing chain for frequency- and time-domain analyses.

the observed bandwidth w of the signal, the STFT provides an approximation of the power spectral density across b frequency bins, over a brief period $t = \frac{1}{F}$ where F is the STFT rate in Hz. The detector maintains a set of the maximum observed power values in each frequency bin for the observation period T. At each STFT output, the maxima are updated, until T has elapsed and they can be analysed further.

The band is split into windows, to combat the substantial variation in emissions across the full bandwidth. The task of the detector is to ascertain the presence of the spectral mask, so local variation is far more important than the total variation across the band. The window size is taken as the smallest size for which the spectral mask has a notch in every window. The spectral mask is represented as a binary vector in which `0` and `1` indicate the low (-80dBm) and high (-50dBm) signal levels respectively. The measured values are then compared to the template by calculating the point-biserial correlation coefficient within each window. The point-biserial correlation coefficient is specifically designed for comparing continuous values against binary classifications; here the signal powers against the two expected levels in the spectral mask. The mean of the correlation coefficients for each window is taken and used as a score for the presence of a power line adaptor. A score above a given threshold α; that is a signal sufficiently similar to the template, is considered to be a detection.

7.2 Time Domain

The time-domain method makes use of the PPDU preamble, which exhibits a reliable structure and redundancy for robustness against noise. EMPower thus performs a similar process to that of a normal receiver, adapted to the conditions of radiated emissions. The power of the time-domain signal is calculated and then passed through a short moving average, to reduce high-frequency noise from the amplifier and receiving radio. The signal is then thresholded to exclude baseline noise and when it exceeds the threshold, a section the length of a preamble is

passed to a matched filter built from a preamble template. A matched filter is optimal in separating a known signal from white noise, so this step assists in pulling weak radiated emissions out of the noise. A peak detection algorithm runs on the output of the matched filter to find potential preambles and each is then passed through the autocorrelation function. As the preamble repeats with a known period, the result of autocorrelation is tested at these points. A genuine preamble displays a strong correlation to a copy of itself shifted by the known interval and a strong negative correlation to a copy of itself shifted by half the interval. If the autocorrelation displays positive and negative values at these points then the system can conclude with confidence that a preamble has been detected.

8 Evaluation

8.1 Experimental Setup

A staged attack was conducted by placing a PLC adaptor in a series of locations within a normal, shared office building. The target (1 - 4) and attacker (A - C) locations can be seen in Figure 4, as can the static position of EMPower. The adaptor (a TP-Link TL-PA511) was placed in a power socket and a Raspberry Pi connected to it. The attacker then inserted another PLC adaptor (a TP-Link TL-WPA281) at publicly-accessible locations on the same floor of the building. EM emissions were tested with the adaptors switched off to leave only *Background*, with the adaptors switched on but the network *Idle*, and finally with the attacker running the `iperf` network benchmarking utility at the *Max* bandwidth the connection would support. The Raspberry Pi acted as the `iperf` client (sender) so as to simulate a bulk data exfiltration. The Background state was observed for a period of two minutes and the Idle and Max states for one minute apiece, to provide equal numbers of observations in positive and negative states.

EM emissions were collected using a USRP N210 software-defined radio, a short wire antenna and a pre-amplifier. The USRP was tuned to a centre frequency of 16.68MHz and collected with 33.3MHz of bandwidth. Samples were captured using a simple `GNURadio` flowgraph and then processed in `R` for each detection method described in Section 7.

For the frequency-domain approach, the STFT rate was 120Hz, corresponding to $t = \frac{1}{120}$ with maxima being tracked over a period of $T = 1$s. The STFT had a width of $b = 16,384$ frequency bins. The band pass filter rejected bins that were outside the HomePlug AV bandwidth (970 below 2MHz and 2,622 above 28MHz) and the remainder passed along the frequency-domain processing chain. The time-domain approach used a 900 sample maximum lag for the autocorrelation function and searched for 4 pairs of peaks and troughs.

8.2 Detection Accuracy

Figures for detection accuracy, precision and recall with each method can be seen in Table 1. EMPower performed well across the tested locations, although differently for each approach. Peak accuracies were 97.8% using frequency-domain

Fig. 4. Floorplan of the target building, showing the public locations (shaded red) and the private locations (white). The markers with dashed lines are on the floor below.

detection and 100% using time-domain detection at close range, whilst minimal accuracies were 74.6% and 50.2% respectively. The wide variation in accuracy is due to the complex factors discussed in Section 6 above; a combination of distance, data rate and noise at the transmitter and receiver. Both approaches were affected by these factors, although the effect was different for each. The frequency-domain approach exhibited consistently high (>89%) precision even at larger distances, although its recall fell as distance increased. In other words, it rarely made a false detection but its ability to detect networks fell at longer range. However, this approach still performed moderately well in the most challenging conditions examined; communication on the floor below. By contrast, the time-domain approach performed near-perfectly at close range, but the performance degraded far more quickly as conditions deteriorated. At distance, even on the same floor, the recall of the time-domain approach had fallen below 32% and for the attack on the floor below it was effectively no better than random. It appears that the two approaches provide complementary properties that can contribute to better combined detection than either method achieves individually.

Higher data rates over the network led to better performance in every case, but even the minimal management traffic on an idle network was enough in most cases. Considering that each result used to calculate the performance metrics

represents a single T period (only 1s of elapsed time), this means that an adaptor within range *would be detected mere seconds after being powered on.*

Table 1. Detection results. Distances are taken from target to detector. Accuracy metrics are shown for each network state and aggregated across all three.

			Frequency Domain			*Time Domain*		
Target	**Attacker**	**State**	**Accu. (%)**	**Prec. (%)**	**Rec. (%)**	**Accu. (%)**	**Prec. (%)**	**Rec. (%)**
1 (at 2.2m)		None	90.1			100		
		Idle	85.2			98.4		
		Max (33.6Mbps)	100			100		
	A	*Aggregated*	*91.2*	*89.8*	*92.4*	*99.6*	*100*	*99.2*
		None	95.5			100		
		Idle	85.5			100		
		Max (36.1Mbps)	100			100		
	B	*Aggregated*	*94.0*	*95.7*	*92.6*	*100*	*100*	*100*
2 (at 2.1m)		None	98.2			100		
		Idle	94.4			100		
		Max (54.8Mbps)	100			100		
	A	*Aggregated*	*97.8*	*98.2*	*97.4*	*100*	*100*	*100*
		None	100			100		
		Idle	69.1			100		
		Max (45.4Mbps)	100			100		
	B	*Aggregated*	*92.5*	*100*	*84.4*	*100*	*100*	*100*
3 (at 12.9m)		None	100			100		
		Idle	82.1			1.6		
		Max (35.6Mbps)	100			51.7		
	A	*Aggregated*	*95.6*	*100*	*90.8*	*63*	*100*	*25.6*
		None	99.1			100		
		Idle	73.2			7.4		
		Max (51.7Mbps)	100			52.5		
	B	*Aggregated*	*93.0*	*99.0*	*86.8*	*66.7*	*100*	*31.0*
4 (at 9.9m)		None	99.2			99.2		
		Idle	4.8			1.6		
		Max (42.1Mbps)	100			1.8		
	C	*Aggregated*	*74.6*	*98.3*	*49.6*	*50.2*	*66.7*	*1.7*

As the frequency-domain method makes use of a threshold (α) in the final decision-making, we analysed the effects of varying this threshold. Figure 5 shows the receiver operating characteristic (ROC) curve for the detector, computed over all the test locations and network states. The ROC curve shows the rate of successful detection against the rate of false detections. Ideal performance is for the true-positive rate (TPR) to reach 1 while the false-positive rate (FPR) is still 0. The best performance on this curve (F-Score = 0.905) is achieved with α set at -0.038. The values in Table 1 are with that threshold value.

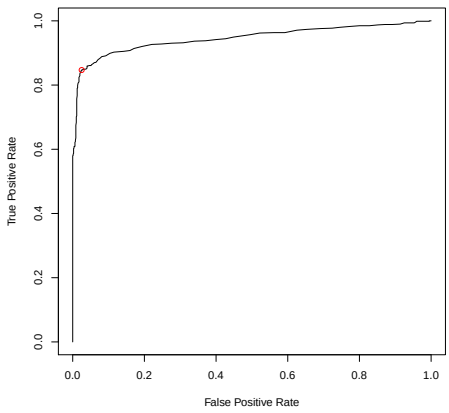

Fig. 5. ROC curve, computed across all test locations and network states.

9 Discussion

An attacker clearly would like to avoid their malicious deployment being detected, and they have two broad approaches. They can reduce emissions by reducing either the transmission power or the utilised bandwidth. There are limitations to both approaches, however. Reducing the signalling power does affect the level of EM emissions produced by the adaptors, but also their signalling range. The further an attacker must be from their target, the less scope they have to minimise the signal power. Alternatively, reducing the bandwidth means disabling further subcarriers than those filtered as standard. In effect the attacker must add additional filtering notches (assuming their adaptors support this) and thereby reduce the correlation between the detected emissions and the template. However, the speed of detection means that an attacker must do this in advance and as [8] notes from a large scale field test, "[n]o line is like the other"; not only does the attack reduces their available bandwidth, they may even risk losing connectivity entirely if the remaining subcarriers are overwhelmed by noise. The attacker might instead increase add noise in the same band in the hope of drowning the spectral pattern. Removing the standard notches, whilst clearly possible, is practically difficult and still unlikely to disturb the correlation greatly. Alternatively they could create a jamming signal to mask the band. This makes a rather easy task for the detector to spot, however, as loud, consistent, broadband noise covering 26MHz is a rare thing. Not only that, but the amenability of power wiring to noise ingress means the attacker would also be jamming themselves.

Detection of power line networks using radiated emissions should naturally be compared to detection using conducted emissions, i.e., a device attached directly to the power network to watch for traffic. Without the need for emissions to radiate and then be received, this approach could reasonably be expected to exhibit

better sensitivity than the ones presented herein. However there are practical difficulties with deploying such a system. Firstly, we have seen that PLC connectivity, while sometimes very far-reaching, can also be severely limited if noise conditions are poor. Furthermore, some electrical devices effect strong attenuation upon a conducted signal (e.g., RCD devices, or transformers in distribution boards as seen in Section 2). A detector that relies upon conducted emissions must be very carefully placed to avoid it monitoring only a small segment of the network. By making use of radiated emissions, EMPower is able to detect networks across any such isolating devices. The second reason is one of practicality; our detection approaches can be implemented with COTS hardware and appropriate software. Nothing need be attached to the power network and no HomePlug compatibility need be developed.

A notable limitation of EMPower is that it does not differentiate between individual networks. In large buildings this is unlikely to be a problem, although users near legitimate networks must account for it, while users operating their own network cannot use the system in this form to detect a new, malicious one.

We focused on HomePlug AV throughout this paper; a standard introduced in 2005 and since superseded. A huge number of contemporary devices still implement this standard, but HomePlug AV2 and G.hn compliant adaptors are also widespread. Both newer standards have enormous PHY-layer similarity to the HomePlug AV design; including the OFDM signalling, utilised (albeit extended) frequency band, filtering notches and preamble structures. Upon testing a pair of Sumvision SVW1000 adaptors (HomePlug AV2), they showed the same spectral usage as the HomePlug AV units, plus additional emissions and notches at various points all the way to the maximum 86MHz limit. As such we are confident that these devices are also detectable by our frequency-domain method. With modifications to accommodate clock rate and preamble changes, our time-domain method could work also. We believe that devices implementing the G.hn standard, again with manifold similarities, will also be detectable, however no devices are currently available in our region so we were unable to test in practice.

10 Conclusion

We have shown how an attacker can easily make use of power networks to establish an unmonitored eavesdropping or bulk data exfiltration capability, or a platform for further network attacks. We have demonstrated detectable EM emissions in real-world settings and argued for the use of these emissions in detecting maliciously-deployed networks. We have introduced frequency- and time-domain detection methods and shown that these can identify the presence of a network with near-perfect accuracy within the same room and still 74.6% accuracy two rooms away and on a different floor. Through an evaluation in a real office environment the methods have been shown to detect an attacker at a maximum distance of 12.9m and in locations from which conducted detection would not be possible.

Acknowledgements

Richard Baker is supported by EPSRC UK as part of the Centre for Doctoral Training in Cybersecurity at the University of Oxford. The authors would also like to thank Mr Michael Webb for his electrical consultancy regarding attacks.

References

1. Carcelle, X.: Power line communications in practice. Artech House (2009)
2. Department of Homeland Security: A guide to securing networks for Wi-Fi (IEEE 802.11 family) (2017). URL https://www.us-cert.gov/sites/default/files/publications/A_Guide_to_Securing_Networks_for_Wi-Fi.pdf
3. devolo AG: devolo dLAN 200 AVminiPCI Datasheet (2009). URL https://www.devolo.com/products/Integrationmodules/dLAN-200-AVminiPCI/data/Product-sheet-dLAN-200-AVminiPCI-com.pdf
4. Dudek, S.: HomePlugAV PLC: practical attacks and backdooring. In: NoSuchCon (2015)
5. HomePlug Powerline Alliance: HomePlug AV specification version 1.1 (2007)
6. HomePlug Powerline Alliance: HomePlug powerline networking technology hits maturation as global broadband standard (2016). URL http://www.homeplug.org/news/member-pr/398/
7. Latchman, H.A., Katar, S., Yonge, L., Gavette, S.: HomePlug AV and IEEE 1901: a handbook for PLC designers and users. John Wiley & Sons (2013)
8. Michael Himmels: Devolo real world field tests (2011). URL http://www.homeplug.org/media/filer_public/25/4f/254f6adb-096a-4913-842b-91e3775da045/devolo_presentation.pdf
9. Newman, R., Gavette, S., Yonge, L., Anderson, R.: Protecting domestic power-line communications. In: Proceedings of the second symposium on usable privacy and security, pp. 122–132. ACM (2006)
10. Newman, R., Yonge, L., Gavette, S., Anderson, R.: HomePlug AV security mechanisms. In: Power Line Communications and Its Applications, 2007. ISPLC'07. IEEE International Symposium on, pp. 366–371. IEEE (2007)
11. PA Consulting Group: The likelihood and extent of radio frequency interference from in-home PLT devices. Tech. rep., Ofcom (2010)
12. Rahman, M.M., Hong, C.S., Lee, S., Lee, J., Razzaque, M.A., Kim, J.H.: Medium access control for power line communications: an overview of the IEEE 1901 and ITU-T G.hn standards. IEEE Communications Magazine **49**(6) (2011)
13. Tasker, B.: Vulnerability: Infiltrating a network via powerline (HomePlugAV) adapters (2014). URL https://www.bentasker.co.uk/documentation/security/282-infiltrating-a-network-via-powerline-homeplugav-adapters
14. TP-Link Technologies Co.: AV200 Wireless N Powerline (2011). URL http://static.tp-link.com/resources/document/TL-WPA281_V1_Datasheet.zip
15. TP-Link Technologies Co.: AV500 Gigabit Powerline Adapter TL-PA511 (2011). URL http://static.tp-link.com/resources/document/TL-PA511.zip
16. Wang, C., Zheng, X., Chen, Y.J., Yang, J.: Locating rogue access point using fine-grained channel information. IEEE Transactions on Mobile Computing **16**(9), 2560–2573 (2017)
17. Zarikoff, B., Malone, D.: Experiments with radiated interference from in-home power line communication networks. In: Communications (ICC), 2012 IEEE International Conference on, pp. 3414–3418. IEEE (2012)