Measurements for Remote Identification of Electrical Equipment

R.O. Lane, T.A. Badran and G.F. Marshall
QinetiQ, Malvern Technology Centre, Great Malvern, Worcestershire WR14 3PS
UNITED KINGDOM

rlane1@QinetiQ.com

This paper investigates the phenomenology of electrical signatures being unintentionally injected onto the mains electricity supply by a range of common household devices. Measurements were made of both current time-series and voltage spectra of a variety of devices, using commercially-available measurement equipment designed for electromagnetic compatibility testing. The measurements give strong evidence that multiple devices can be identified from either their current transient when switched on or off, or their voltage spectrum when running. While some devices are readily identifiable, even in the presence of large amounts of noise, for low power devices it may be necessary to perform filtering of the mains signal to obtain spectra with a suitable signal-to-noise ratio. For certain types of device, spectral signatures of particular instances of the same device appear to be separable in high signal-to-noise ratio environments. For example, it may be possible to identify a specific laptop charger among other chargers of the same make and model. The relevant measurements can be made outside buildings, and when combined with other intelligence, could be used to obtain remote intelligence of building interiors.

INTRODUCTION

Modern electrical equipment, even if tested for electromagnetic compatibility (EMC), still unintentionally injects signals onto the mains electricity supply. An EMC probe typically measures frequencies in the range 9 kHz to 30 MHz. We postulate that information about equipment plugged into the mains can be determined from this signal. Inductive loads produce a noticeable impact on the mains signal; other types of equipment, from individual switched mode power supplies to data centre loads, will also induce a measurable harmonic signal. Particular makes or items of equipment may produce unique detectable signals and useful information about the "pattern of life" of equipment could be deduced from the times-series evolution of the mains current. This would allow information about what is happening inside the building to be determined. This has applications for smart metering or intelligence gathering.

This paper summarizes results of a measurement programme to investigate the phenomenology of electrical signatures. Standard EMC test equipment has been used in a novel manner. Rather than testing whether peak signals of a given frequency are below a specific value, which is the criterion for European certificate (CE) marking, we have measured the spectral content and time domain signature in sufficient detail to determine whether each device has a unique "fingerprint".

EXPERIMENTAL APPROACH

Two types of measurement have been made. Firstly, spectral measurements of the live mains voltage were obtained for different devices. Secondly, current measurements were taken during the process of switching devices on and off, to determine whether this transient signal can be used to identify devices.

To explore the phenomenology of both spectral and time series processes and to identify practical measurement issues, experiments were performed under various conditions. To examine the inherent phenomenology, measurements were taken using a heavily filtered mains input supply, such that any identified characteristics were not due to external conditions. The measurements were made in a shielded chamber with 100 dB of filtering; an additional 40 dB of filtering was used in the spectral analyzer.

In addition to this, measurements were taken with a consumer unit typically used with building electrical supplies and 15 m long length of mains cord, to better simulate a real world measurement scenario. Finally, measurements were also taken with an unfiltered electricity supply, to simulate a non-invasive solution that doesn't modify a building's mains connection.

Measurements were taken for the following devices: filament, halogen, fluorescent, and light emitting diode (LED) bulbs; desk fan; power drill; Nokia and Samsung mobile phone chargers; two different models of Dell laptop charger (including two instances of each model); cathode ray tube (CRT) television; and flat-screen LED television. These devices were selected to cover a broad range of characteristics and have significant variance in the amount of power consumed. Details about the make and model of each device are given in Table 1.

Table 1: Electrical equipment tested.

Item	Make and Model
Incandescent lamp	Status, B11G screw, 15 W
Fluorescent lamp	Megaman, B11G screw, 7 W, BR0107 11W51 GBC01
Desktop fan	Munro, 45 W, PF18, Setting 3
Power drill	Black & Decker, H720H H-15
Mobile phone chargers, loaded by a phone	Samsung SGH-E370 phone (Samsung TAD137UBE charger/ Nokia ACP-8X charger), Nokia 6301 phone (Nokia AC-11X charger)
Laptop chargers, loaded by a laptop	Dell Latitude C640 PP01L Laptop, (2x Dell PA-1900-05D chargers: RevL02/RevL04, 2x Dell AA20031 chargers: G20208/N15773)
Television (cathode ray tube)	Matsui, 14V1R
Television (flat screen)	Technika, LED 22-248COM
LED lamp (230 V)	Diall, GU10 fitting, 6.5 W, GU10HV-4H-WH3
Halogen lamp	Fuxing, GU10 fitting, 50 W, DCL-13W08

A Rohde & Schwarz ESR7 EMI Test Receiver and controller PC were used to measure the voltage supply spectrum while devices were running. The recorded spectra cover the range from 9 kHz to 30 MHz in two bands below and above 150 kHz, with 15 Hz and 2.985 kHz bandwidths respectively. Spectra for each band were computed based on a fast Fourier transform of 10 s of measured data. Typically this device is used to measure peak voltages when testing for compliance with regulations but it can also produce average spectra. It is presumed that average data hold more useful information, so both maximum and average values were recorded. External noise was filtered out using a line impendence stabilization network (LISN). A L289 R&S ESH3-Z5 device was used to perform this function and to provide voltage probes for connection to the spectral analyzer.

Multiple spectra were recorded on different days for each device to determine consistency across time. Ambient measurements were also taken without any devices attached to the system, so that a baseline spectrum could be determined to enable identification of spectral features relative to the background noise. For each run of experimental tests, an ambient signal was measured before and after the run. This process was repeated until enough data was collected to make firm conclusions about the data.

Root mean-square (RMS) live current transient signals were recorded with the Newtons 4th Power Analyzer PPA5531 and a controller laptop. The transient is measured with a time resolution of 20 ms, which is one period of the 50 Hz supply cycle. For devices that have a standby state, transients were recorded for the power cycle: standby, active, standby. For all other devices, data was recorded for the cycle: power off, power on, power off. Each data set was recorded so that the transient cycle was repeated ten times, with a minimum ten second delay between each phase of the cycle. For the flat-screen television, a much longer transient was recorded due to it having a multi stage internal process when activating and deactivating the device. The current signals were recorded without the use of the LISN, with a consumer unit and with a 15 m extension lead. These measurements are representative of what could be obtained in the field with a passive measurement device.

SPECTRUM MEASUREMENTS

Initial experiments compared maximum and average value spectra computed over the measurement window. Average spectra were found to be more consistent and these are presented here. During the first day of testing, measurements for the filament and halogen bulbs and the desk fan were found to be indistinguishable from the ambient spectrum. This is to be expected as these devices are purely resistive or inductive loads and further measurements of these devices have not been made.

Device voltage spectra measured using a LISN filtered mains supply are shown in Figure 1. Both the device spectrum and an ambient spectrum are shown in each subplot. The device spectra are significantly visually distinct from each other, with identifiable features across the whole frequency range from 9 kHz to 3 MHz. Spectra above 3 MHz are not plotted because at these frequencies changing the length of mains cable caused changes in the spectra. Since a practical system must be stable with respect to the length of wiring within a building these frequencies cannot be used. At frequencies below 3 MHz, each device spectrum was relatively stable. The step change visible at 150 kHz is an artefact of the spectrum analyzer, which recorded the signal in two phases with a bandwidth change at the boundary. This is not a feature of the devices themselves. There was a high correlation between measurements made on the live and neutral wires. Differences between these measurements are not critical and, with limited time at the test facility available, it is more important to obtain a large number of live measurements to obtain reliable statistics than compare live and neutral signals. Therefore the conclusions in this paper are based on live measurements only. The presence or absence of a consumer unit in the experimental setup made no discernible impact on the measured signals.

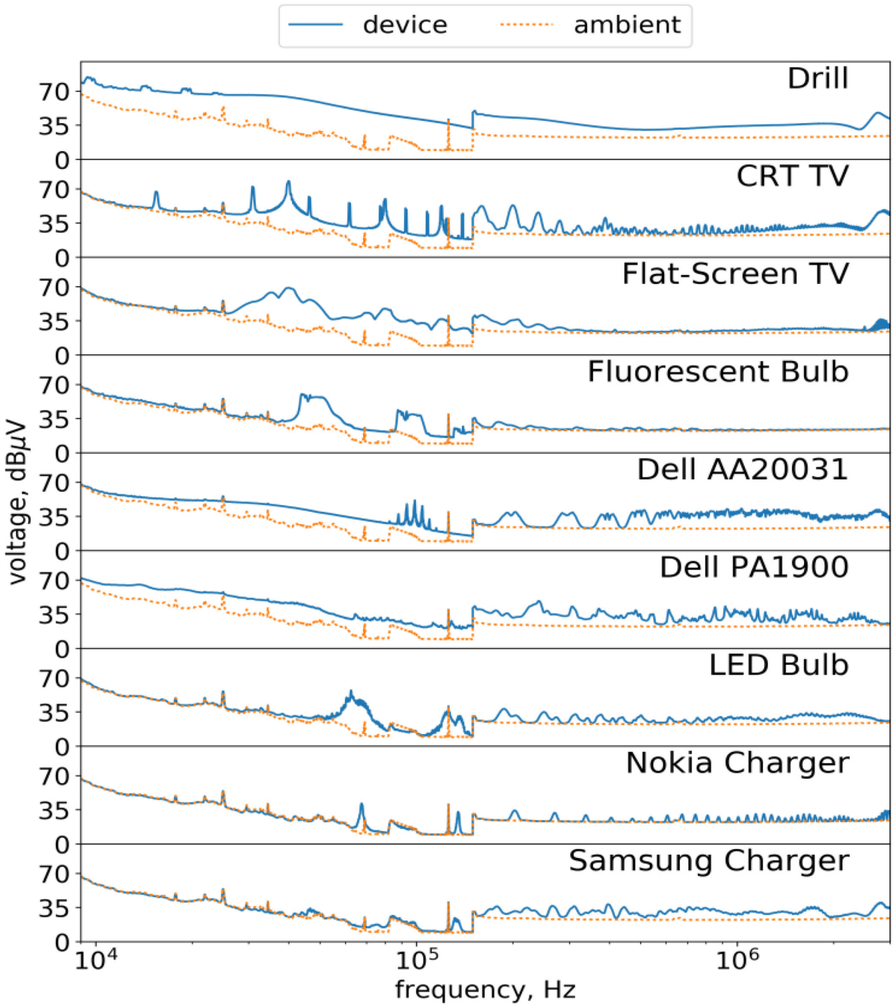

Figure 1: Typical device voltage spectra measured in isolation using a filtered mains supply.

Both device and ambient spectra were consistent across different measurement runs. For most frequencies the variation was almost always less than 5 dB and usually less than 1 dB. The only exceptions were the flat screen television, which showed a slightly higher variation, and the drill spectrum in the upper band, which varied according to pressure on the manually operated trigger. Devices with "on" and "standby" modes had similar spectra in each of these modes, with the "on" mode having additional peaks due to additional circuitry used.

As well as performing analysis under strict laboratory conditions, measurements were made using an unfiltered mains supply. At the time of measurement, the supply simultaneously fed multiple labs running computers and other electronic equipment, as well as control systems for chambers at the test facility and kitchen and bathroom utility rooms. The number and variation of devices in this uncontrolled environment is considerably higher than in a typical residential building. Much of the ambient unfiltered signal is associated with various UK radio broadcast frequencies. There is, therefore, potential to remove these signals using signal processing based on known signal characteristics. However, this has not been performed for the present analysis.

The higher-power drill and CRT devices had spectral components at various frequencies exceeding the unfiltered background signal, meaning these devices should readily be detectable. The laptop chargers, Samsung mobile phone charger, and LED bulb had some above-background power above 1 MHz, with potential for detection. Other devices had no significant above-background energy at any measured frequency. These devices would be difficult to detect based on power spectra alone. However, these observations are based on a particularly challenging scenario, and it may be possible to detect more devices in a less noisy environment.

The above measurements demonstrate that most devices in isolation are identifiable from their spectral signature. However, in practice a multitude of devices are simultaneously powered by the same supply. To explore how spectra combine, measurements were taken with pairs of device connected to a filtered supply. A selection of combinations was chosen to cover the full range of power use, from the highest power device (drill) to the lowest power (Samsung and Nokia phone chargers). Analysis of the data showed that the signals could be approximated by an additive linear system. Although this result was expected from circuit theory, it was useful to test that no non-linear components were introduced by the measurement setup. Since power measurements are recorded on the non-linear logarithmic decibel scale and phase information was not recorded, the spectra in Figure 1 cannot simply be added. However, adding individual device spectra in quadrature in practice produces results similar to combined spectra. For most device combinations at least one distinct frequency peak per device is visible in the combined spectrum. However, when the high power drill is in operation, its spectrum typically dominates those of the other devices.

A set of nominally identical devices would be expected to have small variations in their spectra due to manufacturing differences. For each model of Dell laptop charger, two specific instances of the same charger with different serial numbers were available for testing. Figure 2 shows that the overall spectral shape and location of peaks are very similar for the two instances. For the AA20031 the shapes of the spectra are nearly identical, with small shifts in frequency and power. For the PA1900, spectra share similar peak structures above 150 kHz, but there is more variation in the low frequency range. More measurements are required to determine the extent to which these differences vary with external factors such as temperature, device age, and loading conditions. While it would not be expected that an individual device could be identified from the pool of all manufactured instances, information in addition to the spectrum could be used to provide a more robust identification.

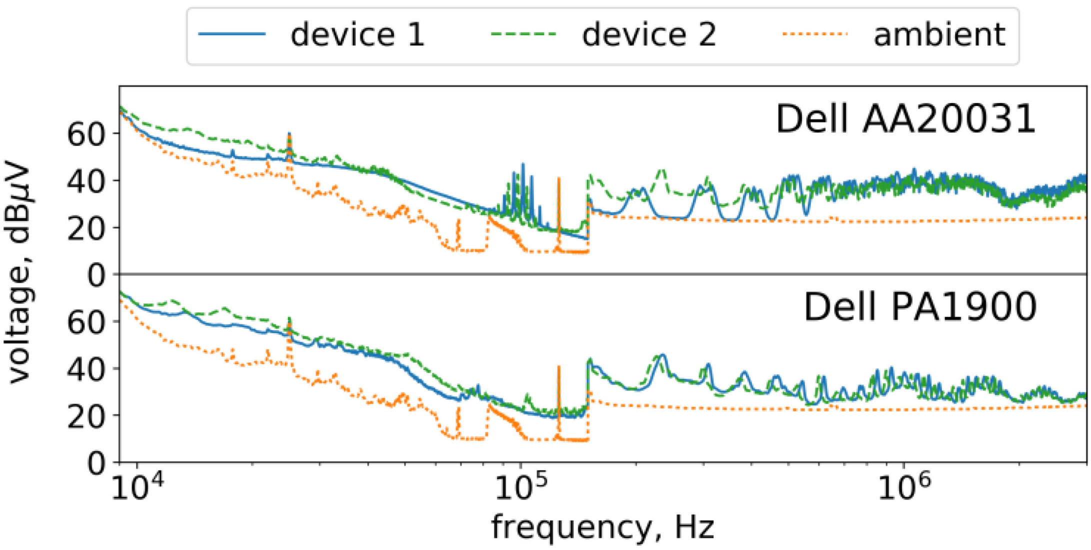

Figure 2: Typical spectra for nominally identical devices.

TIME SERIES MEASUREMENTS

Typical current signals when devices are switched on and then off are shown in Figure 3. The devices follow a pattern of a large in-rush current, followed by a steady state. For some devices, namely the bulbs and Samsung mobile phone charger, the signal does not exhibit other characteristics. In these cases the only differentiating factor between the transitions is the steady state current. However, other devices have more structure to their transient signals. The laptop chargers and the phone chargers have a small region between the in-rush and steady-state period where the current drops back to near the background level. The drill and desk fan both slowly transition to the steady state. Finally, the flat-screen and CRT televisions have more complex, multi-stage transients. In the case of the flat-screen TV there is also a characteristic signal during the powered-to-standby transition. The PA1900 model of Dell laptop charger has a very noisy steady state current compared to other devices.

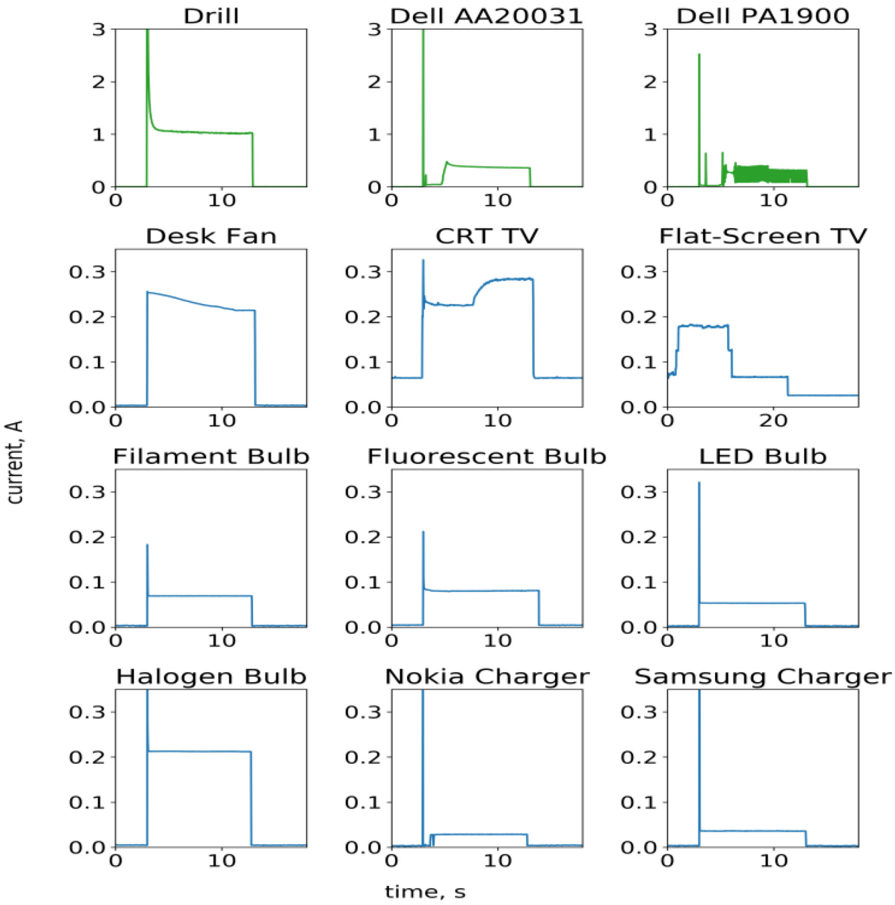

Figure 3: Transient current signals when devices are switched on then off.

Automated classifiers require some consistency between different measurement runs of the same device. The steady state current of the bulbs and Samsung charger are very consistent over time, although there is more variation in the measured peak value of the in-rush current. The drill, desk fan, and flat-screen television transient signals are relatively constant across different measurement runs. Figure 4 shows transient signals for the Dell AA20031 laptop charger. Timings for this signal are fairly constant but for reach run measured there was a decrease in the steady state current from 400 mA in the first run to 300 mA in the tenth run. Two of the later runs saw the current drop to 50 mA for a period of 1 to 2 seconds before returning to normal. This is likely due to different charge states as the battery is charged up.

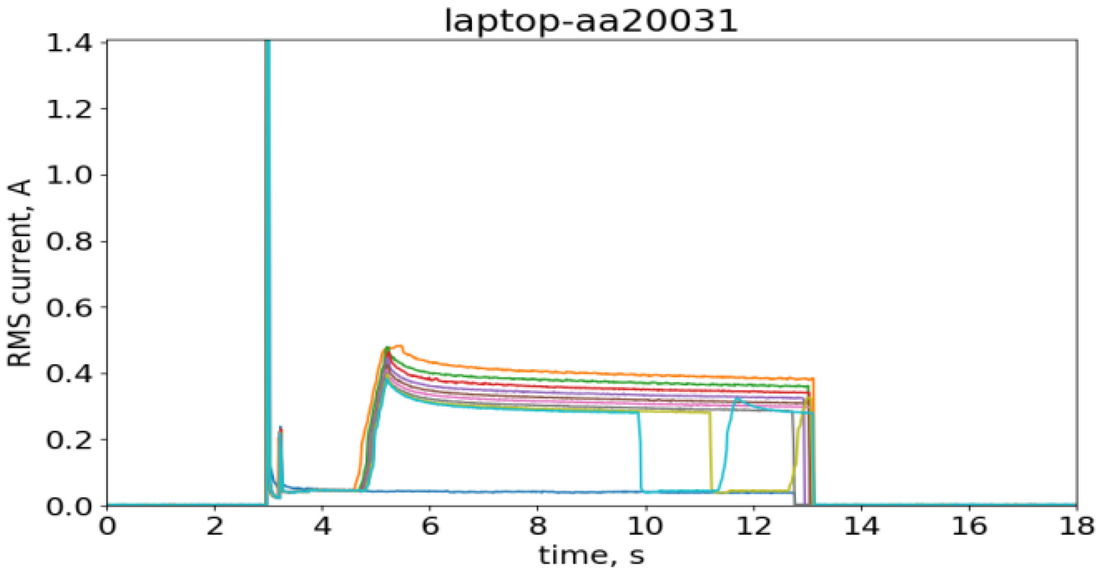

Figure 4: Transient current signals for the Dell laptop charger (AA20031).

Figure 5 and Figure 6 show transient signals for the Nokia phone charger and CRT television. Both devices have in-rush currents followed by a quiet period and then a higher steady state current. The quiet period duration varied between runs from 1 to 6 s for the charger and 4 to 6 s for the television. Figure 7 shows measurements for the PA1900 Dell laptop charger. The duration of quiet period is relatively constant with a spike about 0.7 s after switch on. However, the exact timing of the rapid peaks that start about 2 s after being switched on varies between runs.

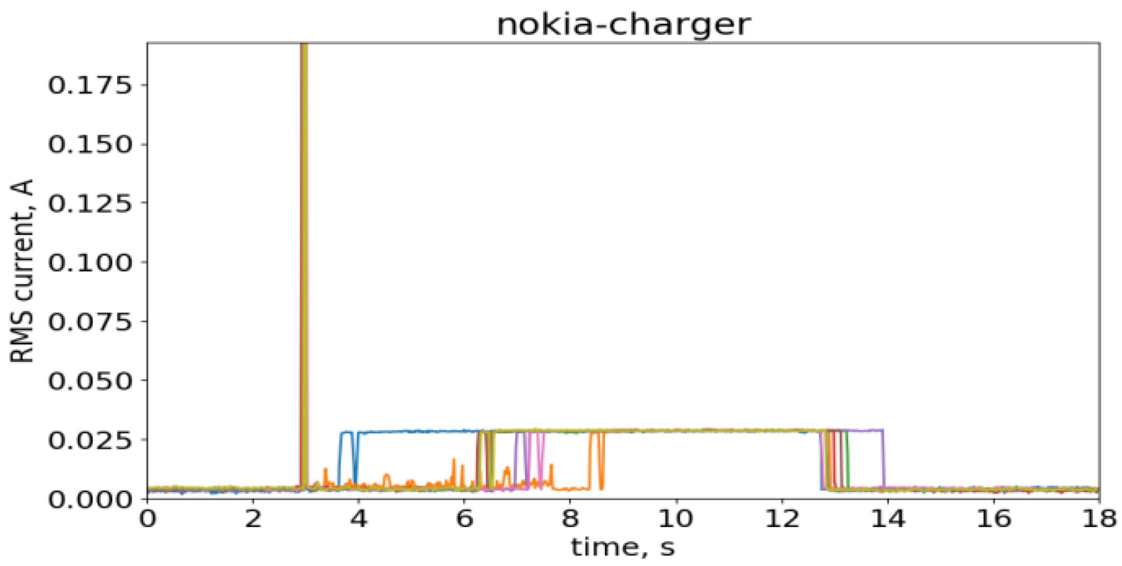

Figure 5: Transient current signals for the Nokia mobile phone charger.

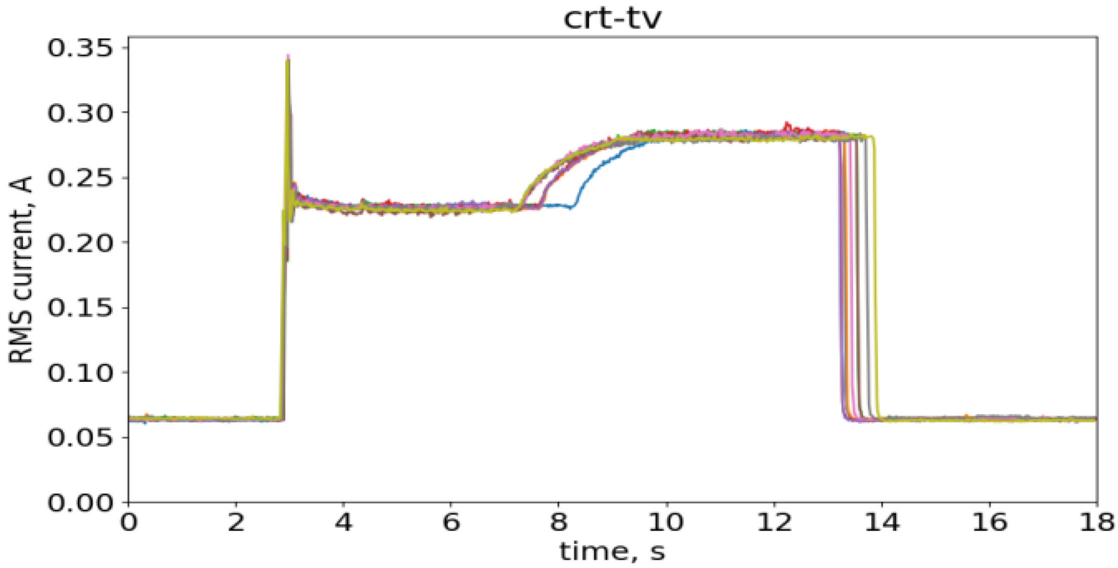

Figure 6: Transient current signals for the CRT television.

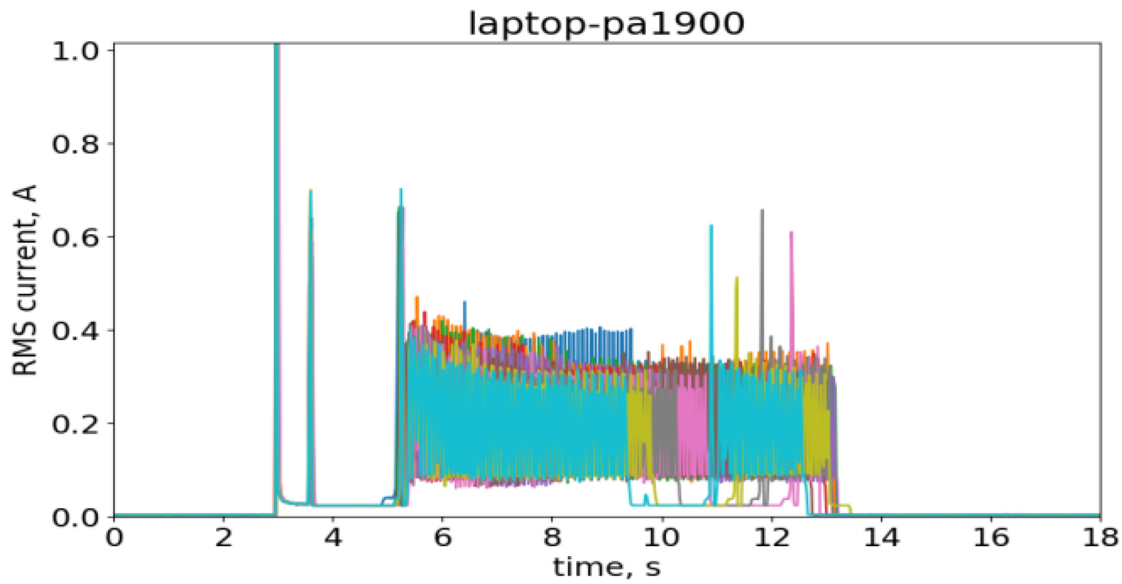

Figure 7: Transient current signals for the Dell laptop charger (PA1900).

MACHINE LEARNING

Supervised learning techniques build a model of labelled data and perform classification predictions for new signals. Typically, multi-dimensional feature vectors are used to represent data. These features can either be the raw data, which would consist of the power in each frequency bin for the spectral measurements, or some other feature constructed from the data. Extracted features could include the mean power, number of peaks, total energy in certain frequency bands, or the distance between certain peaks. Manually defined features have the advantage of using expert knowledge of which aspects of the data provide a discrimination capability but it can be time consuming to determine such features. Raw features can immediately be used as an input to the classifier. However, certain devices, such as the laptop chargers shown in Figure 2 exhibit characteristics based on relative rather than absolute frequencies. Therefore some work needs to be done in defining relevant features. A general detailed discussion about various classifiers with their advantages and disadvantages is given in [1]. A survey focused on specifically on appliance load monitoring is given in [2]. For both spectrum and time series classifiers it is likely that more measurements will be required to train a classifier.

Spectrum classification allows the detection of some steady-state signals. However, for some devices these signals are too similar to each other and the unfiltered ambient noise may be too high at frequencies of interest for classification to take place. The alternative is to analyze the current time series of devices. In addition to detecting the presence of a device, this enables the possibility of determining the precise time at which it is switched on or off. A variety of time series classification techniques exist; a review of these is given in [3]. There are two major approaches. The first is to extract fixed-dimension features from the time series and then use standard machine learning algorithms. The second approach is to use some property of the data arising from the fact that it is a time series. Candidates for this type of algorithm are: template matching, specialized neural nets, and dynamic time warping.

Template classifiers define a windowed template for each class constructed by averaging over aligned training data segments extracted from a long time series. In testing, the cross-correlation between the test signal and each template is calculated. The template with the highest correlation is used to label the data. This procedure is invariant to translations (time delays) of the signal. Implicit assumptions used in this classifier are: raw data have independent identically distributed Gaussian statistics; and signals do not undergo deformations in time other than delays. The identical distribution assumption does not apply to most devices as the magnitude of the measured in-rush current varies significantly more than that of the steady-state current. Also, device signals undergo temporal distortions that are not simple delays. This reduces the practicality of this classifier.

Neural nets are arranged in layers with each layer providing a non-linear map from its input to output, enabling a non-linear decision boundary in measurement space. Invariance with respect to signal translation can be incorporated in convolutional neural nets (CNNs) by imposing a certain structure on the net and using shared weights. When applied to time series they are sometimes referred to as time delay neural nets [4]. A convolutional layer in a CNN learns several “feature maps”, which are the convolution of the layer below with a kernel learnt during training. The network normally includes a pooling layer, which allows for minor temporal distortions in the signal. CNNs have been successful in image recognition, but their application to time series classification has received less attention. A recurrent or recursive neural net has feedback between layers and is typically designed to learn sequential or time-varying patterns. Due to feedback loops, the back-propagation algorithm used to train standard feedforward nets cannot be used and more complex algorithms are required [5].

Dynamic time warping (DTW) is a technique for finding the optimum temporal warping function between two signals with non-linear distortions in time [6]. A distance measure based on the warping function and cost metric is then used as part of a nearest neighbour classifier algorithm. Since DTW allows any monotonic warping function, the standard algorithm may produce unrealistic distortions. These can be limited by setting a window size and only calculating warps within this window, which also speeds up the algorithm. The cost matrix is based on insertion, deletion, or retention of samples as two signals are processed from beginning to end. An example cost matrix for two CRT signals from Figure 6 is shown in Figure 8, along with the optimum warping function. There are four regions of low cost from approximately sample 1 to 50, 51 to 300, 400 to 550 and 600 to 650. These regions are associated with parts of the signal where time distortions have little effect as the signals are similar in amplitude: before switch on, waiting for warm-up, after warm-up, and after switch off. The large increase in cost when one of the signals is held at a sample from 1 to 50 and the other signal is progressed beyond sample 50 is due to the sharp increase in current at sample 50 when the device is switched on. The algorithm has determined that unless both signals undergo this change at similar times then the signals must be quite different. Other regions of slightly increased cost are due to mismatches in the start of the warm-up period and switch-off time. In practice, the time series difference should be used rather than the direct time series to ensure that superimposed signals are correctly dealt with.

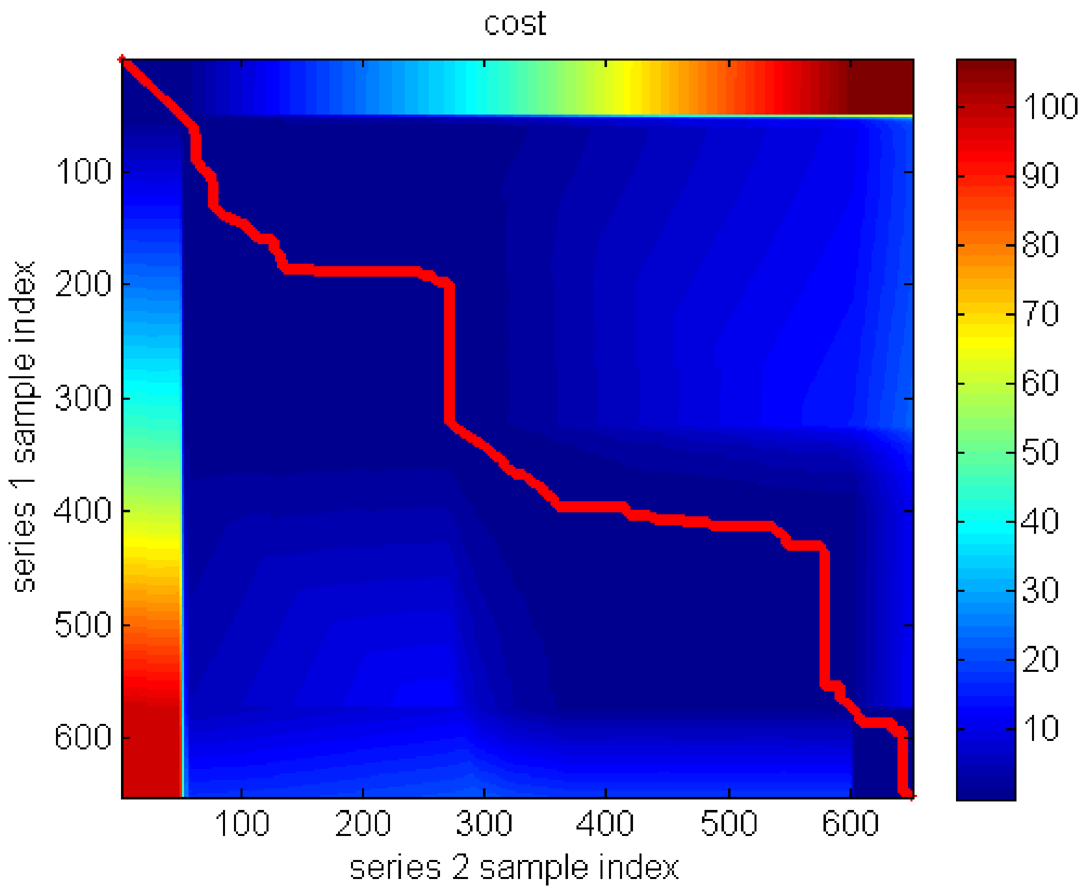

Figure 8: Cost matrix and optimal warping function between two signals.

RELATED WORK

There have been several other efforts to detect operation of individual devices based on the energy supply. These are aimed at cooperative smart meter applications for reducing energy usage in the home. The techniques are referred to energy disaggregation, disambiguation, non-invasive load monitoring, or cognitive metering.

A comprehensive review of the benefits, data, algorithms, hardware, and applications of electricity disaggregation is given in [7]. This type of technique is also possible using gas and water supplies. Most developed countries are planning the introduction of smart meters, which could be modified to have a device detection capability. However the sampling rate of most meters is not likely to be high enough to detect all devices of interest and additional sensors are required. An example bespoke sensing system and associated algorithm to detect and classify electrical devices in the home is given in [8]. Signatures were found to be repeatable and stable over a period of six months.

One example of a system that tries to prevent remote detection of building occupancy by modulating a water heater power supply is given in [9]. However, that system does not mask signals that can be detected using sub-minute measurements.

CONCLUSION

A range of common household electrical devices can be identified from the voltage power density spectrum they superimpose onto the mains. While some devices are readily identifiable when passively measuring these signals, for low power devices it may be necessary to perform filtering of the mains input to obtain suitable spectra. Based on a small sample of devices to date, spectral signatures of particular instances of the same device appear to be separable, such that a laptop charger may potentially be identifiable with access to corroborating information. By looking at current transients, it is also possible to determine when devices are switched on or off, and more complex devices can be identified by the shape of their transient signal. Combined with other intelligence, analysis such as this could provide useful information on the activities happening within a building.

ACKNOWLEDGMENTS

This work was funded by the UK Defence Science and Technology Laboratory, contract number DSTLX-1000098818.

REFERENCES

[1] Webb, A.R., Copsey, K.D.: 'Statistical Pattern Recognition' (John Wiley & Sons Ltd, Chichester, UK, 2011, third edition).

[2] Zeifman, M., Roth, K.: 'Nonintrusive appliance load monitoring: review and outlook', *IEEE Trans. On Consumer Electronics*, 2011, **57**, (1), pp. 76-84.

[3] Wang, X., Mueen, A., Ding, H., Trajcevski, G., Scheuermann, P., Keogh, E.: 'Experimental comparison of representation methods and distance measures for time series data', *Data Mining and Knowledge Discovery*, 2013, **26**, (2), pp. 275-309.

[4] Bishop, C.M.: 'Neural Networks for Pattern Recognition' (Oxford University Press, 1995).

[5] Medsker, L.R., Jain, L.C.: 'Recurrent Neural Networks' (CRC Press, London, 2001).

[6] Liu, B., Luan, W., Yu, Y.: 'Dynamic time warping based non-intrusive load transient identification', *Applied Energy*, 2017, **195**, pp. 634-645.

[7] Armel, K.C., Gupta, A., Shrimali, G., Albert, A.: 'Is disaggregation the holy grail of energy efficiency? The case of electricity', *Energy Policy*, 2013, **52**, pp. 213-234.

[8] Gupta, S., Reynolds, M.S., Patel, S.N.: 'ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home', Proc. 12th ACM Int. Conf. on Ubiquitous Computing, Copenhagen, Denmark, September 2010, pp. 139-148.

[9] Chen, D., Kalra, S., Irwin, D., Shenoy, P., Albrecht, J.: 'Preventing Occupancy Detection From Smart Meters', *IEEE Trans. On Smart Grid*, 2015, **6**, (5), pp. 2426-2434.