

Exploiting the Physical in Cyber-Physical Systems

On the Practical Use of Physical-Layer Information to Enhance
System Security



Richard Baker
Worcester College
University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

Acknowledgements

Institutional

To the EPSRC for funding the full costs of the DPhil programme, along with a stipend to cover my living expenses.

To the CDT in Cyber Security for finally crystallising how and where I wanted to do my doctorate — and for living up to my expectations.

Abandoning a career that did not fulfil me, and doing something that fascinates me instead, was one of the best decisions I have made.

Personal

To Ivan Martinovic, for supervising my DPhil and welcoming me to the systems security community. I have been very lucky to have a mentor who cared deeply about my interests and acted to protect them. I am delighted that we have grown to be friends as well.

To Simon Birnbach and Riccardo Spolaor for being inspiring researchers and good friends.

To Alastair Janse van Rensburg, Andrew Dwyer, Chris Vaas, Kris Wilson, Matt Smith, Louise Axon, Mered Williams, Ilias Giechaskiel and Jan Siloman; my CDT14 comrades. I do not have space to describe the fine qualities of each, but I know.

To Maureen York, who knew when to relax and when to fight — and was always in our corner.

To David Hobbs who, the very day I wrote these acknowledgements, was still dropping everything to rescue me.

To Kasper Rasmussen and Andrew Markham, who advised and assessed this body of work as it developed.

To Andrew Martin and Flavio Garcia, for taking the time to examine this thesis, putting me at ease throughout the viva and providing insightful feedback to improve it. Additionally, to Andrew Martin for all his support and guidance throughout five years of the DPhil.

To Worcester College and the MCR. I've had the privilege to join its ranks and involve myself in its affairs. It has welcomed me since the moment I set foot there and, after five years, I still don't quite believe it.

To those dearest friends I met at Worcester. From those I met first: Dan Hall, Jack Boteler, Alex Morrice, Lisa Bernhardt, Belinda Lo and Michael Dangerfield; to those at the end: Justin Weeks, Kim Füllenbach, Leila Tai, Célia Souque, Susie Wright and Adam Cobb (the AGDC).

To my loved ones, family and old friends; I am blessed with too many to list entire.

To my parents, Jan and Chris, who did not at first understand why I wanted to do this, but would have moved the heavens to help me.

To Claire Parker, I met you here and may we never be apart.

Abstract

This thesis argues that cyber-physical systems are, by their very nature, at risk from physical-layer attacks as well as cyber attacks. The proliferation of cheap and easy-to-use sensing and actuation technologies has drastically lowered the bar for attackers to conduct physical-layer attacks, even with only limited resources. As our reliance upon cyber-physical systems grows, so too does the impact of attacks.

It is argued that the same easy accessibility of technology that equips attackers, also enables the use of physical-layer security techniques in developing defences. A series of work is presented, exploring the practical use of physical phenomena to secure real-world cyber-physical systems.

Timing constraints are used for the verification of aircraft location claims, to inhibit spoofing. This demonstrates a straightforward application of physical-layer techniques, enhanced with mobility, to drastically limit an attacker's capabilities.

Wireless propagation measurements are used to determine the presence of a drone and track it during a privacy-invasion attack; where traffic itself does not provide sufficient insight. The successful results highlight the potential for using even simple, ubiquitous metrics to gain detailed insight into the physical world.

Leaked electromagnetic signals are then used to detect a class of malicious network; exploiting the wireless propagation mode to achieve better performance and more convenient deployment characteristics than are possible with the original signal. This demonstrates the scope for incorporating unconventional physical effects to improve a security design.

The combined results are drawn on to argue that the use of physical-layer features is *practical* in real systems, even those that were not originally designed with due consideration for their tacit physical dependencies.

An eavesdropping attack is also presented against a state-of-the-art electric-vehicle charging system. This attack builds upon the electromagnetic leakage used defensively earlier, which is exacerbated by design choices made in the charging system. The eavesdropping attack is shown to be widely effective against real deployments, with results that suggest various active attacks would also be effective.

Observations from the attack are used to argue that as well as being practical, it is also *necessary* to incorporate physical-layer features in security design, as even emerging modern systems with detailed security models are vulnerable to critical physical-layer attacks.

Contents

List of Figures	ix
List of Abbreviations	xiv
Glossary of Radio and Digital Signal Processing Terms	xvi
1 Introduction	1
1.1 Contributions	5
1.2 Thesis Organisation	6
1.3 Publications	8
2 Secure Location Verification	10
2.1 Introduction	11
2.2 Background	12
2.2.1 Cooperative Surveillance	12
2.2.2 Secure Location Verification	13
2.2.3 Unmanned Aerial Vehicles	15
2.2.4 Air Traffic Management	16
2.3 Related Work	17
2.4 Attacker Model	18
2.5 System Model	19
2.6 Effect of Errors	22
2.7 Security Analysis	27
2.8 Practical Considerations	31
2.9 Prototypes	32
2.9.1 Ground-based Prototype	33
2.9.2 Aerial Prototype	35
2.10 Evaluation	35
2.10.1 Simulation	35
2.10.2 Real-World Data Collection 1 (Ground)	41
2.10.3 Real-World Data Collection 2 (Airborne)	44
2.10.4 Verification of Real Tracks	47
2.11 Discussion	47

2.11.1	Alternative Configurations	48
2.11.2	Potential Applications	49
2.12	Further Work	51
2.13	Conclusion	53
3	Drone Detection	54
3.1	Introduction	56
3.2	Motivation	56
3.3	Contribution	58
3.4	Related work	58
3.5	Background	61
3.5.1	Drones	61
3.5.2	UK Drone Flight Regulations	62
3.5.3	Received Signal Strength Indicator (RSSI)	62
3.6	Attack model	64
3.7	System model	66
3.8	Detection	67
3.8.1	Hallmarks of Drone Activity	68
3.8.2	Statistical Metrics	68
3.8.3	Drone Attack FSM	80
3.9	Evaluation	82
3.9.1	Experimental Design	83
3.9.2	Real-World Tests	85
3.9.3	Results	90
3.10	Discussion	101
3.10.1	Results	101
3.10.2	Observations	103
3.11	Future Work	104
3.12	Conclusion	105
4	Malicious PLC Network Detection	106
4.1	Introduction	106
4.2	Motivation	108
4.3	Security Risk	109
4.4	Threat Model	112
4.5	Related Work	113
4.6	Background	114
4.6.1	HomePlug AV	116
4.6.2	EM Leakage	117
4.7	Designing EMPower	119

4.7.1	Detecting Radiated Emissions	120
4.8	Detector Design	123
4.8.1	Frequency Domain	124
4.8.2	Time Domain	125
4.9	Evaluation	126
4.9.1	Experimental Setup	126
4.9.2	Detection Accuracy	128
4.10	Discussion	129
4.11	Conclusion	134
5	Electric-Vehicle Charging	136
5.1	Introduction	137
5.2	Motivation	138
5.3	Background	140
5.3.1	Combined Charging System (CCS)	143
5.3.2	CCS Security	144
5.4	Related Work	145
5.5	A Near-Ideal Side-Channel	147
5.6	Threat Model	151
5.7	PLC Eavesdropping Tool	151
5.8	Real-World Measurement Campaign	155
5.9	Results	157
5.9.1	Eavesdropped Communications	157
5.9.2	Effects of Location	160
5.9.3	Message Recovery	160
5.10	Security Analysis	162
5.10.1	Unencrypted Communications	162
5.10.2	Private Data	163
5.10.3	Charging Attacks	164
5.11	Lessons Learnt	166
5.11.1	Wireless Threats	166
5.11.2	Reliance on a Non-Existent PKI	166
5.11.3	Available PHY Security Disabled	167
5.12	Potential for Active Attacks	168
5.12.1	Jamming in detail	169
5.12.2	Message injection in detail	171
5.12.3	Relay in detail	171
5.13	Countermeasures	173
5.13.1	Protocol Changes	173

5.13.2	Equipment Changes	174
5.13.3	Physical-Layer Security	175
5.14	Conclusion	175
6	Conclusion	177
6.1	Collected Results	177
6.2	Conclusions	178
6.3	Final Remarks	180
Appendices		
A	HomePlug GreenPHY Receiver	182
	References	187

List of Figures

2.1	Structure of verification system	20
2.2	A two-node TDoA system in two dimensions at three successive points in time (as denoted by subscript).	22
2.3	Diagram of a false location claim being detected. The claimed position is indicated by a red +, but is broadcast by an attacker positioned at the black ×. The two receivers are each indicated by a black • and the hyperbola constructed from the TDoA measurement is in gray. The solid line indicates the branch of the hyperbola that the transmitter must be located on, based on which receiver detected the message first. A claimed position anywhere on the same branch would appear genuine, but not a location off the hyperbola, as attempted here.	23
2.4	Diagram of the diminishing opportunities for an attacker as multiple messages are received. As in Figure 2.3, the attacker is located at the black × and each receiver is a •, but the second branch of each hyperbola has been omitted. For a single message they can claim any position on the hyperbola described by the TDoA measurement. However, here the resulting hyperbolas for two messages can be seen and are substantially different. The only position that lies on both is the attacker’s true location.	24
2.5	Representation of one sheet of the hyperboloid of possible locations for a 3-dimensional TDoA measurement.	25
2.6	A two-node TDoA system in two dimensions at a single moment in time, with the errors for transmission time measurement and receiver localisation visualised.	26
2.7	Effect of TDoA timing error at different offset angles from receivers. A target at uniform distance is shown in each subfigure, but at different angles. The TDoA measurement is corrupted by timing error up to the same maximum in each case. The effect of timing error is most pronounced when the target location sits along the same line as the receivers. It progressively diminishes as the location approaches the perpendicular. In every case it increases with distance.	26

2.8	Effect of TDoA positioning error at different offset angles from receivers. A target at three positions is shown across the subfigure, but at different angles. The TDoA is corrupted by positioning error moving the second focus, up to the same maximum in each case. As with timing error, the effect of positioning error is greatest on the line of the receivers and smallest at the perpendicular. It also increases with distance. Unlike timing error, positioning error has different effects on the two arms of one hyperbola branch; with the target's location experiencing the smallest effect from error.	27
2.9	Mobile receiver node architecture diagram	32
2.10	Classification performance of the verification system. Sub-figure 2.10a shows detection rates against the three attacker classes, while Sub-figure 2.10b shows the receiver operating characteristic for the detector.	38
2.11	Effect of changing measurement noise ($\epsilon_{measure}$) and clock drift (t_{drift}) standard deviation on detection rates	39
2.12	Exterior view of vehicle with prototype installed	41
2.13	Route taken	42
2.14	Maximal detection ranges during route	43
2.15	The prototype mobile receiver	44
2.16	The mobile receiver in flight.	46
2.17	Illustration of aircraft location verification based on TDoA results collected using airborne receiver. The black diamonds are the locations of the receivers (overlapping completely at this map scale). Blue dots indicate validated claims, while orange dots denote verification failures.	48
3.1	Examples of effect of distance on surveillance footage quality.	65
3.2	The attacker launches the drone at launch distance d_l and flies it to surveillance distance d_s to carry out the attack. The detector is installed in the window. The FOV γ of the window limits the area that is in LOS of the detector.	65
3.3	Flow diagram of the detection algorithm.	67
3.4	Illustration of the distance change required to vary observed RSSI values by 15dB, for several closest-approach distances. As the closest point of approach increases, the furthest point of approach increases far faster (exponentially, in fact). This is under line-of-sight conditions in free space.	72
3.5	Metrics computed over time for DJI Phantom drone performing privacy-invasion attack.	75
3.6	Metrics computed over time for mobile device being carried around.	76

3.7	Metrics computed over time for static device.	77
3.8	State machine modelling the privacy-invasion attack. Solid lines indicate state transitions due to test of statistical criteria, dashed lines indicate transitions due to timeouts.	80
3.9	Various example approach patterns	84
3.10	The real-world experiment site	85
3.11	Receiver and transmitter placement in building. Walking routes are also shown.	87
3.12	Examples of receiver deployment.	88
3.13	Drones used in the real-world experiments	88
3.14	Different flight patterns during the approach	89
3.15	Plot of detection accuracy against message rate, for both drones. Only detection of the drones themselves is considered here, other traffic is excluded.	92
3.16	Normalised false negative rates for each approach pattern. Only receivers that experienced false negatives are shown.	94
3.17	Examples of detection during various approaches; all as observed by receiver RX2. The 10s-mean RSSI is shown in red. The vertical lines indicate state changes in the detection FSM: Gray = INIT, Cyan = APPROACHING, Blue = APPROACHED, Purple = SURVEILLANCE, Green = ESCAPE).	96
3.18	Detection timings by receiver (timings indicate the first transition to state APPROACHED). Individual detections are illustrated with a black dot, while the mean for each receiver is denoted by a red cross.	99
3.19	Detection performance as window sizes are varied.	101
4.1	A mock-up of a covert traffic capture and bulk data exfiltration attack using a maliciously-installed power line network.	110
4.2	Spectral usage with standard mask defined in HomePlug and IEEE1901 standards.	120
4.3	Observed radiated emissions from power line adaptors.	121
4.4	Surface plots of relative signal power against the ambient background.	121
4.5	Block diagram of the system, showing the full processing chain for frequency- and time-domain analyses.	122
4.6	Floorplan of the target building, showing the public locations (shaded in red) and the private locations (in white). The detector is marked with a D symbol, the attack locations with numbers corresponding to entries in Table 4.3. The markers with dashed lines are on the floor below.	127

4.7	ROC curve for the detector, computed across all test locations and network states.	131
5.1	Overview of EV charging with V2G and payment options shown. Solid, blue lines indicate power flow whilst dashed, black lines indicate communication.	139
5.2	Illustrations of the physical connectors for CCS charging, along with the network stack used for communication.	141
5.3	An overview of vehicle-to-charger network establishment in HomePlug GreenPHY. If the secure mode is supported by both parties and enabled in initialisation then step 3 occurs, allowing the messages in steps 5–7 to be signed and the one in step 7 also encrypted.	146
5.4	Example single-ended signals, with the radiated emissions that result. As the emissions are the gradient of the signal, the square wave produces only impulses while the OFDM waveform is all but unchanged.	148
5.5	A diagram of the CCS communication circuit. The loads on each line connected to the PLC modem are not balanced. Resistors R2 & R3 alter the voltage in the low-level communication, but also vary the imbalance further.	148
5.6	Observations from PLC signal injection in lab conditions. The gray signal is the emitted signal, the black signal is that observed on the PLC line when communication was idle.	150
5.7	Architecture of PLC monitoring tool. The signal is captured and prefiltered, before moving through a software receiver chain to recover messages. The message following behaviour extracts security-relevant data and stores all messages. Charging traffic can be further processed, while traffic using other protocols will need separate onwards processing.	152
5.8	A composite diagram showing the experiment layout. The five antenna locations are denoted with a dashed × symbol.	155
5.9	Observed signal across the HPGP bandwidth, at each antenna location. The HPGP spectral mask is overlaid to indicate the regions in which transmission occurs, although no valid comparison can be made with its power value as the measurement was not calibrated. Signal degradation and noise ingress is visible in every case, although far more prominently in (b) and (c).	158
5.10	Eavesdropping from the next parking bay (site G), more than 4 metres away on the other side to the charging cable. In this arrangement 91.8% of messages were received successfully.	161

5.11	Two vehicles charging simultaneously. With the eavesdropper between the two vehicles 42.5% of messages were received successfully, including the NMK key establishment for both vehicles.	162
5.12	Tree diagram indicating the potential data available under a range of communication scenarios.	163
5.13	Theorised PHY-layer relay attack on ISO 15118 Plug & Charge. . .	170
5.14	The modified SLAC network establishment. Steps 6.1 and 6.2 are new, while step 7 has been modified.	174
A.1	HomePlug AV adaptors communicating across a short wire. Conducted signals and radiated emissions can be seen on the oscilloscope (top in yellow and bottom in blue, respectively).	183

List of Abbreviations

ADC	Analogue-to-Digital Converter.
ADS-B	Automatic Dependent Surveillance – Broadcast.
AES	Advanced Encryption Standard.
CAA	Civil Aviation Authority (UK airspace regulator).
CCS	Combined Charging System.
CFO	Carrier Frequency Offset.
COTS	Commercial off-the-shelf.
CPS	Cyber-Physical System.
CRC	Cyclic Redundancy Check.
EM	Electro-magnetic.
EV	Electric Vehicle.
FAA	Federal Aviation Administration (US airspace regulator).
FPGA	Field-Programmable Gate Array.
FOV	Field-of-view.
FPV	First-person view (in context of UAV flight).
FSM	Finite state machine.
GDOP	Geometric Dilution of Precision.
GNSS	Global Navigation Satellite System.
GPSDO	GPS-Disciplined Oscillator (indeed any GNSS).
HPGP	HomePlug GreenPHY.
JSON	Javascript Object Notation.
LAN	Local-Area Network.
LLR	Log Likelihood Ratio.
LOS	Line-of-sight.
MSPS	Megasamples per second (in context of ADCs and SDRs).

NEK	Network Encryption Key (in content of HPGP).
NLOS	Non-line-of-sight.
NMK	Network Membership Key (in content of HPGP).
NTP	Network Time Protocol.
OFDM	Orthogonal Frequency-Division Multiplexing.
PKI	Public-Key Infrastructure.
PLC	Power Line Communication.
PPDU	PHY-layer Protocol Data Unit.
PPS	Pulse-per-second (in context of GNSS timing).
QPSK	Quadrature Phase-Shift Keying.
RF	Radio-Frequency.
RPAS	Remotely-Piloted Aerial System.
SCO	Sampling Clock Offset.
SDR	Software-Defined Radio.
SNR	Signal-to-Noise Ratio.
TCB	Trusted Computing Base.
TD_oA	Time-Difference of Arrival.
UAV	Unmanned Aerial Vehicle. Alternatively “RPAS” or “drone”.
USRP	Universal Software Radio Peripheral (an SDR brand name).
XML	Extensible Markup Language.

Glossary of Radio and Digital Signal Processing Terms

Throughout this thesis, a number of concepts related to radio communication and digital signal processing appear repeatedly, a brief description of these concepts and their importance is provided below¹.

analogue-to-digital converter (ADC) an electronic device that measures a continuous, analogue signal and digitises it into a series of observations. Specified by sampling resolution (in bits) and sampling rate (in samples-per-second (SPS)).

auto-correlation a form of *cross-correlation* in which a signal is correlated against a delayed version of itself. The auto-correlation is typically used to detect repetition within a signal.

bandwidth (BW) the difference between the highest and lowest frequency components used in a signal, measured in Hertz (Hz). Human speech occupies approximately 8 kHz, a single 802.11n Wi-Fi channel can occupy up to 40 MHz. Bandwidth is directly related to the fastest transition that can occur in a signal and thereby determines the signalling rate; wider bandwidth signals can convey more information. Bandwidth requirements for a signal are identical irrespective of the central frequency of transmission. By the Nyquist theorem [4], accurately digitising a signal requires a sampling rate of at least twice the bandwidth.

carrier frequency offset (CFO) a manifestation of a difference in the clock frequency between a transmitter and receiver, causing signals to be appear mistuned. A counterpart to sample-clock offset. Accurately tuning to a signal requires that a *tuner* is provided with an accurate reference signal, that is then multiplied up to the desired frequency. Any inaccuracy in the reference signal causes an error in the tuning, which manifests as signals appearing to be at the wrong frequency. For example, if communication takes place between A and B at a nominal frequency of 100 MHz, but A's clock is 0.1 MHz slow, while B's clock is 0.1 MHz fast, then

¹A far more comprehensive treatment of the topics is provided, for radio communication in [1] and for digital signal processing in [2] and [3]

transmissions arriving at B will appear to be at 99.8 MHz, while transmissions arriving at A will appear to be at 100.2 MHz.

channel model a mathematical model of the distortion occurring during propagation of a signal through some medium or *radio environment*. While signals can also experience distortion from other factors, the propagation environment typically dominates. A channel model may have an exact mathematical definition (e.g., for use in simulation), but is more often an approximation derived from observation. Such an observed channel model is commonly represented as a transfer function, in the frequency domain, capturing how components of a signal at different frequencies are distorted in amplitude or phase.

cross-correlation a measure of the similarity of two entire series when one is shifted relative to the other. For time-domain signals, the cross-correlation represents how similar two signals are when one is advanced or delayed against the other. As such, cross-correlation is commonly used to find a known pattern in a received signal and identify the time it was observed; indeed cross-correlation is optimal in detecting a known pattern in white noise (along with a matched filter). One example is a receiver finding a known preamble structure and establishing time synchronisation with the transmitter based on when the correlation peak occurs. A signal can be cross-correlated with itself as the *auto-correlation* to detect repetition.

delay line an artificial delay applied to a signal, either by use of a signal path that is physically long (e.g., a very long wire), a change of medium (e.g., converting a radio signal to ultrasound, such that it propagates slower) or a digital memory that stores samples for a defined period before replaying them.

digital signal processing (DSP) the processing and manipulation of signals represented as time series of discrete observations, as opposed to via some continuous physical process such as analogue electronics. Computational approaches can be used to mirror physical effects upon signals, along with any arbitrary processing that is desired. Depends upon the use of analogue-to-digital converters (ADCs) and digital-to-analogue converters (DACs).

digital-to-analogue converter (DAC) an electronic actuator that controls a continuous, analogue signal based on the input of a time series of amplitude values.

filter in general, any system that distorts a signal in an intended and beneficial way. Primarily, however, used to refer to components that filter out parts of a signal

at specific frequencies. *Low-pass* and *high-pass* filters attenuate high-frequency and low-frequency parts of a signal respectively, as determined by some designed *cutoff frequency*. A *band-pass* filter combines both properties, to allow only signals within a frequency range to pass. A *band-stop* or “notch” filter has the opposite effect; removing any signals that fall within the specified frequency range. A filter is characterised by its impulse response and can thus be either a finite impulse-response (FIR) filter or an infinite impulse-response (IIR) filter. In this thesis, all meaningful filtering is performed using FIR designs, with IIR filters only used tangentially, as a computationally-efficient approximation to a moving average.

→ **matched filter** a filter exhibiting an impulse response that matches a defined pattern, such that its output value is maximised when that pattern appears as input signal. An optimal detector for the defined pattern in a signal that also contains white noise (along with the cross-correlation).

forward error correction code (FEC code) a code applied to data before transmission in which some data redundancy is provided, such that a receiver can correct up to a certain number of bit errors, without requiring re-transmission. The level of redundancy, and thereby the maximum number of correctable bit errors, is a parameter of the coding scheme used [5]. Common examples are Viterbi codes (802.11a/b/g Wi-Fi), Turbo codes (4G LTE, HomePlug PLC), low-density parity check (LDPC) codes (802.11n Wi-Fi, 5G NR).

Fourier transform the transformation of a time-domain signal to a frequency-domain one. Used throughout signal processing whenever analysis or communication is more easily performed by handling the frequency representation of a signal. The opposite transformation is the inverse Fourier transform. The *fast Fourier transform* (FFT) and inverse fast Fourier transform (IFFT) algorithms [6] perform the transformation in $O(N \log N)$ time and are used universally throughout DSP.

→ **short-term Fourier transform** (STFT) the Fourier transformation of one section of a longer signal, that provides frequency-domain information for the specific period of time covered by the signal section. Often used to analyse time-varying behaviour of signal, such as when a transmission takes place intermittently or the waveform representation of a data packet changes structure throughout.

gain/loss the increase or decrease in strength of some signal or component thereof, used variously for signal amplitude or power. The use of “gain” is more common than “loss”, such that a signal is sometimes said to experience negative gain.

impulse response the effect of passing a single, infinitesimally-brief input (i.e., a Dirac pulse) through a system (e.g., an electronic component or propagation channel), as a function of time. The impulse response of a system fully characterises any distorting effect it has on input signals. It is the time-domain counterpart to the transfer function. The impulse response of some systems (e.g., a propagation channel) are out of the control of a designer and must be accommodated. However, signal processing systems are built specifically to exhibit a given impulse response, so that they affect an input signal in a beneficial way. This can be achieved with electronic components in an analogue radio system, or mathematically in digital signal processing.

noise unwanted and irrelevant signals measured by any practical receiver, distinct from the signal that is intended to be received. Noise can be a product of other transmissions (although if these become dominant they are referred to as *interference* instead), emissions from electrical devices nearby or internal effects in a receiver's circuitry. The *noise floor* for a particular set of observations thus varies substantially by location, environment, frequency band and equipment. In practical uses, a noise floor between -135 dBm and -85 dBm might be expected.

orthogonal frequency-division multiplexing (OFDM) a modulation technique that mitigates some challenging channel effects that occur for signals of large bandwidths. In most practical cases, the effect of environmental distortion during propagation varies substantially with frequency. For wideband radio transmissions this causes different effects across the various frequencies involved and makes accurately correcting for the channel challenging. OFDM splits a single wideband channel into a large number of narrowband *subcarriers*, each of which is narrow enough that environmental distortion can be considered consistent across its small slice of bandwidth. By appropriately selecting the number of subcarriers and sampling rate, the subcarriers can be made orthogonal, such that they do not interfere with adjacent subcarriers and can therefore be packed tightly without guard bands between each. OFDM is widely used in wireline and wireless digital communications, underpinning 802.11 Wi-Fi, HomePlug powerline communication, DVB-T digital television and downlink communication in 4G LTE and 5G NR. OFDM communication modulates data in the frequency domain, before converting it into a time-domain waveform using an inverse Fourier transform.

peak detection a method for identifying a spike in the amplitude of a signal, usually implemented by evaluating the gradient of the signal before and after a

point in time. The measured value of a peak is often less important than when it occurs, for example to acquire time synchronisation by detecting a peak in the correlation of a preamble with its known template

phase the offset of a periodic wave from some defined starting point. For example, with two signals $X = \sin(t)$ and $Y = \sin(t + \pi)$ (where t is some time unit), both waves have the same amplitude and frequency, but are not identical due to having commenced at different points in their periodic behaviour. X and Y could be described as being 180° (π radians) out of phase with each other.

preamble a known sequence that forms the beginning of a radio transmission, usually in a communication system that permits random access to the medium. The preamble waveform is typically known to all users of the technology and thus it can be used to establish time synchronisation or measure the distorting effects of a communication channel and derive a channel model.

radio environment the environment through which radio communication is taking place; a combination of factors including the medium itself, the landscape, the geometry and properties of nearby objects and the presence of other transmissions, either intentional or interfering. In an empty *free space* environment, only attenuation of a signal occurs and signal power at a receiver can be calculated by the Friis equation:

$$P_r(d) = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d^2 L}$$

where P_r is the received power at distance d , P_t is the transmission power, G_t and G_r are the antenna gains at transmitter and receiver respectively, λ is the wavelength and L is an additional loss factor to model other sources of power loss if they have been analysed.

In most situations, the radio environment is not accurately modelled as free space, and so communication is subject to the reflection and diffraction effects generally known as *multipath*, as well as further attenuation from any objects on the propagation path. The attenuation a signal experiences not only varies widely based on frequency and location, but also changes over a period known as the *coherence time*, which is often less than a second. Signal power measurements in a dense multipath environment with no line-of-sight follow the Rayleigh distribution, while environments that afford line-of-sight as well yield measurements following the Rician distribution with a scalable parameter K indicating the dominance of the line-of-sight component [1].

radio reciprocity the principle, deriving from Lorentz electromagnetic reciprocity,

that if a transmitter and a receiver were swapped and the same signal transmitted, the same signal would be received as before the swap. From a security perspective, this can imply that eavesdropping a signal allows injection of a signal, although it is not sufficient alone as requirements on the waveform and the necessary power at the receiver must also be met.

sample-clock offset (SCO) a manifestation of a difference in the clock frequency between a transmitter and receiver, causing signals to desynchronise in time. A counterpart to *carrier frequency offset*. The *analogue-to-digital converter* in a receiver takes samples of a signal at regular intervals, driven by some reference signal that indicates when these intervals occur. If a transmitter and a receiver do not have identical reference signal frequencies, then samples will be observed at different rates and will appear to slowly fall out of time synchronisation.

signal power as radio receivers typically measure voltage, the power of a radio signal is proportional to the square of the signal amplitude, either instantaneously as $P = \frac{V^2}{R}$ where V is the measured amplitude and R is the resistance of the system (commonly 50Ω , by design), or averaged over a period as $P = \frac{V_{RMS}^2}{R}$ where $V_{RMS} = \sqrt{\frac{1}{n}(v_1^2 + v_2^2 + \dots + v_n^2)}$. Signal power in radio communication is typically quoted in dBm; a decibel ratio related to a reference power of 1 mW.

signal-to-noise ratio (SNR) the ratio of the amplitude of the intended signal to the noise. Communication reliability is often expressed in terms of SNR, or SNIR (signal-to-noise-and-interference ratio).

software-defined radio the use of DSP to implement radio communication or analysis systems. Characterised primarily by the aim of minimising the analogue circuitry in any transmission or reception chain and extending the ‘digital boundary’ such that software control and configurability can be enjoyed throughout as much of the system as possible.

thresholding flattening of a signal to a binary 0 or 1 based on comparing its amplitude to a reference value.

transfer function the distortion applied to an input signal as it is passed through a system, as a function of frequency. It is the frequency domain counterpart to the impulse response. Can describe systems out of the designer’s control (e.g., propagation channels) or intentional distorting elements (e.g., filters). The transfer function is often more amenable to human analysis than the corresponding impulse response.

triggering or *squelching*, in which signals are discarded unless some trigger condition is met. A trigger condition might commonly be an absolute signal power level, a substantial change in signal power, or the detection of a known pattern such as a preamble.

tuner a component in a radio system that performs configurable frequency translation, allowing a transmitter or receiver to operate at frequencies that exceed its maximum operating bandwidth. For example, a transmitter may produce a signal with maximum bandwidth of 1 MHz, which if transmitted directly could only occupy the band 0 Hz – 1 MHz (*baseband*), if the signal is passed through a tuner configured for 500 MHz transmission, then the signal will occupy the band 500 MHz – 501 MHz (*RF band*). A tuner makes use of a local reference signal of some known frequency and multiplies that frequency to the desired tuning point.

turbo code a type of *forward error correction* code that uses an iterative decoding scheme. Data bits are encoded by *convolutional codes* using a known boolean polynomial function in which several output bits are affected by each input bit. A decoder observes the (potentially erroneous) received bitstream and estimates the most likely input that could have resulted in the received bits. A turbo code differs from a plain convolutional code in the use of multiple encoding stages, that are each applied to different orderings of the input bits. As channel errors often occur in localised bursts, combining the probabilities output from each decoder allows remaining errors in one ordering of the estimated input bits to be corrected by another stream in which the errors affected different bits [5].

wireless communication strictly-defined, any communication that does not use a wire or waveguide, whether it use electric fields, magnetic fields, visible light, infrared, acoustics, ultrasonics or another physical phenomenon. In practice, the term usually refers to electromagnetic communication by manipulation of electric fields, or “radio”². Radio systems make use of a large swath of the electromagnetic spectrum, with frequencies typically ranging from 30 kHz (*LF* or “long-wave”) to 300 GHz (*EHF* or “millimetre-wave”).

²It is also used to refer to manipulation of magnetic fields instead, such as for near-field communication (NFC), although they are based on different physical phenomena.

1

Introduction

Contents

1.1 Contributions	5
1.2 Thesis Organisation	6
1.3 Publications	8

All computer systems fundamentally rely on physical phenomena. For most, this reliance is purely as a means to realise the computational model; building on mechanical procession [7], digital electronics [8] and latterly quantum processes [9].

For security, this dependence is in the concept of a trusted computing base; those elements of a system which must remain untampered if the system is to meet its security promises. The malicious use of physical effects¹ against the trusted computing base has given us a long history of famous security compromises, primarily in the form of side channels and fault injections. Some of the most well-known have been side-channel attacks on cryptographic implementations, for which every effort is made to minimise the trusted computing base. The literature is rich with attacks exploiting differences in timing [10], power consumption [11], acoustic

¹Here and throughout this thesis, “physical” is meant in the sensing, actuation and networking sense, rather than the sense of walls, locks and guards (“physical-layer” rather than “physical security”). Physical-layer attacks are hence those that undermine assumptions made about physical phenomena in a system’s design, rather than the many security breaches that can be achieved through grabbing or smashing.

cues [12] and electromagnetic emissions [13]. Fault injections have again used electromagnetic signals, controlled temperature changes or high-intensity visible light to change stored data [14, 15]. These instances are but the tip of an iceberg, yet the scope for attack even here is testament to the range of physical effects that are used or produced, even by systems that are not intentionally designed to interact with the physical world.

Cyber-physical systems, however, *are* intentionally designed to interact with the physical world. Their reliance on physical effects for sensing and actuation is not only to underpin their computation, the phenomena are instead an integral part of the analysis and decision-making process. As such the trusted computing base necessarily expands and a consistent challenge faced in securing cyber-physical systems is in ensuring that the logical model of the world that is being processed is consistent with physical fact. Manipulating sensors can lead to disastrous consequences [16, 17] and the problem was brought into stark focus with discovery of the famous Stuxnet attack [18].

Indeed some security properties that are highly desirable for a system simply cannot be realised at a logical level alone.

An aircraft, reporting its position to air-traffic control, cannot certify just by the movement or transformation of data that the position it reports is its actual location. It must be trusted that it will not report incorrectly; whether due to malice or subversion. Trust in the aircraft can be bolstered through its proven possession of a secret, trust in the implementation by a tamper-resistant root-of-trust in its hardware — but the trust must also extend, without proof, to the assumption that whatever sensing was used to derive its location in the first place accurately reflected the physical world.

The problem can also be more subtle. The example just described concerned the provenance of a value that is used directly, but it can also concern ancillary attributes of a value or a process. Wireless local-area-networks, meant to blanket a space with connectivity, often cover areas that the operators did not wish them to; at best sharing the resource unintentionally and at worst expanding their attack

surface for no benefit. Cryptographic means can again attest to the identity of a network participant but cannot say anything of other properties that may be the determining factor in security decisions, such as whether it is in the building or not — a problem that plagues free café Wi-Fi the world over.

Relay attacks take the problem further, often quite literally. A system that controls access to a resource can take all manner of precautions at a logical level. It may authenticate its counterpart with cryptographic primitives and a provably-sound protocol. It may perform all appropriate checks on validity and revocation. It may be free of flaws in its implementation. Yet when it authorises a payment [19], unlocks a door [20] or activates a charger (§5.12), it still has proven nothing about the arrangement of the physical world, only of the data.

Common among all the examples above is the presence of communication, particularly wireless communication, and this reflects the existing identity of the field of physical-layer security. This appears to be primarily for two reasons: immediate need and the focus of early work.

Firstly, while wireless communication isn't a requirement for the described problems to exist, the extensive use of wireless networks in cyber-physical systems and the weak constraints on attackers in a broadcast environment originally put them at the forefront of attempts to enhance security by physical-layer means. Secondly, wireless communication (specifically electromagnetic communication) has the allure for theoreticians of one particularly-strong physical-layer guarantee. Faster-than-light travel is (as far as we know) *impossible* and therefore a worthy primitive with which to build security protocols. Seminal work to exploit this exciting physical limit [21] and attempts to formalise physical phenomena for protocol analysis [22] thus focused on radio communication — the impact of that work appears to have stuck, despite the capability to realise such systems only appearing years later [23].

Realistically though, factors that are possible-but-difficult are usually suitable as well, for some required level of security². Over time, a wider range of techniques have come to be included in physical-layer security. Some examples are environmental,

²Ultimately, even systems working with immutable physical laws can still have implementation issues [24]

such as it being infeasible to fully control a real-world radio environment [25] or obtain an identical channel model to a target [26]. Many are implementation-focused, such as the difficulty in manufacturing components to arbitrarily-low levels of variation and how this manifests itself in communications behaviour. Still others are processing challenges such as the challenge of retrieving unknown information from below the noise floor [27]. The term of physical-layer security still denotes a field primarily concerned with radio communication however.

An enormous array of other factors are available outside of wireless technology too though and while they are not yet popularly represented under the term of physical-layer security, *the conceptual application of physical factors to stymie an attacker is the same*. Some shared concepts have appeared in new environments [28], some creative uses of wireless networks have yielded ways to reason about the physical world [29, 30], and the general use of sensors is abundant across security study, from biometrics to secure hardware. Formal modelling of physical features is even beginning to incorporate more variety such as the inclusion of fluid levels and pressure in attack modelling in [31].

Indeed, measuring physical phenomena and incorporating those measurements into systems has, of course, long been possible and specialist security systems have existed to monitor individual domains for years. However, it has generally been too expensive, inefficient or difficult-to-scale for the techniques to be applied on a wider scale. Similarly, physical-layer attacks had previously remained the preserve of the well-resourced adversary. But the growing availability of commodity sensing and actuation technologies — that are cheap, interface widely and offer rich control from software — has drastically lowered the threshold for mounting physical-layer attacks. Attacks that previously required specialist technical knowledge are now accessible to the layperson with a shopping list and a download.

The same technological trends also now permit the wider use of physical factors for security and the change in practice from developing bespoke systems for specific scenarios, towards incorporating physical considerations as part of standard design practice. The work herein does still use electromagnetic effects; radio transceivers

are still one of the most widely-deployed sensing and actuation technologies and are abundant in cyber-physical systems, so their use is prudent. But there are many ways in which electromagnetic effects can provide insight into the physical world [32, 33], *outside the context of communication itself*, and the use of these broader effects is explored. Herein, some measurements are used directly (e.g., timing, power), while others form convenient proxies from which to derive other physical properties (e.g., proximity, distance, movement, environment) that are ultimately being used.

This thesis presents a series of novel contributions in support of the argument that the exploitation of physical factors to secure systems is *possible* across a range of fields using commodity hardware and that it is *necessary* to ensure that cyber-physical systems are not exposed to critical attacks.

The argument is demonstrated with systems to help secure legacy technologies for air-traffic monitoring and in-home networking, as well as to realise a practical security system for defence against the modern threat of drone surveillance. It is further argued that a lack of consideration for physical effects is not only a problem for legacy systems, but immediately haunts brand-new systems as well. This is shown with a practical eavesdropping demonstration on a widely-deployed vehicle charging system and a proposed active attack that exploits a discrepancy between the logical state and the physical world — the core problem that motivates physical-layer security.

1.1 Contributions

The work comprising this thesis makes the following novel contributions to the field of systems security:

- Presents a practical, secure location verification system for air traffic position claims that operates with fewer receivers than static-receiver systems and requires no change to existing aircraft hardware or communication protocols
- Presents a drone detection system, to mitigate privacy-invasion attacks that uses only widely-available radio signal-strength metrics and requires

no specialist knowledge to deploy, making it the first applicable to normal individuals and domestic environments

- Introduces the security relevance of unintended wireless coupling in powerline communication
- Presents a detection system for malicious powerline networks, exploiting unintended wireless radiation to improve the performance and practicality over a system that uses conducted signals
- Argues that exploiting physical phenomena to secure systems is practically achievable across a wide range of fields using commodity equipment
- Describes the first software receiver for broadband powerline communication technologies
- Demonstrates an eavesdropping attack on state-of-the-art electric-vehicle charging communication and describes latent physical-layer attacks to disrupt communication or conduct billing fraud
- Argues that the incorporation of physical-layer information is necessary in securing modern cyber-physical systems

1.2 Thesis Organisation

The thesis is presented with chapters that each describe work on an individual system. Given the disparate nature of the systems involved and the little overlap of background description among them, a chapter on shared preparatory material is omitted and relevant background is instead presented within each chapter.

In Chapter 2, location claims made by aircraft are verified by multilateration using mobile receivers. While much can be done to protect the confidentiality, authenticity or integrity of location claim messages at the logical level, no amount of doing so can guarantee that the claimed position matches the real location of the aircraft. While multilateration systems have existed for many years, the proposed system enables non-cooperative verification using fewer receivers than the minimum

$dimensions + 1$ required of the common, static-receiver architecture. The use of randomised mobility also increases the difficulty for an attacker attempting to fool the system by observing the positions of the verifiers. The system is further able to localise a transmitter with similarly-few receivers, albeit only where the receivers have significantly greater mobility than the target.

In Chapter 3, a lightweight system for drone detection is described, which uses easily-available signal-strength measurements and a set of statistical tests to identify whether a detected signal is being emitted by a drone. An examination purely of the data that are exchanged between two devices cannot determine that they are a drone and its controller, however the incorporation of physical measurements allow this identification to take place.

Chapter 4 discusses the threat of rogue access point attacks using powerline communication, an attack which can be mounted with similar ease by non-technical attackers as the well-known Wi-Fi equivalent. The propensity for powerline communication to leak signal radiatively is exploited to enable wireless detection of the attack. Wireless detection, in turn, eases deployment by removing the need to follow complex wiring runs and minimising the effects of localised noise sources that inhibit conductive sensors. Furthermore, the use of wireless techniques for detection enables familiar localisation techniques that operate in 3D space to be applied to find offending devices. This would be far more error-prone in conductive detection where distances can only be measured along wiring runs, rather than by straight-line distance.

In Chapter 5, the effects of powerline radiative leakage are shown instead in an attack context. Properties of powerline communication (particularly the use of orthogonal frequency-division multiplexing and unbalanced wiring) are shown to permit simple message recovery from wireless emissions. This observation is applied to a state-of-the-art electric vehicle charging system and shown to be effective at receiving communications between the vehicle and charger. Analysis of the charging communication standard shows that signal leakage is exacerbated by design compromises and that, despite the modernity of the system, physical-layer

security is abandoned by choice for implementation convenience. A number of active physical-layer attacks are then also theorised against a new automated-billing system that is being deployed to protect electric vehicle charging payments. The attacks can prevent a driver from charging, or enable one driver to charge their vehicle at the expense of another. These attacks further highlight the impact of physical-layer attacks, while the theft attack, in particular, depends on the same problem of matching physical and logical state, that has been discussed throughout.

In concluding remarks, the thesis argues that systems design must give greater consideration to the effects of manipulating physical phenomena and that binding logical state and physical state are necessary for secure cyber-physical systems design.

1.3 Publications

The work comprising this thesis has given rise to several peer-reviewed outputs, which are detailed below. Each publication is derived from a corresponding chapter herein.

Publication: R. Baker and I. Martinovic, “Losing the Car Keys: Wireless PHY-Layer Insecurity in EV Charging”, in *Proceedings of the 28th USENIX Security Symposium (USENIX Security 2019)*, 2019.

Publication: R. Baker and I. Martinovic, “EMPower : Detecting Malicious Power Line Networks from EM Emissions”, in *Proceedings of the 33rd IFIP International Conference on ICT Systems Security and Privacy Protection (IFIP-SEC 2018)*, 2018.

Patent: R. Baker and I. Martinovic, Oxford University Innovation Ltd., “Power line communication detection”, UK Patent Appl. GB1805071.6, 28th March 2018.

Publication: S. Birnbach, R. Baker, and I. Martinovic, “Wi-Fly? : Detecting Privacy Invasion Attacks by Consumer Drones”, in *Proceedings of the 2017 Network and Distributed System Security Symposium (NDSS 2017)*, 2017.

Publication: R. Baker and I. Martinovic, “Secure Location Verification with a Mobile Receiver”, in *Proceedings of the 2nd ACM CCS Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC 2016)*, 2016, pp. 35–46.

2

Secure Location Verification

Contents

2.1	Introduction	11
2.2	Background	12
2.2.1	Cooperative Surveillance	12
2.2.2	Secure Location Verification	13
2.2.3	Unmanned Aerial Vehicles	15
2.2.4	Air Traffic Management	16
2.3	Related Work	17
2.4	Attacker Model	18
2.5	System Model	19
2.6	Effect of Errors	22
2.7	Security Analysis	27
2.8	Practical Considerations	31
2.9	Prototypes	32
2.9.1	Ground-based Prototype	33
2.9.2	Aerial Prototype	35
2.10	Evaluation	35
2.10.1	Simulation	35
2.10.2	Real-World Data Collection 1 (Ground)	41
2.10.3	Real-World Data Collection 2 (Airborne)	44
2.10.4	Verification of Real Tracks	47
2.11	Discussion	47
2.11.1	Alternative Configurations	48
2.11.2	Potential Applications	49
2.12	Further Work	51
2.13	Conclusion	53

2.1 Introduction

This chapter considers a relatively direct application of physical-layer security methods to a large cyber-physical system; a cooperative surveillance system. These are systems that allow the monitoring of node positions by relying upon voluntary reporting. Such systems have existed for many years and are in widespread, production use in critical systems such as aviation traffic management. A key challenge with cooperative surveillance systems is the reliance upon the reported state matching the actual physical state. While non-cooperative surveillance systems make direct measurements of the physical state and are robust to false reporting, cooperative systems that rely upon a logical protocol are not able to determine the physical state. There is instead a reliance on the trustworthiness and accuracy of the reporting party. This leaves such systems vulnerable to incorrect reporting of false locations. This can be due to benign errors that are not detected before reporting, such as sensor failure or misconfiguration. It can also be the result of malicious spoofing attacks in which reports are intentionally falsified.

This chapter demonstrates a secure location verification system for aircraft position reporting. The system enhances an existing logical-layer reporting system with physical-layer information that can be used to detect claims that are not consistent with actual positioning. The system makes use of a mobile receiver to increase the challenge for an attacker. The movement of the receiver progressively reduces the number of locations from which an attack can be conducted, until eventually the only viable location to report is the true location. The system represents a non-cooperative adaptation of the model by Čapkun et al. [34] and additionally addresses deployment practicalities.

The presented technique performs secure location verification of aircraft position claims by measuring the time-difference of arrival (TDoA) between a fixed receiver node and a mobile one. The mobile node moves randomly in order to increase substantially the difficulty for an attacker to make false messages appear genuine.

An exploration is made of the performance and requirements of such a system in the context of verifying aircraft position claims made over the Automatic Dependent Surveillance – Broadcast (ADS-B) system. Through the use of simulation it is found that the system correctly detects false claims with a peak accuracy of over 97% for the most complex attack modelled; requiring only 75m of deviation between the reported position and the actual position in order for a false claim to be detected. A design for a mobile receiver is then presented and the construction of a prototype using low-cost COTS equipment is reported. The additional benefits of incorporating a mobile node, the difficulties to be overcome and the applicability of the approach in other location verification use-cases are then explored.

2.2 Background

2.2.1 Cooperative Surveillance

Cooperative reporting systems rely on individual nodes in the system to actively provide an indication of their state, such as their location, to receivers. In contrast with non-cooperative or “primary surveillance” systems such as traditional radar, which actively produce a means of tracking nodes, cooperative systems must rely upon each node to be able to determine its state accurately of its own accord and then report it. For position claims this means the node must be able to localise itself and then provide that location.

Cooperative systems are widely deployed. Air traffic management makes use of Automatic Dependent Surveillance - Broadcast (ADS-B) to track aircraft (discussed in more detail below). The Automatic Identification System (AIS) is a marine tracking system in which vessels determine their location and then broadcast it over a VHF radio link [35]. Sensor networks in which the geographic location of sensors gives meaning to their measurements need the sensors to report that location along with their collected readings. Similarly, knowledge of the locations of nodes is crucial when employing geographic routing or to assess distribution of sensors to ensure appropriate coverage [36]. Connected vehicles can report status including position via Vehicle-to-Infrastructure (V2I) links in appropriately-equipped cities,

these reports can be used to enhance awareness of traffic patterns, as well as be used to implement road pricing and pay-as-you-drive insurance schemes [37, 38].

Systems vary widely in their location accuracy, transmission range, frequency of reporting and between those that openly broadcast information in the clear (such as ADS-B or AIS), via those that establish transient links (such as connected vehicles forming platoons) through to those that provide only anonymous reports using secure channels. Where transmission is made by open broadcast, a significant security consideration is the possibility for a malicious party to report false information; to mimic other nodes, to invent fictitious nodes or to alter legitimate reports. Even where the channel is cryptographically secured to ensure message integrity, the problem is not alleviated; knowing that the received message is the same as the one that was sent still does not guarantee that it is an accurate representation of the physical state. Attacks can affect any dependent systems, whether they are centralised control or reporting systems, other nodes in the system or even human operators. Attacks on operators include overloading them with information so that genuine reports go unhandled, or even simply placing them under high levels of stress to increase the risk of human error occurring [39, 40]. The open provision of information about individual nodes in the system gives rise to privacy concerns as well; particularly the easy tracking of nodes en masse.

2.2.2 Secure Location Verification

It is often necessary to have a means of verifying that location claims made by nodes in a cooperative reporting system are genuine. Verification can exploit any property of the transmission that is difficult or expensive for an attacker to influence maliciously; ideally prohibitively so. Common techniques include fingerprinting, distance-bounding, received signal strength measurement, angle-of-arrival and time-of-arrival.

Fingerprinting identifies operating features of the transmitter or protocol implementation that differ between classes of transmitter, or even manufacturing imperfections that are unique to individual units (e.g., clock offsets, rise/settle time

differences). If expected values are known for a given transmitter then computed features from a particular transmission can be compared to assess whether it is genuine. Distance-bounding techniques place an upper bound on the distance of a communicating party by using a challenge-response mechanism and measuring the round-trip time. With appropriate hardware the travel time of the signal is the dominant factor and the responder can be constrained to within a certain distance of the challenger. Received signal strength approaches are based on understanding of the attenuation of radio signals during travel. For a given type of environment the expected attenuation over various distances and from various locations can be estimated. With known transmission power the power of the signal at the receiver allows the verifier to identify ranges or locations from which the transmission could have been made. Techniques based on “angle of arrival” can be used to determine a direction to the transmitter and can be used either with ranging techniques or cooperatively in a “triangulation” arrangement to determine the transmitter location. Either mechanically-revolved directional antennae must ‘sweep’ an area or an antenna array must be employed and arrival times compared at each antenna to compute the angle of arrival. Time-of-arrival techniques measure the time that is taken for a signal to propagate through a medium and use those measurements to constrain possible transmission locations. In a pure time-of-arrival (ToA, or “trilateration”) system the time of transmission is known and so the time-of-arrival at a receiver indicates the range from which the transmission was made. Where the time of transmission is not known, instead the can verifier note precisely when a signal arrives at a set of receivers at known locations. These “time-difference of arrival” (TDoA, or “multilateration”) readings can then be used by the verifier to identify a set of locations from which the transmission must have come. With more receivers the verifier can constrain the set of locations further, eventually down to a single point [38].

2.2.3 Unmanned Aerial Vehicles

Unmanned Aerial Vehicles (UAVs or ‘drones’) have become extremely popular in recent years; employed in military and industrial roles and flown recreationally by individuals. A multitude of aircraft are available, being variously fixed-wing or rotor-wing, electrically or chemically powered and remote-controlled or autonomous. The continuing development of compact, inexpensive control and sensing equipment has made construction of UAVs far more widespread [41]; indeed the low cost of many designs has made UAV usage feasible in situations where an aerial system has previously been unavailable or too expensive to run, such as internal building survey, hobbyist photography or emergency rescue in hazardous areas [42]. Additionally, operating a UAV requires far less skill than a conventional aircraft so usage requires only minimal training and experience. Indeed the prevalence of UAV operators engaging in risky behaviour has prompted regulatory responses. The United Kingdom’s Civil Aviation Authority (CAA) mandates that general operations must remain within 500m of the operator, below an altitude of 400 feet and in direct line-of-sight at all times. In December 2015 the U.S. Federal Aviation Administration (FAA) mandated registration of all operators of UAVs with a mass of over 250g [43], to improve operator accountability. The implementation of ‘no-fly zones’ for UAVs is already widespread, with various technical means such as ‘geofences’ being employed in an attempt to enforce them, albeit with mixed success. In addition, the CAA specifies tighter controls on “any aircraft which is equipped to undertake any form of surveillance or data acquisition”; mandating a separation of 50m from any person, building, vehicle or vessel not under the operator’s control, in an attempt to mitigate the associated privacy risks [44]. Many UAV systems already report telemetry to their controllers that includes position and movement details and it appears likely that future regulations will require this information to be reported to controlling authorities as well.

Many UAVs are also capable of performing autonomous flights and there is considerable work to enhance this capability. Incorporating autonomy can not only reduce the risk of harm from an inexperienced operator, but also supplement

operator judgement in hazardous situations (using extra sensors and collision avoidance systems) and reduce operator workload to enable several UAVs to be managed simultaneously by a single human controller. Sufficient technical advances in UAV autonomy and regulatory changes could enable use at greater distances and in far more complex environments, such as for delivery or survey purposes [45, 46].

2.2.4 Air Traffic Management

Air traffic management attempts to ensure safety is maintained during aircraft manoeuvres in the governed airspace, whilst maximising the efficiency of traffic movement. Air traffic controllers maintain communications¹ with aircraft as they operate and provide advisory information or mandatory instructions to pilots. In addition, air traffic control (ATC) operations have long made use of surveillance technologies to assist in tracking aircraft for which they are responsible. Originally this function was performed by non-cooperative means via “primary surveillance radar” and this is still common in military uses. However primary surveillance technologies provide limited information beyond the range and bearing of an object and require considerable infrastructure and expense on behalf of the ATC operators [47]. Practice is moving away from traditional active surveillance and towards a cooperative reporting model in an effort to enhance situational awareness for all parties and accommodate greater traffic; so-called “secondary surveillance radar” (SSR). In this model, aircraft broadcast a variety of status information, which can be received by ground stations and other aircraft. One such SSR, already widely deployed in Europe and the United States (and mandated for use by 2020 by regulators in both jurisdictions) is the Automatic Dependent Surveillance - Broadcast (ADS-B) system [48]. Aircraft are equipped with a transponder that periodically reports details of the aircraft’s status in the ADS-B format. ADS-B makes use of one of two possible data links; with the standard for civil aviation being 1090MHz Extended Squitter (1090ES). Earlier SSR systems made use of the 1090ES data link and so its use for ADS-B allows integration with existing

¹The primary communication method is currently voice, although this will be replaced by data communication in the near future

transponder equipment. Messages can variously indicate speed, position, callsign, climb rate and emergency status², although only position reports are considered in this work. All messages contain an ICAO identifier, a unique 24-bit number identifying the aircraft. The interval between messages varies depending upon the message type; for position messages it is approximately 0.5s [49].

2.3 Related Work

Location estimation with mobile nodes is described by Luo et al., using a similar approach to that described herein (termed “Mobility-Differentiated Time Difference of Arrival”) and applied to the issue of surveying sensor networks for node displacement [36]. However the approach is node-centric (i.e., the node is attempting to estimate its own location) and does not aim to be secure, in that it does not consider adversarial behaviour in the location estimation process.

Perazzo et al. describe a roving verifier that determines whether nodes in a sensor network have been displaced and describe an algorithm to construct a near-optimal route for the verifier to take to conserve fuel. However the localisation approach assumes cooperation from the sensor nodes and hardware to enable a distance-bounding protocol to be used [50]. Čapkun et al. proposed a system that exploits the difficulty for an attacker in claiming a false location when the verifier is moving in an unknown way [34]. However the location verification protocol is again cooperative.

Strohmeier et al. have proposed a location verification system that makes use of widely-distributed, low-cost receivers to improve the coverage of TDoA verification over existing, professionally-deployed systems. They overcome the poorer accuracy of low-cost receivers by collecting multiple messages and testing them together; comparing them statistically to expected values. However, they depend on widespread deployment of receivers and a prior training phase to produce fingerprints that message timings can be compared to. We instead attempt to make the verification process secure solely through the use of mobility; with no prior system training and fewer receivers.

²Among many other pieces of status information

Schäfer, Lenders and Schmitt describe a static multi-receiver ToA approach that considers a ‘track’ of position claims and uses it to compensate for individual errors to improve verification accuracy. However the model described therein assumes an accurate prover-local timestamp is sent with the message [51] and the authors note that while this capability can be realised from the ADS-B standard, deployment of systems that broadcast on such a dependable timescale is not widespread. Position claims are otherwise sent with much less accurately-measured periodicity.

2.4 Attacker Model

The attacker intends to influence the ATC system by sending false location claims; to invent ‘ghost’ aircraft in the skies. The very act of transmitting ADS-B position messages is an implicit statement that there is an aircraft at the claimed location which must be dealt with by other airspace users. For a meaningful attack, the attacker not only needs to do this, but must be able to perform the attack for some claimed location *of their choosing*, so as to impact a particular target.

Their attack depends upon their messages being received successfully, either by a particular target or by all nearby receivers as the attack requires, and so the attacker must continue to transmit at or above some minimum rate for the duration of their attack. We assume that attackers broadcast with an identical capability to legitimate sources and on a defined interval. As the system being attacked is designed to receive public broadcasts, no covert, point-to-point communication takes place between the attacker and the reporting system. Messages are further assumed to be broadcast and so are receivable by correctly configured receivers that are aware of the protocols involved, within some maximal range and with a proportion of messages being lost.

Attackers are considered as belonging to one of three classes, each modelling a different attack case. The attacker classes are as follows:

Static attacker The simplest attacker class; a static attacker is broadcasts from a fixed location but claims to be in another. This class exemplifies an attacker inventing a ‘ghost’ vehicle from their home or a parked vehicle nearby.

Mobile attacker A mobile attacker is a more complex version of the static attacker. In this case the attacker broadcasts from an airborne mobile platform, modelled as a UAV. This class covers an attacker using a UAV to attempt to fool the verification by legitimately moving the transmitter, but while claiming to be another (perhaps much larger or more important) vehicle, or simply mounting their receiver on a UAV to avoid being caught and the equipment confiscated. It could also model a drone operator flying outside of a permitted area whilst claiming to be inside.

Course deviation attacker A course deviation attacker initially broadcasts correct position claims, but then selects a new course and attempts to hide this behaviour by falsely claiming to still be following their old course. For air traffic this attacker class can be seen as an aircraft that has been seized by force or maliciously diverted. Less morbid alternatives exist in other scenarios, such as an attacker in a road vehicle briefly exceeding a speed limit or using a prohibited lane, while continuing to report seemingly law-abiding behaviour.

We further assume that the attacker does not possess the capability to locate the mobile receiver(s). The validity and implications of this assumption are investigated later in Sections 2.7.

2.5 System Model

In essence the approach consists a number of receivers $R1...Rn$ and a centralised processing unit. In general there can be any number of receivers, within which at least one must be mobile.

Although the system generalises to three dimensions straightforwardly, for simplicity we describe the concept in two dimensions throughout this section, with two receivers of which only one is mobile (termed the *Fixed receiver* and *Mobile*

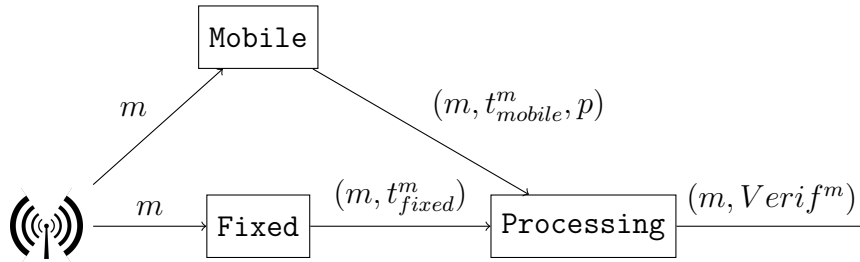


Figure 2.1: Structure of verification system

receiver). Both receivers collect the broadcast messages with claimed locations and record them along with the precise time of reception. The mobile receiver additionally records its location at the time of reception. The recorded (message, time, location) triplets are then sent via an out-of-band channel to the processing unit, which matches claims from each receiver by means of a short-term unique identifier for the message³. Upon receipt of matching claims from both nodes, the processing unit performs a TDoA calculation.

Figure 2.2 demonstrates an example operation of the system in two dimensions as it evolves through time. In this case, the source $S_{0 \rightarrow 2}$ moves on a linear course, broadcasting messages at a regular interval, while the mobile node $Mobile_{0 \rightarrow 2}$ moves randomly during the same period. When messages are received by both the fixed and mobile nodes, given that the locations of each node are known, the difference in elapsed time for a message to reach each receiver can be used to compute the possible source locations, which are given by the hyperbolae shown. In each case, only one side of each hyperbola is plotted, as it is determined by which node receives the message first (the fixed node in this example). The expected time difference from the claimed position can also be calculated easily and any discrepancy from the measured difference determined. If the discrepancy is within a defined acceptance threshold α then the message is considered genuine, otherwise it is taken to be false and flagged as such; enabling it to be reported to operators or downstream control systems.

³The identifier is situation-dependent, but is usually simply the message content itself, perhaps augmented with reasonable bounds on the time of arrival to avoid matching two subsequent messages with identical data against each other when their time difference is unreasonable.

More explicitly, for a position claim m , the distance from it to the two known verifier node positions can be calculated as the Euclidian distance in each case; d_{fixed} and d_{mobile} . Then the expected time difference for the claim $\Delta_{expected}^m$ can be calculated by finding the absolute distance difference and dividing by the propagation speed c (i.e. the speed of light).

$$\Delta_{expected}^m = \frac{|d_{fixed}^m - d_{mobile}^m|}{c}$$

Meanwhile the actual measured time difference Δ_{actual} can be obtained easily from the recorded values at each node.

$$\Delta_{actual}^m = |t_{fixed}^m - t_{mobile}^m|$$

Verification of the claim x is then a matter of comparing the deviation of the actual time difference from the expected time difference against an acceptance threshold α , to produce a result $Verif^x$ that can be output.

$$Verif^m = \begin{cases} Accept & |\Delta_{expected}^m - \Delta_{actual}^m| < \alpha \\ Reject & |\Delta_{expected}^m - \Delta_{actual}^m| \geq \alpha \end{cases}$$

The mobile node is, of course, mobile. It randomly and independently makes changes to its course. These are not pre-determined nor known by any other party. Furthermore, they are not transmitted prior to the course being followed. Upon initialisation, the mobile node begins listening for position claims and commences its movement; it selects a random waypoint and begins moving towards it, upon reaching the waypoint it selects a new one and repeats the process.

It is highly situation-dependent what action should be taken in the case of a claim being falsified. Various reporting, filtering or enforcement options exist but these are beyond the scope of the verification system itself. Some instances are discussed in Section 2.11.

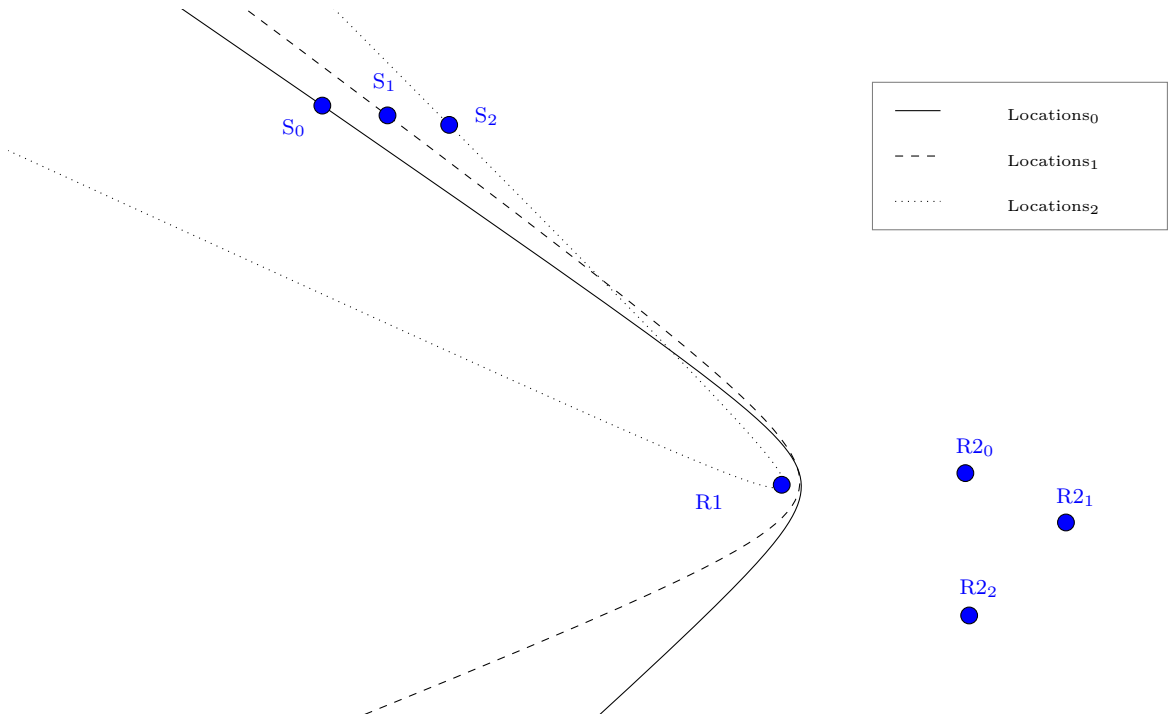


Figure 2.2: A two-node TDoA system in two dimensions at three successive points in time (as denoted by subscript).

2.6 Effect of Errors

The system is affected primarily by two sources of error, as shown in Figure 2.6:

Timing error Errors in the TDoA calculation can arise due to inaccuracy in the measurement of the arrival time, or due to poor clock synchronisation between receivers and the processing unit; altering the eccentricity of the hyperbola. Timing error affects the Δ_{actual}^m value directly. Intuitively, these errors can be thought of as giving the hyperbola some ‘thickness’. The magnitude of the error increases with distance and affects the whole span of the hyperbola branch. The scale of the error is also heavily impacted by the angle between the receivers and the target. Figure 2.7 visualises the effects of varying the angle, whilst applying the same level of error to a TDoA measurement. The greatest uncertainty results when the target is close to the line connecting two receivers and progressively diminishes as the target approaches the perpendicular.




Figure 2.3: Diagram of a false location claim being detected. The claimed position is indicated by a red +, but is broadcast by an attacker positioned at the black \times . The two receivers are each indicated by a black \bullet and the hyperbola constructed from the TDoA measurement is in gray. The solid line indicates the branch of the hyperbola that the transmitter must be located on, based on which receiver detected the message first. A claimed position anywhere on the same branch would appear genuine, but not a location off the hyperbola, as attempted here.

Positioning Error Incorrect localisation of a receiver can introduce error in the TDoA calculation by altering the points that the hyperbola is constructed about. Figure 2.8 shows the effects of positioning error at various angles to the target. As with timing error, a shallow angle and greater distance both exacerbate the problem. Unlike timing error, positioning error only affects one arm of the relevant hyperbola branch. Indeed, while positioning error does lead to an incorrect hyperbola, the resulting hyperbola still passes through the target. This is in contrast to timing error that can cause the target to be missed completely. The positioning error




Figure 2.4: Diagram of the diminishing opportunities for an attacker as multiple messages are received. As in Figure 2.3, the attacker is located at the black \times and each receiver is a \bullet , but the second branch of each hyperbola has been omitted. For a single message they can claim any position on the hyperbola described by the TDoA measurement. However, here the resulting hyperbolas for two messages can be seen and are substantially different. The only position that lies on both is the attacker’s true location.

is at its lowest passing through the target.

We consider the effects of errors throughout later analysis, although the primary focus is on timing error as its effects are more pronounced than those of positioning error and in our practical observations it is generally of far greater magnitude.

We also discuss controlling the phenomenon of errors being amplified as the target approaches the receiver line. This is referred to as “geometric dilution of precision” (GDOP) and appears as errors from each measurement overlap and cannot provide any refinement over the other.

Timing errors can be further broken into *clock drift* and *measurement error* [36,




Figure 2.5: Representation of one sheet of the hyperboloid of possible locations for a 3-dimensional TDoA measurement.

51]. In our simulations over long time periods, clock drift ϵ_{drift} is modelled as a linear progression with coefficient t_{drift} applied over any interval between times t_i and t_j , while measurement error $\epsilon_{measure}$ is applied as independent normally-distributed noise to each measurement.

$$\epsilon_{drift}^{j-i} = (t_j - t_i) \cdot t_{drift}$$

$$\epsilon_{measure} \sim \mathcal{N}(0, \sigma_{measure}^2)$$

So an observation of the current time t_{obs} will include clock drift error applied since the start of the system at $t = 0$:

$$t_{obs} = t + \epsilon_{drift}^t + \epsilon_{measure}$$

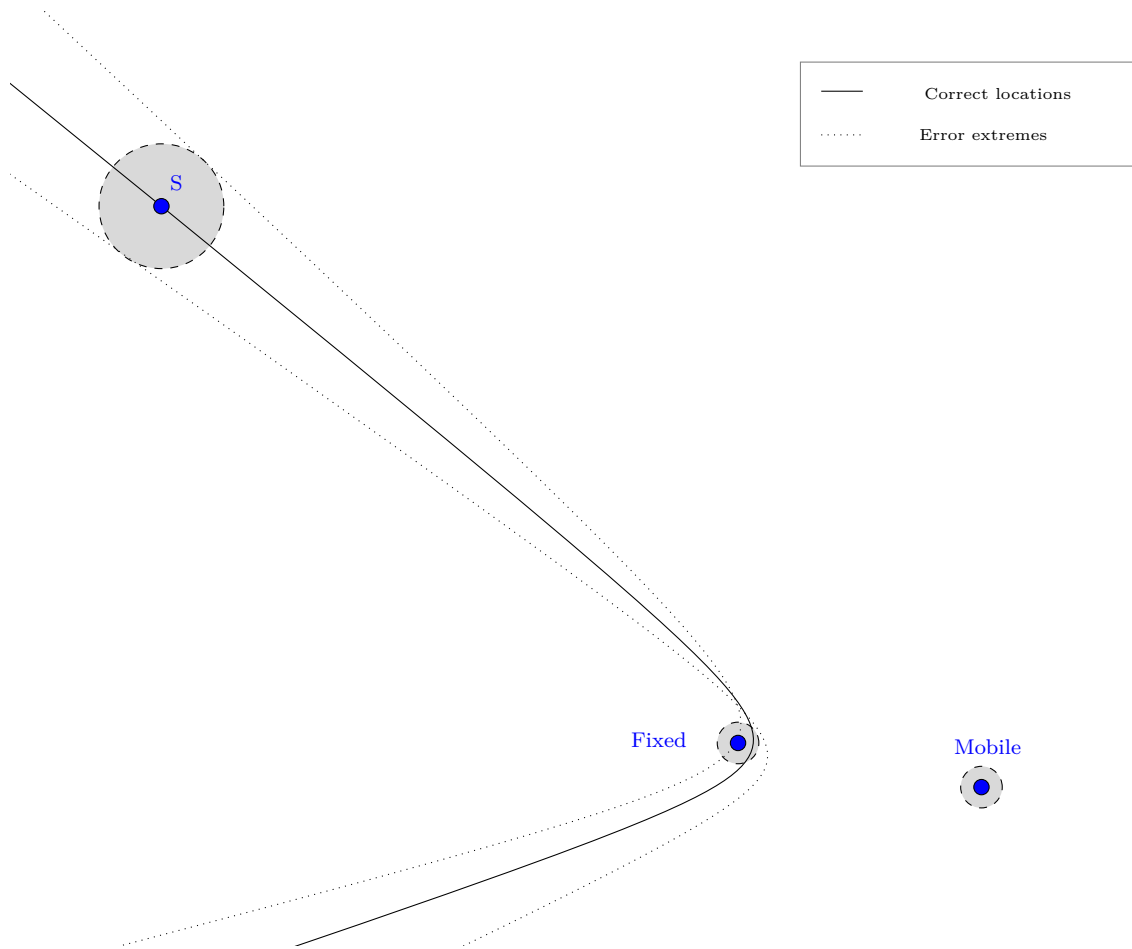


Figure 2.6: A two-node TDoA system in two dimensions at a single moment in time, with the errors for transmission time measurement and receiver localisation visualised.




Figure 2.7: Effect of TDoA timing error at different offset angles from receivers. A target at uniform distance is shown in each subfigure, but at different angles. The TDoA measurement is corrupted by timing error up to the same maximum in each case. The effect of timing error is most pronounced when the target location sits along the same line as the receivers. It progressively diminishes as the location approaches the perpendicular. In every case it increases with distance.




Figure 2.8: Effect of TDoA positioning error at different offset angles from receivers. A target at three positions is shown across the subfigure, but at different angles. The TDoA is corrupted by positioning error moving the second focus, up to the same maximum in each case. As with timing error, the effect of positioning error is greatest on the line of the receivers and smallest at the perpendicular. It also increases with distance. Unlike timing error, positioning error has different effects on the two arms of one hyperbola branch; with the target’s location experiencing the smallest effect from error.

2.7 Security Analysis

As with any verification system, an attacker can attempt to fool the process, subvert the system or disable it completely. In this section we consider each in turn and discuss the implications.

An attacker attempting to fool the verification process clearly wishes their messages to appear genuine. The capabilities of an attacker are considered in great detail in [39]; noting that in general they can perform message injection, deletion or modification.

In an attempt to have their injected messages accepted as genuine, the attacker can modify their transmission timing. This requires that they know the location of each receiver, such that they can broadcast messages that arrive at each receiver with a time difference consistent with the claimed location. If location verification is performed at a single point in time by more than two receivers, wherein the TDoA measurements all refer to the same time of transmission, it can detect transmissions that are not made at a single point when there are sufficient ($n \geq 3$ in 2 dimensions, $n \geq 4$ in 3 dimensions) receivers. Where there are fewer receivers, such as the minimal two considered in much of this work, the system cannot constrain an

attacker to less than an entire branch of a hyperbola. A message transmitted from any point on that side will appear to be genuine. The security of this approach lies in the fact that as the mobile node moves relative to the fixed one, the hyperbola of possible genuine positions sweeps across a substantial distance, this is particularly noticeable between the second (dashed) and third (dotted) lines in Figure 2.2. For each message the attacker has a very small chance of remaining on the valid hyperbola as the receivers move. An attacker reacting to the system must accurately determine the location of each receiver, construct the hyperbola and then transmit from some location on it in order for their message to be accepted as genuine. As the movement of the mobile node is randomised, the attacker cannot predict its location and so must monitor it and move accordingly, with sufficient speed to catch the arc of the hyperbola. Where the attacker is additionally required to transmit regularly, the maximum permissible interval between transmitting claims bounds how quickly the attacker must complete this process. In three dimensions with only two nodes, the location is constrained instead to one sheet of a hyperboloid (see Figure 2.5). The attacker therefore has another degree of freedom for their transmission location, but must still engage in the same reactive behaviour and must now contend with the movement of the mobile node in another axis as well.

As the effect of errors is most pronounced in areas of high geometric dilution of precision (GDOP), an attacker will benefit from locating themselves in such an area, where discrepancies in their positioning will be tolerated the most. However in our approach the areas of high GDOP also change randomly as the hyperbolas do — just as an attacker has difficulty in predicting where they must locate themselves, they also experience difficulty predicting where the areas of high GDOP will be at any point in the future. However, an attacker that is able to make the angle between receivers small relative to themselves can maximise the overlap for a given distance and error tolerance, whilst an attacker that can move further away from the receivers with the angle between them kept small can expand the overlapping region in absolute terms. With an airborne mobile node, our system does have a distinct advantage over purely ground-based systems in combating this situation.

The greater variation between receiver altitudes reduces the GDOP in the vertical dimension, thus somewhat restricting an attacker's use of altitude to engineer sufficient distance for high GDOP. The vertical range of the mobile node may well be limited, but as long as it is greater than the surrounding terrain elevation, the system benefits from the mobility.

Ultimately, an attacker transmitting from only a single point requires greater mobility than the verification system in order to fool it. The system poses a trade-off for the attacker: the further they are from the midpoint between the two receivers, the more space they are given by error factors in which to make false claims, however the faster they must move to remain on the hyperbola of valid locations as the receivers move.

Message deletion is possible primarily by selective jamming. As [39] notes, destructive interference (wherein the signal is inverted and superposed onto the original) is extremely difficult to accomplish for a moving aircraft, so we do not consider it here. Constructive interference (wherein a noise 'spike' is injected during message transmission such that the message is corrupted) is feasible and the system presented herein provides no defence against such attacks. Message modification is a derivative of the injection and modification approaches. Again, applying interference to alter a message in-flight is hard to achieve, however a message can be observed, interrupted with constructive interference and re-injected. The normal operation of the verification system for the message injection then applies, with the attacker's timing additionally delayed by having to observe the original message before transmitting.

Aside from attempting to fool the verification system, an attacker can attempt to subvert its operations or render it inoperable entirely. Jamming of the reporting link between the mobile node and the processing unit represents an appealing prospect for an attacker (likewise the reporting link from the fixed node if it is wireless). This would prevent the system from performing verification until either the jamming ceased or the mobile node moved to a location from which it could overpower the jamming signal (e.g. returning to the fixed node, if such functionality were

implemented). Moreover the attacker exposes themselves to detection by doing this. Alternative transmission means such as free-space optical (FSO) communication, which are far harder to intercept and jam, could provide a practical countermeasure to this attack in the future. An attacker that is willing and able to take such a blunt approach could instead jam position reporting outright, although this is expensive and risky to do at scale. This system cannot (and indeed makes no attempt to) stymie such an attack.

Spoofing of messages on the reporting link is not a realistic concern however; messages can be encrypted and signed well within the constraints of affordable, portable hardware and the establishment of keys carries little operational cost in a system of only three components. An attacker who spoofs a GPS signal to a legitimate aircraft and causes it to report a false position will be detected indirectly as the aircraft's claims will be falsified, but the ultimate cause of those false claims will not be directly revealed. Detecting GPS spoofing is an area of active research and the likelihood of feasible solutions being found is high [52].

An effective attack against TDoA systems is the use of directional antennae to break the assumption that each receiver is detecting a message broadcast at the same time. By transmitting each message to only a single receiver, the attacker can apply different time offsets to each identical copy, such that the arrival time differences are consistent with the claimed location. This task is made far more difficult with a mobile node as the attacker must track the mobile node with one antenna and alter their timings based on its position. While logistically this is easier than having to physically move the transmitter for each message, the underlying problem of locating the mobile node remains the same. In this way our system provides an additional benefit that cannot be realised with static receivers alone.

A mobile node, of course, needs power and as such its activity will be limited by energy constraints. If the mobile node simply stops while it refuels then the attacker has an open window in which to launch an attack. Some policy must be employed to overcome this, such as using a number of identical mobile nodes (albeit at greater cost).

The length of the mobile node's reporting period determines the maximum time in which an attacker is guaranteed not to have been detected. As such the selection of an appropriate reporting time is crucial in limiting how long an attacker can perpetrate a falsely-claimed track. It also determines how often an attacker can attempt to localise the mobile node using the transmission to assist with a timing attack. A short reporting period keeps the maximum 'guaranteed-undetected' period short, whilst a long reporting period gives the attacker only low-resolution location-estimation capabilities, as well as reducing transmission overhead and associated power consumption. The system is agnostic to any selected reported period and so this factor could be scaled with traffic as necessary and randomised to make the task of the attacker more difficult. The selection of this value represents an important factor in adjusting the performance characteristics of the system to suit the use-case.

2.8 Practical Considerations

Typically, the fixed node would be an existing surveillance system receiver, with the processing unit co-located, while the mobile node would be a UAV. There is clearly a requirement that the fixed node and the mobile node should have overlapping reception areas, but there are no other location restrictions on components.

Messages do of course need to be successfully received in order for the system to work. As such it is suitable only for broadcast systems in which the receiver can directly sense messages on the transmission medium. If the report is, for example, relayed en-route to the receivers then the verification results will clearly be incorrect. Similarly while the verification system does not in theory need to know the content of messages in order to be able to record their reception times, in a situation where only a small fraction of traffic is relevant to the verification system (e.g., location claims being made by a mobile device over a 4G connection along with other traffic), the system cannot determine which messages are relevant and must timestamp each one, substantially reducing efficiency.

It must also be possible to match messages using some unique identifier. The parameters of uniqueness in this context are very weak however; the identifier




Figure 2.9: Mobile receiver node architecture diagram

need only be unique for the period between collected claims being reconciled at the processing unit. A target-specific identifier such as a MAC address, callsign or results of transmitter fingerprinting and a message-specific identifier such as a sequence number, hash of message contents or even a sufficiently-precise position claim itself can be appropriate here.

The mobile node will seldom follow the defined course exactly; both due to limitations on the movement accuracy of the mobile platform itself and environment factors such as high winds blowing it off-course. Crucially, tight adherence to the prescribed course is not necessary however, as only the position at message arrival is required for the verification step. The course-following behaviour is simply to create a challenge for an attacker to predict the location of the mobile node and use that information in their attack.

2.9 Prototypes

We developed a receiver architecture for the mobile node and implemented two prototypes using inexpensive commercial-off-the-shelf (COTS) equipment. Figure 2.9 shows the architecture. There are three main detection subsystems: a Radio receiver to detect position messages from the channel, a Clock to provide a timestamp on each message arrival and a Localisation system to determine the position of the mobile node when a message is received. A central Collector is responsible for recording these values together and storing them until the next reconciliation, at which point they are passed en masse to the processing unit via a Downlink.

The Radio subsystem could be implemented with off-the-shelf message reception hardware for the protocol being observed (e.g., a Wi-Fi or ZigBee receiver) or with a customisable software-defined radio (e.g., a USRP or bladeRF unit) for increased adaptability. The requirements here are simply that messages are provided to the collector in a timely and predictable fashion in order to avoid introducing additional error and that the detection resolution of the receiver is sufficiently high that the measurement error is kept small. The Clock subsystem can be any sufficiently-precise clock available, either internally or via an external device. The clock should display minimal drift even in the presence of changing environmental factors as a result of movement (such as variations in temperature and pressure), or at least display predictable drift and a means of reporting it. As per the clock error model detailed in Section 2.6, multi-stage apparatus are a potential choice, such as a local oscillator that is periodically disciplined by a time signal from a global navigation satellite system (GNSS), such as GPS. Any localisation approach that provides sufficient availability and precision is a suitable candidate for the localisation subsystem. Use of a GNSS is the most obvious choice. The sensor subsystems need not be implemented completely independently; for example in the air-traffic management scenario a Radarcape ADS-B detector could be employed as it incorporates message decoding, high-resolution timestamping and positioning capabilities together, requiring the Collector to simply record the output.

The Collector subsystem runs on an onboard computer attached to the sensor subsystems and to the mobile platform itself, denoted as the Controller. The Controller also runs a movement planning algorithm that is responsible for orchestrating the random movement of the mobile node in conjunction with the underlying mobile platform's control systems. The Downlink can be any suitably long-range, secure connection with the fixed node, such as a common 4G modem.

2.9.1 Ground-based Prototype

Our first prototype implementation was constructed using off-the-shelf components at a cost of less than £100 (\$155). The hardware consisted of a Raspberry Pi 2

Model B acting as Controller and providing the Clock subsystem, an RTL-SDR⁴ (specifically a NooElec NESDR Mini 2 dongle) to form the Radio subsystem and an Adafruit Ultimate GPS Breakout v3 (using an MTK3339 GPS module) acting as the Localisation subsystem and assisting the onboard clock. The Raspberry Pi ran the Raspbian Wheezy Linux distribution with a v3.18 kernel and the `pps_gpio` kernel module to accept a pulse-per-second (PPS) signal. With this capability, the GPS unit not only provided accurate location data for the mobile node, but also acted as a timing source to discipline the Raspberry Pi's internal clock by being configured as a Stratum-0 source for the local Network Time Protocol (NTP) daemon. At the start of every second the PPS signal raises an interrupt to correct the clock if necessary. This corresponds to the time drift model discussed in Section 2.6, with the manufacturer of the GPS unit quoting a 10ns jitter for the PPS timing signal [53].

Capture of ADS-B messages was performed using a Linux port of the `dump1090` utility⁵ with extremely minor modifications. Upon receipt of a message the standard behaviour of the utility is to log the raw message, along with the decoded data and the value of a rolling sample count, the utility was modified to also log the system time in this case. The sample count represents an incrementing counter of each sample provided by the RTL-SDR adaptor. With the sampling frequency set at 2MHz the counter has a nominal resolution of 500ns, dependent upon the accuracy of the oscillator in the RTL-SDR adaptor and the avoidance of any lost samples.

A Python script was used to monitor the output of the GPS unit and log the position and system time on each update. Statistics were also captured from the NTP daemon to monitor the drift of the system clock against the PPS output provided by the GPS unit. No live downlink was implemented, instead messages along with timestamps and receiver locations were recorded locally and retrieved later.

⁴RTL-SDR refers to a range of DVB-T digital television receivers that use the Realtek RTL2832U demodulator chipset, which allows samples to be streamed to a host via a debug mode, creating a low-cost receive-only software-defined radio.

⁵Originally <https://github.com/antirez/dump1090>, forked to a Raspberry Pi-compatible Linux version and extended at <https://github.com/MalcolmRobb/dump1090>

2.9.2 Aerial Prototype

We constructed a second receiver prototype, as shown in Figure 2.15. It was constructed using a Radarcape device that incorporates ADS-B decoding, high-resolution timestamping from a GPS-disciplined oscillator and positioning capabilities. Messages were collected by a Raspberry Pi 2 Model B and stored locally at the receiver. The receiver was mounted on a 3D Robotics X8+ multirotor platform, as seen in Figure 2.15a. The low-tech mounting arrangement for the receiving equipment consisted simply of an off-the-shelf electronics project box bolted on struts to the bottom of the airframe. With the exception of protruding antennas the collection equipment itself, even in experimental form, is compact and weighs less than most contemporary laptops. It is substantially less heavy than the drone with its battery attached.

The Radarcape decodes ADS-B messages with an onboard FPGA decoding implementation, timestamping each message in the process, then provides a selection of high-level interfaces providing received messages in a convenient format. We made use of a JSON format accessible from a web server hosted by the device itself, which provides a listing of the most recent message for each aircraft. This service can be polled at a high enough rate to miss no position updates (2 Hz) without consuming large quantities of processing resources.

2.10 Evaluation

2.10.1 Simulation

Primary evaluation of the proposed system was conducted in a simulation, testing verification performance against various attack types and with a selection of property values employed. The use-case was taken to be verification of aircraft position claims broadcast in ADS-B messages to be received by air-traffic control stations.

The simulation modelled the operation of the static and mobile verifiers and a number of attackers. Position claim data obtained from the OpenSky Network⁶ [54]

⁶A collaborative ADS-B reception and recording initiative, intended to provide data for ATC research

were used to provide real flight tracks of legitimate aircraft. The data were captured from a receiver at the University of Oxford Department of Computer Science over a 24-hour period on the 13th June 2012. There were 392,549 position messages in total, covering 1,088 ICAO identifiers with anywhere from a single message to 4,240 messages per identifier. The locations were all treated as genuine and the original transmission times were computed from the reception times by subtracting the propagation delay between the claimed location and the receiver. An attacker was then added, selected from one of three classes. The attacker's false messages are combined with the genuine ones and the resulting dataset fed into the processing unit.

For the purposes of the simulation a static attacker is placed at a random location on the ground, with a randomly-selected flight track that they attempt to mimic. A mobile attacker is airborne, using a rotor-wing UAV, and following their own random flight path. A course deviation attacker is instantiated as an airborne aircraft with a randomly-selected course along with a randomly selected deviation from it up to a maximum of 10% alteration in heading or pitch. The attacker begins by correctly reporting their position but then moves further away from their claimed track as the simulation progresses.

The mobile node was modelled as a rotor-wing UAV and assumed to have already completed its initialisation phase and was instantiated at a distance of 2km from the fixed node, stationary at 250m above the ground. Throughout the simulation the mobile verifier changes course on a randomised interval, selecting a new direction, speed and interval until the next change.

While it is wholly achievable to construct receivers, whether fixed or mobile, with extremely accurate clocks, we consider substantial error levels in keeping with the focus on off-the-shelf equipment and potentially extreme operating conditions (e.g., airborne in very cold or hot weather). A fixed node is assumed to have a high-accuracy clock with negligible drift, while a mobile receiver experiences noticeable but still linear drift. A mobile receiver's clock is disciplined at a regular interval c_{sync} by a more accurate time synchronisation signal and returns to the correct time, with only some far smaller error in measuring the synchronisation signal

$\epsilon_{sync}^t \sim \mathcal{N}(0, \sigma_{sync}^2)$. Such a model is consistent with inexpensive GPS-disciplined oscillators that output a pulse-per-second (PPS) output. The clock drift is only in effect during the interval between clock corrections.

As such the clock drift error for a time interval between t_j and t_i is limited to the drift rate applied since the last synchronisation interval plus the synchronisation error:

$$\epsilon_{drift}^t = t \bmod c_{sync} \cdot t_{drift} + \epsilon_{sync}^{c_{sync}}$$

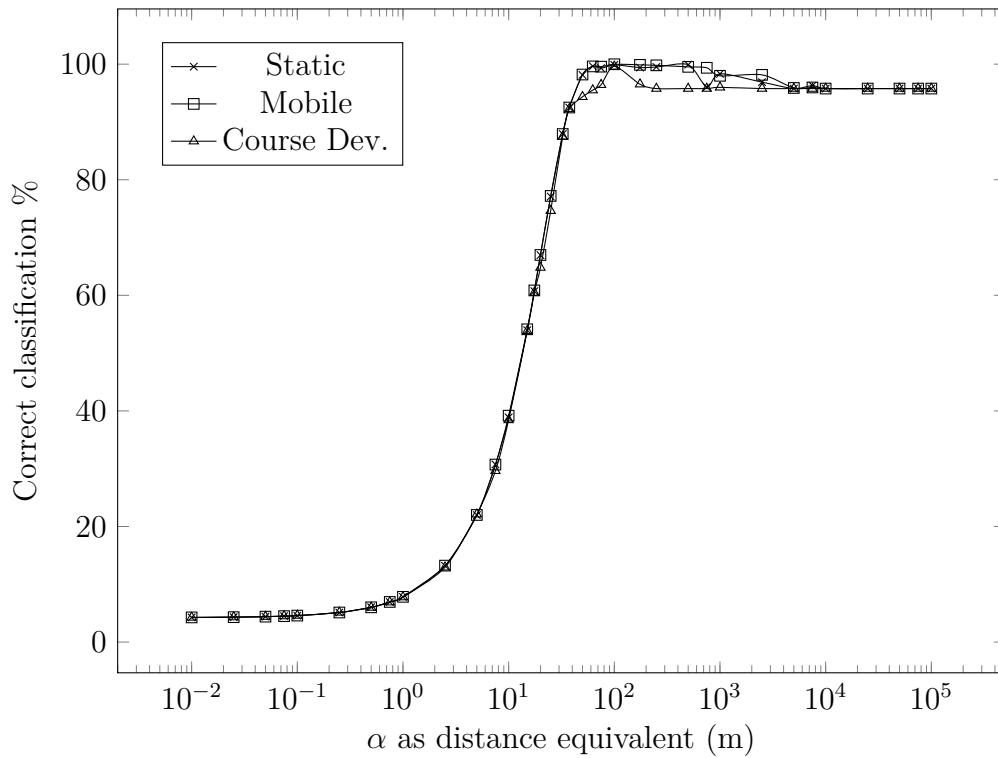
Each time the clock is disciplined, the measured drift ϵ_{drift}^t is included in a moving average $\overline{\epsilon_{drift}}$. When the mobile receiver receives a position claim message, it applies a drift compensation based on the average measured drift and the resulting value t_{rec} is recorded.

$$t_{rec} = t_{obs} - t \bmod c_{sync} \cdot \overline{\epsilon_{drift}} + \epsilon_{measure}$$


All aircraft were assumed to be in range of the receivers for the mobile node and the fixed node at all times, no reception range limits were applied. Similarly, message loss is substantial for Mode-S transmissions; the ADS-B data used here would suggest only a 7.54% detection rate for position claims. However message loss was not modelled in this simulation. This is not as strong an assumption as it may first appear; each claim is verified individually so a reduction in the number of messages received does not affect the verification itself, only the number of times that it happens.

The acceptance threshold was initially set to allow 100m of deviation; that is $\alpha = \frac{100}{c}$ (for convenience, thresholds are quoted as distance equivalents throughout this paper). All claims were treated individually; no aggregation of data at a flight level was undertaken.

Figure 2.10a shows the detection performance of the verification system in simulation against a single instance of each attacker class, as the detection threshold α is varied. For all attacker classes, peak detection occurs at α values equivalent to distances between 50m and 100m. The peak correct classification rates are



(a) Effect of changing α on detection rates for each attacker class



(b) Receiver operating characteristic (ROC) curve for detector

Figure 2.10: Classification performance of the verification system. Sub-figure 2.10a shows detection rates against the three attacker classes, while Sub-figure 2.10b shows the receiver operating characteristic for the detector.

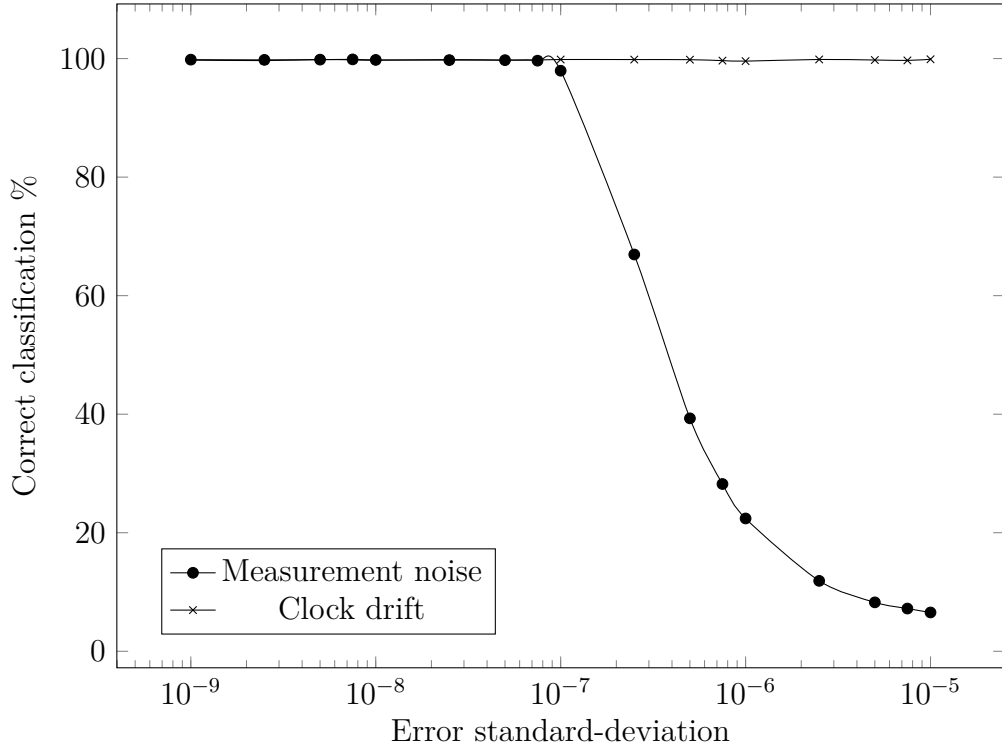


Figure 2.11: Effect of changing measurement noise ($\epsilon_{measure}$) and clock drift (t_{drift}) standard deviation on detection rates

99.8% at 100m for static attackers, 99.8% at 100m for mobile attackers and 97.3% at 75m for course-deviation attackers. The most obvious result is that a more permissive detection threshold increases the correct classification rate in almost every case until the threshold becomes unreasonably large. To understand this behaviour one must consider the scale of the air-traffic scenario; ADS-B position claims can be detected hundreds of miles away. Even a threshold on the order of several kilometres still gives an adversary comparatively little area in which to mount their attack compared to that available without the verification system. In this scenario the primary determinant for α is overcoming the sources of error in the system to avoid false negatives. In a deployment around a specific target (such as an airfield), the selection of α would be more constrained. Unsurprisingly, for this reason, performance for static and mobile attackers is near-uniform throughout. While using a mobile platform may make an attacker harder to physically stop, it does little to assist them in making false claims potentially many kilometres away appear genuine. Peak classification of course-deviation attacks is notably

lower. This can be attributed to the far smaller variation in position that appears with this attacker class; when the course-change is sufficiently small it will not breach the threshold until the new course has been maintained for some time, hence causing more position claims to be misclassified. In all cases the performance begins to diminish again as α grows to the kilometre scale and beyond. At this level of permissiveness, more false claims are treated as genuine and the higher false positive rate is the cause of the reduction in accuracy.

The relationship between the false positive rate and the true positive rate is shown in Figure 2.10b, a plot of the receiver operating characteristic (ROC) curve for each attacker class. Here the difference between the static and mobile attackers and the course deviation attacker is particularly noticeable. The potentially enormous difference between a static or mobile attacker's real position and an arbitrarily-chosen one in the sky causes them to be detected reliably until the detection threshold is made unreasonably large. The systematically smaller position difference for course deviation attackers makes avoiding detection more achievable, especially where the selected course change is small.

The effects of the primary sources of error were also explored and are visualised in Figure 2.11. The measurement error $\epsilon_{\text{measure}}$ was varied between 1ns and 10 μ s, with a single static attacker and an acceptance threshold equivalent to 100m. As would be expected, lower measurement error helps the system to make correct classifications, but only up to a point; further reductions in value below 100ns have little additional effect. Above this level performance falls away substantially, almost entirely due to a sharp increase in the false negative rate as the measurement error approaches the scale of the detection threshold. By contrast, there is almost no effect upon classification performance as the clock drift grows. This is explained by the clock error being sampled only at the start of the simulation and then assumed to drift consistently by the same amount. Under these assumptions the compensation strategy proposed in Section 2.6 can easily accommodate the drift.

2.10.2 Real-World Data Collection 1 (Ground)

The first prototype mobile node was taken on a representative-scale but ground-based collection route. Figure 2.12 shows the prototype installed in the vehicle. The equipment as tested weighed 330g, of which 130g was casing. Even with the additional mass of a battery (approximately 300g for a 10,000mAh example at time of writing) mounting the unit on a UAV platform is completely feasible. UAVs with lift capacities in excess of 1kg are widely available at low cost⁷.



Figure 2.12: Exterior view of vehicle with prototype installed

The mobile node was mounted on a car and travelled along a journey of approximately 37.5km, as shown in Figure 2.13. During a 65-minute collection period between 19:00 and 20:05 (BST) on 11th September 2015, 18,464 position claims were received and successfully decoded. A large number of non-ADS-B Mode-S replies were detected but excluded. Similarly, 679 position claims were not decoded correctly and were also omitted. In Figure 2.14 the maximal detection

⁷<http://www.dji.com/product/spreading-wings-s900/spec>

ranges are overlaid on a map of the southern United Kingdom, demonstrating that the low-cost receiver hardware and off-the-shelf antenna are capable of receiving position messages at considerable distance, even at low elevations. The maximum distance for a received position claim was 207km.




Figure 2.13: Route taken

No corresponding data were collected for the fixed node during the test, so live verification performance could not be assessed. Only equipment performance was examined.

The clock drift was observed at intervals of 16s throughout the capture journey. Clock drift was initially large; in the tens of microseconds as the equipment warmed up. After approximately 1,000 seconds, the conditions had stabilised, enabling the NTP PPS module to compensate for drift more effectively. The clock offset became more predictable; settling to single-digit microsecond values and continued as such for the remainder of the test. This suggests that clock drift could be compensated for adequately, but also highlights the needs for a mobile verifier to

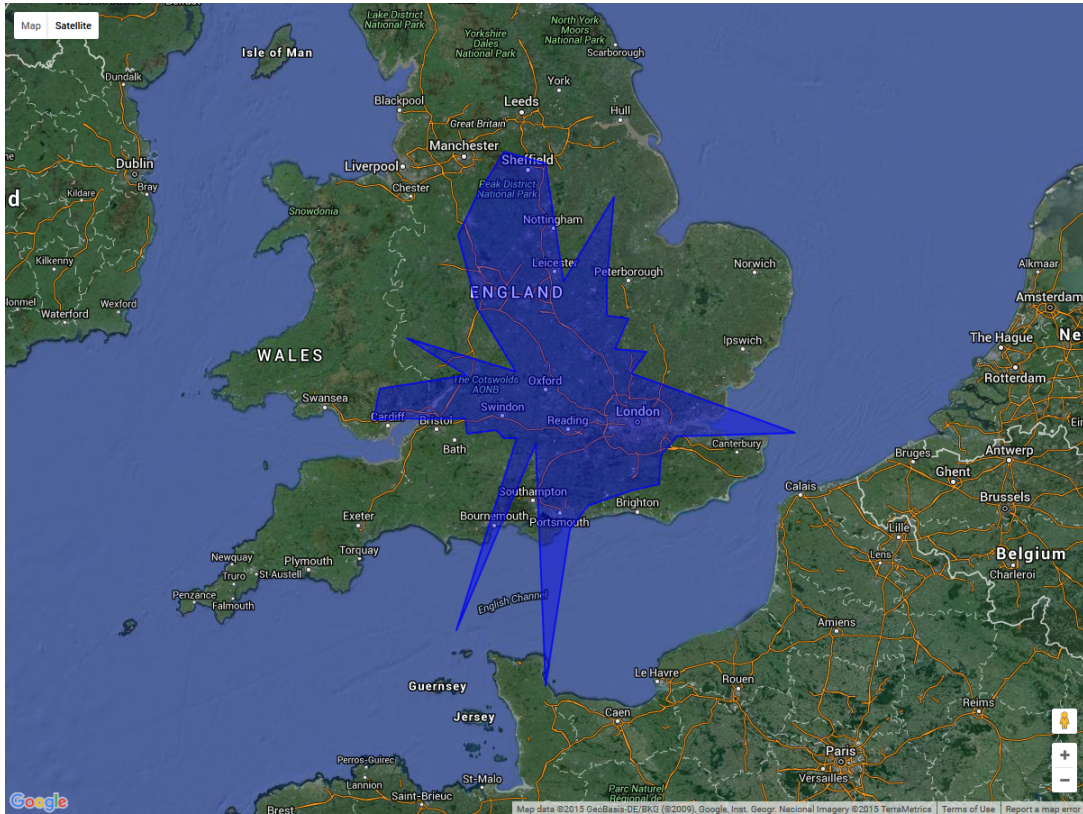


Figure 2.14: Maximal detection ranges during route

have an ‘acclimatisation’ period included in its initialisation phase to ensure its clock is stable before it is used to measure message reception times. Furthermore, the effects of pressure on the clock were not demonstrated by a ground-based journey. An airborne mobile node would need to be shielded from the effects of pressure on the clock drift, or a compensation strategy employed.


The measurement error presents the most notable problem however. The maximum stable sampling rate for the RTL2832U chipset has been observed at 2.4 million samples per second (MSPS) [55], so a maximum accuracy of 417ns could be achieved with this hardware, corresponding to approximately 125m of position inaccuracy due to measurement error. This is a lower bound however as the RTL-SDR passes samples over USB 2.0 with unpredictable delay and local timestamps are only recorded once samples have been transferred and are being processed for decoding by the dump1090 program. Unfortunately the combined effect means that the measurement error in this case is of the order of tens of microseconds

and dependent upon the scheduling of the interrupt handler on the Raspberry Pi. As Figure 2.11 suggests, this level of measurement error would severely impact accuracy. We discuss options for remedying this problem in Section 2.12 below.

2.10.3 Real-World Data Collection 2 (Airborne)



(a) 3DR X8+ drone platform, with prototype receiver underslung



(b) Prototype receiver

Figure 2.15: The prototype mobile receiver

The airborne mobile receiver was, this time, allied to a fixed receiver on the top

floor of a Department of Computer Science building. The airborne receiver was flown in a nearby area of open access land⁸. Figure 2.16 shows the mobile receiver in flight.

Data captures were conducted using both receivers simultaneously. The mobile receiver was flown in a series of random patterns by the operator, while it collected ADS-B messages and logged its position continuously. The collection took place between 08:13 and 08:56 (UTC) on 7th June 2017. A total of 165,200 messages were collected by the mobile receiver, from 237 unique aircraft. Correspondingly the fixed receiver collected 346,341 messages, from 369 unique aircraft.

Sadly, once again severe limitations with the timing accuracy were noted; within the Radarcape itself in this case. The device is intended for static use and, while highly accurate in both positioning and timekeeping in that context, once the unit is moved beyond a threshold of approximately 250 metres it registers an error state and attempts to survey its location again. Unfortunately, the handling of this error state was for reported message timings not to degrade in accuracy, but instead to default to a timing value initialised upon startup, with no bearing on the true time of reception. As such, while the collection was illustrative of the reception capabilities of a flying receiver, the flight was restricted to a hemisphere of up to 250m around the launch site; severely impacting the variation in time differences achievable.

The short-term unique identifier used to match up messages collected by each receiver was a combination of the aircraft ICAO identifier, the position values (latitude, longitude, altitude) and the time of reception to a precision of 1ms. The inclusion of a time value allows messages to be uniquely identified even when multiple messages indicate the same position. As position messages are broadcast approximately every 500ms and a message will cover the most generous reception range expectations (500 nautical miles = 926km) in about 3ms, coarse-grained timing can be used to match messages, before they are then localised using fine-grained timing. Messages that were affected by the Radarcape timing problems noted above were unable to be matched due to the wild differences in reported

⁸As the mobile receiver carries data collection equipment, it is subject to the same regulations as camera-equipped drones. For this reason, and prudent safety considerations, an area of open and empty land was selected.



Figure 2.16: The mobile receiver in flight.

time-of-arrival, however as steps had been taken to restrict flight distance, this did not affect many messages.

Message loss was a considerable factor however. Based on the earliest and latest messages collected for each aircraft by either receiver, multiplied by the ADS-B position message broadcast rate of 2Hz. There were a total of 779,172 messages broadcast within the full extent of the tracks we observed. Our collection would suggest message loss of 55.5% for the fixed receiver and 78.8% for the mobile receiver. This is above our expectations from literature, which placed loss between 40% and 50% for the time of day and distances we received over [56]. We attribute the greater message loss, particularly for the mobile receiver, to the requirements for compactness necessitating a far smaller antenna. The aerial position, even at very low altitudes proved a useful benefit over collections performed at ground-level, wherein the same equipment received as little as a tenth of the messages that the well-placed fixed receiver collected. More interestingly, the messages collected by the mobile receiver were not a strict subset of those collected by the fixed receiver, each receiver collected some messages that the other did not. There

were 8 aircraft that were only ever seen by the mobile receiver and 140 that were only ever seen by the fixed receiver.

The two message sets, each in the hundreds of thousands, ultimately resulted in only 7,247 matched messages; a joint reception rate of 0.93% of broadcast messages. Of these, a total of 49 messages were obviously spurious; displaying a time difference that corresponded to a distance greater than the total separation between the receivers (of 3.3km). These readings were attributed to equipment bugs and discarded, leaving 7,198 messages. Messages were collected from a wide area however, with the furthest position claim that was jointly received being over 250km from the fixed receiver.

2.10.4 Verification of Real Tracks

We examined the accuracy of our verification approach by applying it to the aircraft position claims collected with the airborne prototype, which were taken to be uniformly correct. With a parameter for α corresponding to 100m, we successfully verified 7,102 messages, corresponding to 98.7% of the 7,198 total messages.

Figure 2.17 shows the messages with their verification status, overlaid on a map of the area. As can be seen, the failed verifications were mostly for particular tracks, with a handful of verifications failing for otherwise fully-verified tracks. The longest unverified streak was 51 position claims, although even that aircraft did have some claims verified (indeed, every observed aircraft verified at least once). Given this, we attribute the failed verifications to errors in our system rather than incorrect reports from aircraft; specifically to smaller-scale clock errors from the receivers and the effects of GDOP as the receivers moved relative to one another. Nevertheless, as the receivers were only 3.3km apart in this test and some claims were from hundreds of kilometres away, the level of angular precision was quite substantial.

2.11 Discussion

We have presented a system that can be constructed today using widely-available components. However the capabilities of the system only look set to be enhanced




Figure 2.17: Illustration of aircraft location verification based on TDoA results collected using airborne receiver. The black diamonds are the locations of the receivers (overlapping completely at this map scale). Blue dots indicate validated claims, while orange dots denote verification failures.

by greater prevalence of cooperative reporting systems, widespread mobile data link provision and progress in mobile platform technology such as UAVs. Advances in flight duration and regulatory developments allowing autonomous operations beyond line-of-sight will enable mobile nodes to provide location verification capabilities over a wider area for a longer time. The equipment required to implement our location verification system could certainly be minimised substantially if intended for widespread deployment, allowing it to be carried by ultra-light UAVs such as those using thermal-hunting, gliding techniques [57].

2.11.1 Alternative Configurations

We have thus far described a simple configuration of the system, making use of a strictly fixed node and a single mobile node. This approach is suitable for many situations, but is by no means the only configuration. The system could also be implemented with more than two nodes, either to provide tighter verification constraints, or to enable greater coverage. The nodes could be part of a consistent

group, or form temporary verification groups from a larger set of mobile nodes, wherever coverage overlaps and the same message is detected by more than one. Similarly the system could be altered to use only mobile nodes and have no fixed node at all. In this case each mobile node would send reports to some remote processing unit to perform the verification, rather than having it co-located at the fixed node. This configuration would greatly enhance the mobility of the verification system, with only the mobile nodes' travel range and the availability of a suitable downlink limiting coverage. Alternatively a hybrid approach is possible, wherein some mobile base station such as a large road vehicle, equipped to act as both fixed node and processing unit, moves to an area and deploys the system ad-hoc to meet a temporary need.

2.11.2 Potential Applications

The air traffic management scenario modelled in our simulation represents only one use-case for the verification system presented herein. The approach is also applicable to many other scenarios as well, albeit with variations in configuration and different challenges in each case. The system as presented translates easily for other large-scale traffic management instances such as for marine traffic near a busy port. With a suitable VHF radio capable of receiving AIS messages and a long-range airborne mobile receiver (such as an ultralight fixed-wing), the system could be deployed in almost exactly the same configuration. Additionally, the near-constant altitude of marine vessels combined with a comparatively high-altitude mobile receiver substantially reduces the effect of GDOP in the vertical plane. Some other notable scenarios are explored below:

Connected vehicles The development of 'connected vehicles' is an active area of research. In this model vehicles report their status, both vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I). Judicious use of this information centrally can allow far more detailed traffic management, while drivers themselves can benefit from information about road conditions provided by vehicles further ahead (one of the

claimed advantages of ‘platooning’ approaches on busy roads, whereby vehicles form an ad-hoc network forwarding relevant information along the column). Such schemes can particularly benefit emergency vehicles; both in obtaining information about routes with lower traffic density and in warning other road users of their approach so traffic can be cleared more quickly. The advent of autonomous vehicles seems likely to only increase the usefulness of all such systems, and the strict adherence to their instructions. In all cases it is vital to detect any false position reports, or indeed any fake claims of the presence of an emergency vehicle in order to clear traffic for a selfish attacker. The fixed node in the verification system proposed here could be co-located with the V2I detection equipment, while the mobile node moves overhead. Alternatively a small number of mobile nodes could be directed reactively to provide location verification in response to unusual or suspicious activity.

Civilian UAV operations An attractive use-case is for the policing of civilian UAV operations. Proposals for structured usage of UAVs by commercial operators are often suggesting defined airspace or routes for operations. Inner city ‘drone lanes’, where vehicles can move across the city in defined airways that keep them away from pedestrians or ground vehicles seem a very distinct possibility [46]. Monitoring traffic is as important in this situation as on road networks; enabling the same reactive management, policing and safety control capabilities. Existing UAV operation proposals note requirements for collision-avoidance systems to be fitted to UAVs, broadcasting the location of the unit publicly so that others can ensure separation. The system then need only detect these position claims. Practical difficulties exist as transmission distances in purely collision-avoidance systems are shorter than those for traffic management. However the low cost of the system would permit widespread deployment of the system in the configuration described here. We discussed above the possibility of replacing the fixed node with another mobile node. If this configuration is viable then it could easily be deployed in a roving model to cover large areas of city, either in a deterministic manner to survey traffic patterns or in an unpredictable way for enforcement purposes.

Sensor network survey Security against localisation attacks is a concern in any sensor network where the position of individual sensors is crucial to the validity of the collected information [36]. If individual sensors are programmed to report the location they believe themselves to be at then the verification system herein can be used to detect if any sensor's localisation subsystem is being misled (e.g., by GPS spoofing). Again, a configuration with only mobile nodes would allow the verification to cover wide areas extremely easily.

2.12 Further Work

There are many avenues of future work that would enhance understanding of the verification system.

Firstly, the general model could be explored in greater detail. Other use-cases with more stringent constraints on broadcast range and radio propagation could be considered, along with alternative configurations. Localisation error could be included and its effects analysed. Alternative movement patterns could be modelled to include additional practical restrictions or take advantage of other effects. A random movement pattern maximises the difficulty for an attacker in predicting the mobile node's location at a given point in time. However other patterns could be employed in an attempt to improve verification, such as selecting waypoints that move areas of high geometric dilution of precision away from claimed aircraft locations, or that move them by a large amount for each subsequent message. Work to explore the performance characteristics of the system with different routes would help determine whether sacrificing some randomness could be rewarded by a greater increase in detection accuracy. A random movement pattern also ignores limitations that would be placed upon the movements of a real mobile node. For example, a UAV operating in a city or near an active runway would need to contend with prohibited airspace and meet logistical needs such as returning to its base before its fuel is depleted. Studying the effects of respecting these restrictions on the verification performance would be a great step in understanding the practical limitations of a deployment.

Similarly, modelling a more advanced attacker would also enhance the security analysis. A mobile attacker that moves randomly demonstrated very little improvement over a static one. Modelling an attacker that knows the mobile node's movement to some variable accuracy and reacts accordingly, would allow better evaluation of the security level provided by this approach. Better yet, considering a well-equipped attacker with a number of mobile transmitters or with directional transmission equipment would help explore the theoretical boundaries of our approach.

On more practical matters, tighter bounds on measurement error would overcome a critical barrier for the prototype as described here. Use of an alternative radio receiver, ideally with a direct input for a GPS-disciplined oscillator, would be the easiest and most effective approach, while remaining within the realms of commodity hardware and low cost. The Radarscape was capable of the required timing accuracy but did not respond well to movement and, as an integrated device, could not easily be modified to do so. One approach with our first prototype's hardware would be to use the sample counter for samples delivered by the RTL-SDR device, instead of the onboard clock. In this manner the variable USB transfer latency, while present, would only affect the timeliness of delivery and not the estimation of arrival time as each sample number identifies one capture period on the device, even if that sample does not actually arrive at the Collector until many milliseconds later. This approach is used to implement multilateration techniques in hobbyist air-traffic communities [58]. Only an initial calibration is then required to determine the actual time of the first sample, the Clock component is then fulfilled by the Radio receiver's onboard oscillator and a suitable mapping. This would be the simplest change with the existing hardware, but would be outperformed by a more capable Radio implementation, such as a USRP or bladeRF software-defined radio device. A device with a higher sample rate would also reduce the lower bound on timestamp resolution; if combined with a more reliable Clock implementation this could substantially decrease the measurement error. Alternatively, considering

tracks of claims instead of individual messages would allow the processing unit to filter over the noise parameter with sufficient samples.

2.13 Conclusion

In this chapter, a classic physical-layer security property, propagation delay, was applied in an extant cyber-physical system.

The use of TDoA to verify location claims already places difficult constraints on a successful attack under this threat model. The incorporation of mobility into the secure location verification, as described herein, further increases the difficulty for an attacker to falsify position claims without being detected; with the system demonstrating >97% accuracy against message injection attacks by the most complex attacker described by our threat model. It is additionally adaptable to many use-cases and by using only commodity hardware, it is both inexpensive and easily-deployed.

As with all TDoA systems, the performance is dependent upon errors being minimised or compensated; the level of timing measurement error is critical to the useful performance of the system and dictates implementation choices. Appropriate selection of a detection threshold depends upon the usage scenario and the level of accuracy provided by the implementation. As the system can be practically constructed in a variety of configurations using widely-available equipment, it appears to offer an attractive approach to enabling secure location verification in many contexts.

3

Drone Detection

Contents

3.1	Introduction	56
3.2	Motivation	56
3.3	Contribution	58
3.4	Related work	58
3.5	Background	61
3.5.1	Drones	61
3.5.2	UK Drone Flight Regulations	62
3.5.3	Received Signal Strength Indicator (RSSI)	62
3.6	Attack model	64
3.7	System model	66
3.8	Detection	67
3.8.1	Hallmarks of Drone Activity	68
3.8.2	Statistical Metrics	68
3.8.3	Drone Attack FSM	80
3.9	Evaluation	82
3.9.1	Experimental Design	83
3.9.2	Real-World Tests	85
3.9.3	Results	90
3.10	Discussion	101
3.10.1	Results	101
3.10.2	Observations	103
3.11	Future Work	104
3.12	Conclusion	105

Statement of Authorship

Much of the work included in this chapter was conducted in collaboration, primarily with Mr Simon Birnbach. The problem formulation, attack model, general system model, experimental design and data collection are all areas in which both the author and Mr Birnbach consider there to have been equal contribution by each of them. There was also assistance in the data collection by Dr Simon Eberz.

In the publication derived from this chapter, the detection tests themselves and much of the evaluation were contributed by Mr Birnbach. As such, a separate detection method designed solely by the author is described herein and evaluated anew. Other elements of the chapter mirror the published work.

3.1 Introduction

This chapter addresses a scenario that bears some similarity to that of Chapter 2; keeping track of where aerial vehicles are, albeit with drones instead of aircraft. The scenario is far newer however, with drones and their misuse arising only in the last few years. It is also one with a very different threat actor; namely the misbehaving owner of an off-the-shelf drone acting improperly, rather than a concerted effort to disrupt or endanger aviation safety. Nonetheless, both are examples of pervasive cyber-physical systems that are a part of daily life. They are also both situations in which the use of physical-layer features is crucial in building robust security systems. In this instance, the desire to use everyday equipment created tight restrictions on the physical information a security system could incorporate; relying only on widely-available received signal-strength indicator data. Nevertheless, extensive information about the state of the cyber-physical system could still be derived in order to protect the user.

3.2 Motivation

Drones have become enormously popular in recent years. Since the release of the Parrot AR.Drone in 2010, examples of the technology have exploded into the public consciousness. The use of drones, or variously unmanned aerial vehicles (UAVs) or remotely-piloted aerial systems (RPAS), is no longer restricted to military and professional domains. It has opened up to laymen as well. Self-stabilising multirotor designs that automatically hover upright have made piloting drones easily accessible for newcomers, meanwhile advances in sensor design and image processing have brought sophisticated flight-assist features to further ease the burden. Modern consumer drones can be flown using a handheld controller and a smartphone; and while safety records do vary wildly it is completely achievable for a novice to operate this type of aircraft without incident despite no prior training whatsoever.

The opportunity to take to the skies without arduous preparation has clearly appealed to many people. Industry analyst firms have reported sales of drones

increasing every year since 2014, with the global drone market valued at an estimated \$4.9 billion for 2019 [59, 60]. The US Federal Aviation Administration reported in January 2018 that there were over a million drones on their national register [61].

Sales by the market-dominant DJI Technology were reported at 375,000 units in 2017, with the majority intended for photography [62]. Indeed, photography is one of the most popular uses and the inclusion of high-quality, onboard camera equipment is a recognised selling point for a device.

However, the rise of consumer drones has brought many problems with it. The use of drones to deliver weapons, drugs and other contraband into prisons is now commonplace [63], political figures have been targeted in drone-based protests [64, 65] and the repeated appearance of a drone closed Gatwick airport for two days in December 2018 [66]. More localised abuses are often reported as well and there has been an increasing unease in the general population about privacy invasion by drones carrying high-fidelity camera equipment. As most consumer models are outfitted with cameras for live-view video during flight, drones can fly over fences to see into nearby gardens or even look into windows to observe the interior. In Seattle, a woman called the police after spotting a drone flying in front of her 26th-floor apartment window; observing her partially-undressed inside [67]. In a famous landmark case, a man from Kentucky was arrested after shooting down a drone over his property. He explained this act by saying that the drone was spying on his sunbathing daughter [68].

While systems for detecting drones are available, they are specialised and costly; intended for law enforcement use to protect secure sites and VIPs, rather than for normal citizens. By contrast, we focus on these threats to private individuals. An attacker can launch a drone from a distant location and fly it to an otherwise-inaccessible private area, such as outside a top-floor window within a walled garden, and freely observe the occupants. The occupants require a system that alerts them to the presence of an impending privacy invasion, with sufficient warning to allow them to react accordingly. The system must be easy to deploy without specialist knowledge and detect a wide variety of drones based on intrinsic properties of UAV

activity; both to be robust to the rapid pace of advancement in drone technology and to avoid spurious alarms. To be most readily-applicable, the system should avoid additional hardware requirements and take advantage of existing sensing capabilities.

3.3 Contribution

A system is presented here that detects the presence of a live-streaming UAV that is being used to invade someone's privacy in their home. We make use of readily-available commercial, off-the-shelf (COTS) hardware to measure the signal strength of the communication between the drone and its controller; developing metrics that identify signal properties inherent to drone flight. In doing so we make the following specific contributions:

- Development of a model of UAV-based privacy-invasion attacks
- Identification of statistical metrics that react to drone movement and surveillance
- Proposal of a state-model detection method using combinations of statistical tests
- Evaluation of the proposed system in extensive real-world experiments using popular consumer drones

3.4 Related work

There are a range of methods that can be used to detect drones. Research efforts have been dedicated to improving techniques for each method, although all have notable limitations and many are costly and difficult to operate. Commercial systems generally adopt hybrid methods that combine several techniques together, although these only exacerbate cost problems further.

The most traditional approach is through the use of radar; a staple of military and aircraft control for a long time. However, the design of radar systems (particularly the frequency of transmitted radio waves) must be optimised for a given aircraft size, so radars used for conventional aircraft monitoring cannot track objects as small as

drones. Instead, high-frequency radars well into the GHz range are required [69, 70] and such systems are prone to exhibit performance degradation in adverse weather conditions. More importantly, deploying any radar system is an undertaking that is out of reach of normal individuals; requiring expensive equipment, transmission licensing and trained operators.

Image and video analysis can be used to detect UAVs by using both their appearance and motion cues [71], although the growing variation in drone shapes is challenging for appearance-based approaches and methods based on motion cues struggle with similarities between drone and bird movement [72]. More generally, visual detection requires line-of-sight and is affected by light levels and visibility conditions. The use of infra-red emissions instead of, or as well as, visible light improves performance in darkness, but still suffers from confusion of drones and birds as well as reduced performance where background temperature levels are similar to drone surface temperatures [73].

LIDAR systems have showed recent promising results in accurate detection, but again require line-of-sight conditions and are a far more expensive technology than conventional cameras [74].

Acoustic detection, using microphone arrays or acoustically-augmented cameras, circumvents many of the limitations of visual methods by being effective at night or with no direct line-of-sight [72]. However, such systems suffer from acoustic variants of the same problems; namely that sound distortion in bad weather or noisy environments can limit performance. Systems also rely upon known acoustic signatures in order to detect drones and must be kept up to date. While acoustic systems can be built with off-the-shelf equipment, even comparatively cheap examples are still beyond the means of most private individuals. The low-cost acoustic array system in [75] requires an array of 24 microphones and cost \$3768 to build.

Hybrid approaches such as CSUAV [76] combine radar, acoustic arrays and video cameras to profit from the benefits of all approaches. But while they improve

the detection performance, they also increase the complexity and cost of the system substantially.

Boddu et al. [77] propose an approach that uses humans as sensors by building a collaborative smartphone app that allows users to share drone sightings. Their approach is more appropriate to target large-scale threats, and is unlikely to be helpful for the defence of a single property as no reliance can be made on the general public to provide a report.

We focus instead on methods that use radio frequency (RF) signal processing for detection. Some RF-analysis based methods use localisation to detect drones simply as flying transmitters. However, these methods generally require multiple receivers and costly, precise time-synchronisation mechanisms between the receivers [78]. As our work is intended for use by private individuals, we focus on the use of inexpensive consumer technologies, and the reuse of existing domestic communication systems, to permit easily-accessible drone detection. Academic work in this area has used protocol signatures from the drone's Wi-Fi connection and MAC address prefixes to recognise drones [79]. These methods depend upon known manufacturer lists and unencrypted communications however. They are also unable to distinguish a neighbour turning on a drone in their house from an actual privacy-invasion attack. More recent systems have used a visual stimulus at a potential target site in order to inject a controlled signal into the live-streamed video feed and correlate that against observed radio traffic. Upon detecting the injected signal in camera traffic, it can be deduced that the drone must be observing the target. These systems require greater deployment outlay than ours however and rely on specific digital video encoding behaviours in the drone [80].

A large number of commercial drone detection systems are available. Systems such as the confusingly-similar Drone Detector¹ and Drone-Detector² are comparatively accessible, but still far too specialist for domestic use. Meanwhile, systems

¹www.dronedetector.com

²www.drone-detector.com

from manufacturers such as Thales³, SAAB Group⁴ and Chess Dynamics⁵ are military equipment and completely unsuited.

Once a drone has been detected acting maliciously, some set of countermeasures may be deployed. A vast array of countermeasures have been proposed; from shotguns, through flying nets [81], live eagles [82], control-signal jamming [83], Wi-Fi de-authentication attacks [79], flight-controller exploitation [84] and GPS spoofing [83, 85], to interference with on-board gyroscopes via acoustic resonance [86]. In a domestic setting, the course of action could be as simple as automatically shutting the blinds. The relative merits of each approach warrant careful consideration, however they are beyond the scope of this work and we do not discuss countermeasures further — focusing solely on the mechanics of detection.

3.5 Background

3.5.1 Drones

The market for consumer drones is contested by many different manufacturers, but at time of writing it is dominated by only a few companies, namely DJI Innovations (75% market share), Yuneec (5%), 3D Robotics (3%) and Parrot (2%) [87].

The majority of drones provide a live video stream back to the operator; so-called “first-person view” (FPV). In all but the simplest of flights it is a great benefit to the pilot to be able to see from the perspective of the drone itself. The user can commonly connect to the drone with their smartphone, tablet or a dedicated controller to receive the video, and often record the footage in higher resolution as well. Interoperability with existing user devices has played a great role in enhancing the appeal of these consumer UAVs. While some higher-end models use proprietary communication technologies for their video downlink, most drones rely on common 2.4GHz or 5.8 GHz Wi-Fi instead. Some even use Wi-Fi for the control channel of the drone. In these cases, the drone can be controlled completely by an app.

³www.thalesgroup.com/en/squire-ground-surveillance-radar

⁴<https://saab.com/air/sensor-systems/ground-based-air-defence/giraffe-amb/>

⁵<http://www.chess-dynamics.com/auds/>

Otherwise, a separate, dedicated controller is needed and the Wi-Fi connection is only used for the live-view video. While we focus on Wi-Fi throughout this study, the approach could be applied effectively for any receivable RF technology as it depends solely on measurement of received power.

3.5.2 UK Drone Flight Regulations

Airspace regulations differ between jurisdictions in terms of the requirements placed upon drone operators. Mandatory registration is already in force in the US and similar legislation is due to come into force in the UK in November 2019. Almost all authorities mandate that minimum separation distances are observed between the drone and any surrounding persons or property. In the UK, the Civil Aviation Authority (CAA) classes drones equipped with cameras as “small unmanned surveillance aircraft” and requires that they [88]:

- fly no closer than 50m (150ft) away from any persons, buildings or vehicles that are not under the control of the operator
- stay below a 120m (400ft) flight ceiling
- stay within line-of-sight of the operator

3.5.3 Received Signal Strength Indicator (RSSI)

The received signal strength indicator (RSSI) is a value calculated by a receiver that relates to the measured power level of the received signal. Conceptually, any signal can have an RSSI calculated from it, although in practice the term is only applied in wireless communication systems — predominantly IEEE 802.11 Wi-Fi. While a relationship to the observed signal power exists, only weak inferences can be made about power levels from RSSI values. Notable limitations are:

- No standardised methodology for power measurements
- No mandatory bound on error tolerances
- No required resolution specification
- No accepted reference level

As such, RSSI measurements and calculations can vary widely among implementations. While far less can be claimed with confidence from RSSI values than from standardised power measurements, RSSI nevertheless does correlate to absolute signal power and can usually be treated as internally-consistent for a given receiver. In practice, modern Wi-Fi receivers usually provide measurements that are close to true power values. RSSI has indeed been used in distance-estimation and localisation for many years [89], although effective use requires some knowledge of the idiosyncracies of the transmitter in order to derive the best correlation of RSSI to distance. As the properties of a drone transmitter cannot be known in advance, distance-estimation approaches cannot be applied directly here. However, the broad relation of RSSI to distance can still be used, even if absolute values are not available. Of particular relevance in a drone-detection scenario however, is its use in determining whether a transmitter is in line-of-sight to a receiver and/or whether it is moving [90] [91].

Much of the value of RSSI, despite its limitations, derives from its availability. While a wealth of RF localisation knowledge can be applied in commodity wireless networks, methods generally require more sophisticated hardware. Timing and lateration methods require precise clocks and synchronisation between receivers, angulation approaches require directed antennas or similarly-precise timing. These underpinnings are not generally available in commodity hardware and so exclude the associated methods from common use. Substantial improvements have been made in recent years in localisation performance by incorporating features of channel estimation (and derived channel state information) to extract a detailed picture of the radio environment [92]. Further still, the use of measurements from several spatially-separated antennas in multiple-input-multiple-output (MIMO) transceivers has been shown to enhance accuracy [33]. In both cases however, access to the necessary measurements is far from universal in commodity systems and has shown little sign of becoming more readily available. CSI collection tools have existed in the research community since at least 2011 [93]⁶ and the advantages that CSI

⁶<http://dhalperi.github.io/linux-80211n-csitol/>

analysis can bring have been demonstrated repeatedly [94]. Yet today, only a handful of devices provide access to CSI values and that access is only through custom drivers or laborious reverse-engineering [95]. In contrast, RSSI values are made easily-available to userspace from the vast majority of Wi-Fi chipsets and as such, their widespread use continues.

3.6 Attack model

The attacker uses a radio-controlled drone, with an onboard camera and a live video feed, to invade the privacy of individuals inside a building. Specifically, the attacker uses a drone to look into a window that they would otherwise have a poor view of due to distance and/or height of it. Whether they are motivated by voyeuristic pleasure, a desire to observe private information or even simple curiosity, high-quality camera footage is the means to satisfy their aim. It is also necessary to stream live footage back to the operator throughout, in order for them to pilot the drone accurately⁷.

The attacker does not modify their drone to help achieve their goal, although they can fly it in any manner that is required. This is consistent with an opportunist who acts inappropriately with their equipment but lacks the skill, resources or premeditated will to optimise for an attack.

We model the attack as follows; visualised in Figure 3.2. The attacker does not have access to the premises and is thus positioned some launch distance d_l away from the target window. They fly the drone from their starting position towards the target window until it can establish line-of-sight (LOS) into the room and is close enough, at surveillance distance d_s , to observe a significant portion of the interior at the required quality. The quality of the footage is determined not only by the distance, but also the camera resolution and the field-of-view (FOV) that the window allows, taken as γ . The speed of the approach is limited only by the drone’s capabilities.

⁷In our experience during data collection, this requirement was surprisingly stringent. Even at small distances it quickly becomes difficult to estimate the relative positions of the drone and obstacles in 3D space from only a single viewpoint. Most relevantly, we found it particularly challenging to know whether a drone had a clear shot of a window without watching its ‘first-person view’ (FPV) feed. Indeed, not paying due attention to this feed and flying based on the pilot’s view alone led quickly to the most serious crash we encountered, wherein the drone contacted the



(a) View into window at distance. The FOV is narrow and the reflections mask view. Little of the room is visible.



(b) View into window at close-range. The FOV is wider and the reflections less dominant. The door, curtains, red walls and bed can be seen.

Figure 3.1: Examples of effect of distance on surveillance footage quality.




Figure 3.2: The attacker launches the drone at launch distance d_l and flies it to surveillance distance d_s to carry out the attack. The detector is installed in the window. The FOV γ of the window limits the area that is in LOS of the detector.

An attack is divided into three distinct phases:

Approach The drone is launched and flies from the start location towards the target window via some flight pattern. The flight pattern may be dictated by obstacles between the attacker and the target, or at the attacker's discretion in an attempt to avoid detection. At some point during the approach, line-of-sight must be established in order to permit surveillance and the drone must get close enough to have a detailed view inside the building.

Surveillance The drone establishes a good surveillance position, it hovers for a period of time and records video of the build interior. The movement of the drone in this phase is kept minimal in order to increase the quality of the recorded footage.

Escape Once the surveillance has been conducted, the drone is piloted away from the window and returns to the launch site.

While the attacker cannot alter the fundamental operation of their drone, they can vary the flight pattern and speed in an attempt to avoid detection. For example, they may try to approach as quickly as possible or as slowly as possible. They may attempt to mimic hobbyist flight or stay out of sight of the window for as long as possible during approach. The attacker is limited only by the bounds of the drone's manoeuvrability.

3.7 System model

An RF receiver is mounted in the window. It is capable of receiving from the radio communication technology used by the drone's video stream and taking received signal-strength indicator (RSSI) measurements periodically. It further has an ability to separate flows from different transceivers, using identifiers in the packet or stream data. Such a model is consistent with commodity IEEE

building wall and fell 9 feet to ground, immobilised.

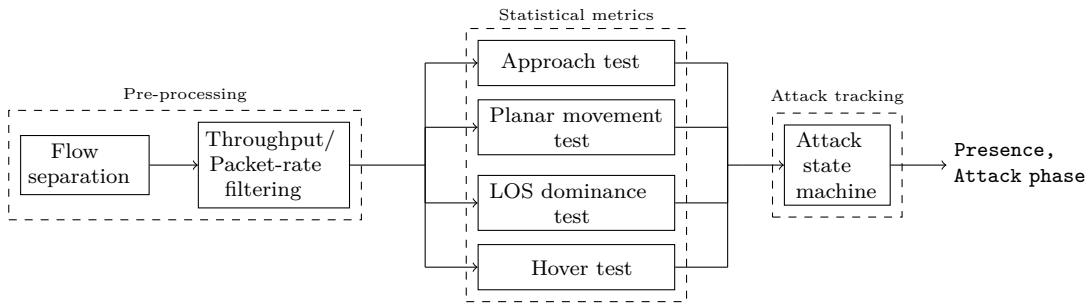


Figure 3.3: Flow diagram of the detection algorithm.

802.11 Wi-Fi receivers, but applicable to a wide range of other technologies. We focus on Wi-Fi in this work and use a receiver in Monitor mode, hopping between channels periodically. Separation of flows is performed using the MAC address in the `Transmitter Address` header field.

Figure 3.3 illustrates the overall operation. Each flow is analysed to determine whether its transmitter is a live-streaming drone conducting a privacy invasion attack. For each flow, the system computes a set of statistics over a rolling time window. It also maintains a finite state machine (FSM) modelling a drone attack. The computed statistics are used by a set of tests to determine progress through the FSM.

3.8 Detection

A drone privacy invasion attack consists of several distinct phases and displays some hallmarks of high-throughput, mobile, line-of-sight communication throughout. While many parameters of the attack can vary, such as the speed of approach, the message rate, the transmitted signal power and the length of surveillance, all attacks follow the same fundamental structure. As such we consider the detection task as effectively a signal-processing problem, in which appropriate calculations on the measured RSSI values can indicate these recognisable characteristics.

We define a set of *hallmarks* of a drone privacy-invasion attack; properties that if detected identify a drone visit distinctly from other transmitter behaviour. We then introduce a set of statistical *metrics* which we use as indicators of those properties. Finally, we describe how the appearance of hallmarks, in sequence, advances the FSM and indicates progress through the drone attack.

3.8.1 Hallmarks of Drone Activity

Live Stream The real-time requirements of live-streamed FPV footage from a drone necessitate a consistent, high message rate.

Approach To conduct a privacy-invasion attack, an attacker must reduce the drone's range to the target window substantially, such that it can see clearly inside.

Planar Movement When a drone pilot instructs the drone to move, they are effectively setting an angle-of-inclination control, such that the drone no longer hovers flat but instead pitches in one direction. A portion of the thrust generated by its rotors is now directed to one side, instead of straight down, causing the drone to accelerate in the opposite direction. The inclination of the drone then stays largely consistent as it flies, until it is again instructed to change speed or direction. As such, the antenna in a drone maintains a relatively stable orientation and the movement is largely smooth.

LOS Dominance For a drone to see into a window, it must have line-of-sight.

Hovering For surveillance, the drone must adopt a position outside a window and slow substantially; nearly to a halt, hovering in place while observing the inside of the building.

3.8.2 Statistical Metrics

The system operates over a time series of RSSI samples, computed as messages arrive. Detection makes use both of statistical functions and time-variant behaviour. The statistical functions require multiple samples in order to become valid, while the time-variant behaviour tests depend on a running series of values in order to be meaningful. As such, we make use of rolling time windows that capture all samples received over some period. A time window covers a fixed period, but can contain a variable number of values depending on the number of messages received in that period.

Function	Meaning
$count(W)$	Number of elements in window W
$duration(W)$	Time period that window W covers
$min(W)$	Minimum value in window W
$max(W)$	Maximum value in window W
$diff(W)$	Difference between latest and earliest values in window W
$ac(W, i)$	Autocorrelation of window W at lag i
$mean(W)$	Mean of values in window W
$var(W)$	Variance of values in window W
$kurt(W)$	Kurtosis of values in window W

Table 3.1: Functions used in statistical metrics

As new raw samples arrive, rolling windows are updated to contain the new sample and expired samples are dropped. Statistical functions can be evaluated over the rolling windows, such that a new result is produced for every new raw sample. The results of the statistical functions can thus be considered as live, time-variant statistics (albeit slightly delayed from real-time by the period covered by the window). The statistical functions are described in Table 3.1.

We first introduce our statistical metrics conceptually and then later describe the specific time window applied in each case. In the descriptions below, we denote a rolling window as W_X , where X is the time series of raw RSSI values over which the window operates.

The statistical metrics used by our system are as follows:

Message Rate The message rate is trivially computed from the count of messages received in some window:

$$\text{Message Rate} = \frac{count(W_X)}{duration(W_X)}$$

Total Mean Change As, for a constant transmit power⁸ in a LOS environment, the RSSI relates to the distance of the source, observing its value over a long period can provide insight into change of distance. As noted in Section 3.5.3, no absolute distance values can be derived because the system does not know the transmission power of the source. However, change relationships still hold even without absolute numbers: a 15 dB increase means the distance has reduced by 5.6×. The further away the closest point of approach is, the greater the distance that must have been travelled to achieve that multiple. As such, large changes in the RSSI, with a range-limited radio link, are more likely to be observed from transmitters that end up close to the receiver than those moving around some distance away. Figure 3.4 illustrates this.

The total mean change is simply the maximal change of any two received power values:

$$\text{Total Mean Change} = \max(W_X) - \min(W_X)$$

Trend of Mean For constant transmission power, a positive trend in the RSSI over the time window is taken as an indicator of the distance to the target decreasing. As such the value acts as a proxy for movement speed. The value is calculated over a running mean to smooth the output.

$$\text{Trend of Mean} = \text{diff}(\text{mean}(W_X))$$

Trend of Variance Over a given time window, the variance expresses the scale of the RSSI change, suitably smoothed. In an open, LOS environment this change relates primarily to movement and thus the variance correlates with speed. The trend of the variance, in turn, relates to the rate of change of speed; the acceleration.

⁸802.11 Wi-Fi contains a dynamic power scaling capability; intended primarily to ease co-existence with other networks. We expect this would cause the system to overestimate distance changes, however we did not observe dynamic power functionality in any drone we tested.

In a dynamic, NLOS environment, the variation in power could also be due to movement in the radio environment, which we test for separately.

$$\text{Trend of Variance} = \text{diff}(\text{var}(W_X))$$

Autocorrelation of Mean The autocorrelation measures the extent to which a signal correlates with itself when offset by a certain amount of time. Over a short period, a signal that autocorrelates well indicates consistent behaviour of the transmitter and environment. If, through other tests, it is determined that a signal change is due to movement, the autocorrelation can be used to indicate if that movement is smooth and sustained, or momentary and jerky. If there is no dominant change in the signal, small error variations in measurement are the main factor varying the signal and hence the autocorrelation value becomes low.

$$\text{Autocorrelation of Mean} = \sum_{i=0}^N ac(\text{diff}(W_X), i)$$

where N is the number of elements in W .

Kurtosis of Per-Message RSSI Differences A LOS environment and an NLOS environment exhibit different sample distributions. The radio environment can be tested by examining properties of the observed distributions. Doing so in such a way as to largely avoid the influence of movement, requires measuring over a short interval during which little movement could occur with the speeds involved in drone flight. The shortest measurable time with our system model is from one message reception to the next. We take the difference of RSSI between successive messages, gather differences over a short period (in order to have a reasonable sample size) and apply the kurtosis function to those values⁹.

$$\text{Kurtosis of Differences} = \text{kurt}(W_Y)$$

⁹The kurtosis is a measure of the relation between the peak and tails of a distribution, or the ‘peakedness’ of it. A higher kurtosis indicates a narrow distribution, a low kurtosis indicates a wide, flat one.

where Y is a series formed of the differences of consecutive elements in the raw RSSI measurements (i.e., $Y = (x_1 - x_0, x_2 - x_1, \dots, x_n - x_{n-1}), x_i \in X$). Samples from LOS signals are generally normally-distributed, with a kurtosis value around 0¹⁰, while samples from NLOS signals have either lower or higher values, but rarely settle at 0.




Figure 3.4: Illustration of the distance change required to vary observed RSSI values by 15dB, for several closest-approach distances. As the closest point of approach increases, the furthest point of approach increases far faster (exponentially, in fact). This is under line-of-sight conditions in free space.

The metrics just described are used in combination to test for the hallmarks of drone activity that were described above. Table 3.2 summarises how metrics are mapped to hallmarks.

The message rate is used first, to determine whether a transmitter displays sufficient signs of live-streaming to be considered a potential drone at all. The total mean change, mean trend and variance trend are used together to determine if a transmitter is moving towards or away from the receiver. The autocorrelation

¹⁰Kurtosis is sometimes described with a bias value; centred on 3.0, but we do not follow that convention here and instead consider the unbiased value centred on 0.

is used to indicate whether movement is smooth and consistent. The kurtosis of differences is used as the primary indicator of LOS dominance in the signal. At appropriate times, the mean trend, variance trend and autocorrelation are again used to determine hovering behaviour.

Various time horizons are used when computing metrics, according to what physical features are being measured. Total changes in RSSI are the maximum increases or decreases ever observed. RSSI mean, RSSI variance and message rates are smoothed over a long window W_{long} . Trends of those values, along with the autocorrelation, are then computed over the smoothed values and across a shorter window W_{short} . The LOS test takes message-to-message differences, which are subsecond, but statistics are again computed over W_{short} , such that there is a reasonable sample size.

Intermediate processing is applied for some metrics. Tests for movement, or lack thereof, are most useful when they indicate behaviour that is sustained over a period, rather than momentary satisfaction of some criteria. We use derived metrics that abstract away from specific measured values and instead represent behavioural characteristics only in terms of time. The primary derived measure is a running count of positive (or negative) values that increases while values display that trend, but resets when the trend becomes negative (positive). With knowledge of the message rate, a test can then be made for consistent behaviour over a given time window (e.g., W_{short}). This is applied for mean trend, variance trend and autocorrelation. The autocorrelation is also used separately, after a level-change detection algorithm has been applied, in order to detect drastic changes in autocorrelation brought about by the switch between movement-dominated measurements and static, noise-dominated ones.

The level-change algorithm is a common double sliding-window trigger, configured to check for a sudden fall in the level of autocorrelation, without resorting to specific, absolute values. The algorithm maintains two sliding windows $w_{early} = x[i - 2w : i - w]$ and $w_{late} = x[i - w : i]$, for some signal x and window size w , and computes an output value $y = \frac{\sum w_{late}}{\sum w_{early}}$ to detect increases in the signal or

Hallmark	Statistical Metric	Derived Test(s)	Timescale
Live Stream	Message Rate	—	W_{long}
Approach	Total Mean Change	Total increase	—
	Total Mean Change	Total decrease	—
	Trend of Mean	Running count of consistent movement towards the receiver †	W_{short}
	Trend of Variance	Running count of consistent acceleration away from the receiver †	W_{short}
Planar Movement	Autocorrelation	Running count of consistent positive autocorrelation	W_{short}
LOS Dominance	Kurtosis of Differences	Mean kurtosis value	W_{short}
Hovering	Trend of Mean	Running count of consistent lack of movement	W_{short}
	Trend of Variance	Running count of consistent lack of acceleration	W_{short}
	Autocorrelation	Magnitude of relative local decrease to average value	W_{short}

Table 3.2: All the statistical metrics indicators used by the system, with a description of how they are used to determine a hallmark of drone activity and any postprocessing performed on them. A † indicates metrics that are computed over values smoothed by a moving average of time W_{long} .

$y = \frac{\sum w_{early}}{\sum w_{late}}$ to detect decreases. The latter configuration is used here, to detect when a drone stops approaching and settles to a hover.

As little can be assumed about the speed, flight pattern or transmission strength of the drone, attempts are made to avoid the use of ‘magic’ absolute values in test of movement. Tests made on values that could be directly affected by attack variations have no threshold values; instead being strictly positive/negative comparisons that depend only on the window size selections. The kurtosis of differences value is compared against an absolute value, although it does not seem feasible to influence this within our threat model. The total mean change and message rate also have tests with a fixed value. However, these are long-term parameters for a given deployment, representing respectively the distances that are considered ‘far’ vs. ‘close’ and the average message rate for domestic devices.

These metrics can be seen plotted in Figures 3.5, 3.6 and 3.7, for respectively a drone privacy-invasion attack, a person walking with a mobile phone streaming video and a static wireless access point.




Figure 3.5: Metrics computed over time for DJI Phantom drone performing privacy-invasion attack.




Figure 3.6: Metrics computed over time for mobile device being carried around.

The drone visit in Figure 3.5 has a distinctive, flat-topped appearance to the mean. It rises rapidly, before settling flat during surveillance and then falling again




Figure 3.7: Metrics computed over time for static device.

during escape. The message rate looks very similar, largely due to the stronger signal at close range leading to lower message loss at the receiver. There are

also two noticeable plateaus, at start and end, which hint at another factor in the system. These plateaus correspond to take-off and landing, during which the camera is very close to the ground and the live video therefore changes substantially frame-to-frame for any movement. By contrast, when the drone is airborne and further away from objects in the scene, there is far less parallax between frames and so the image changes less quickly. As the video processing in this drone uses commonplace digital-compression algorithms prior to transmission, the message rate gives some indication of this increased movement in the scene. Similarly, it is notable that the message rate shows more variation during the surveillance period, when the drone again has objects close in view. We do not make use of this phenomenon here, although it has been incorporated in other work [80]. The kurtosis of message-to-message RSSI differences does exhibit large peaks, especially when near the target window (and thus seeing inside and experiencing some indoor multipath), but returns to zero often and indeed stays close to zero on average during the flight¹¹. The gradient of the mean shows the expected correlation with speed towards the receiver and then away, although the greater signal change at close-range warps the speed indication somewhat. Likewise, the gradient of the variance correlates well with the acceleration of the drone, both coming up to speed on approach or escape and also slowing to a hover once the target or launch site has been reached. The autocorrelation clearly indicates the smooth movement towards and away from the target, even accentuating the subtle change in approach speed initially. The substantial drop of the autocorrelation during hover is clear and consistent, enhanced further by the high message rate at this point allowing a comparison over many values. In this simple approach, the maximal RSSI gain is just a rebased replica of the mean, although the increase of more than 20dB means the distance must have changed tenfold during the attack. This is useful in avoiding a false alarm caused by a drone that moves and hovers in the distance, but never approaches the building closely enough to be a threat.

¹¹The spikes at start and end are aberrations due to the trace being padded with zeros fore and aft for clearer plotting here, the transition between zeros and real values leads to a brief deviation in the kurtosis value that is not present in real, unpadded traces.

By contrast to the drone attack, an approaching device inside the building (Fig. 3.6) displays less rapid rises and falls in mean and message rate that fluctuates more consistently. The kurtosis of differences, while displaying lower peaks and appearing only subtly different at first glance, never returns to zero and is in fact far higher on average. In particular, we note that when the value rises it remains high for much longer periods than in the Figure 3.5. The gradients of mean and variance show unpredictable and jerkier changes than for the drone. This is a combination of two effects of the phone being carried by hand; namely the bounce of footsteps and the rapid change in antenna orientation. This is supported by the autocorrelation which remains low for the majority of the movement and only rises at the end when the message rate drop has reduce the measurement resolution¹². The maximum RSSI change here is very similar to the drone case, highlighting the challenge of false positives due to devices at very close range inside the building.

A static device (Fig 3.7), even in a dynamic environment, is a far simpler case for the system. The mean RSSI is plotted here to the first decimal place, unlike the drone and mobile device cases which exhibited a range an order of magnitude larger. The message rate varies throughout; dominated by message loss in the dynamic, indoor environment. The kurtosis of differences is this time notably below zero, and remains so throughout. There is little gradient in the mean or variance, as indeed would be expected for a static device. As such, the autocorrelation remains low throughout, rising for small periods when the environment is still, but never remaining so for long. The maximal RSSI gain here, if it were in LOS conditions, would be consistent with a distance change of only around 5-6%.

While some metrics more obviously separate transmitter classes or behaviours than do others, each displays different behaviour in various conditions and provides distinct insight into what a transmitter is doing. Judicious use of all the metrics, in concert, enables a system to perform detection in a way that is robust against confusing edge-cases in single metrics.

¹²The initial spike in autocorrelation here, is again an aberration due to padding of the trace with zeros zeros fore and aft for clearer plotting here, the zeros lead to a brief perfect autocorrelation, but are not present in real processing

3.8.3 Drone Attack FSM

We use a state machine to model the attack phases, shown in Figure 3.8. Tests based on combinations of the metrics described above are used to advance the machine state. This enables the system to capture a notion of ordering in a drone privacy-invasion attack that would be lost if tests were applied independently. For example, the drone cannot escape until it has approached and hovering at a distance is not considered a problem, but hovering after it has approached close to the receiver constitutes surveillance and warrants an alarm.




Figure 3.8: State machine modelling the privacy-invasion attack. Solid lines indicate state transitions due to test of statistical criteria, dashed lines indicate transitions due to timeouts.

When a transmitter is first seen, a new instance of the FSM is created for it, beginning in the initialisation state `INIT.`. If it meets a minimum message rate test, that suggests it could be transmitting live video, then the transmitter becomes a potential drone and accordingly enters state `POTENTIAL`. The presence of drone-flight characteristics, along with sustained movement towards the receiver, prompts a transition to state `APPROACHING`. This state indicates a target that appears to be acting like a drone commencing a privacy invasion, but is not intended to be a final decision. If the pattern of behaviour continues and the transmitter appears to approach closely, the FSM progresses to state `APPROACHED`. Progress to this stage is considered to be confirmation that the transmitter is a drone and that it is engaging in a privacy invasion attack. An alarm output would be declared at this point. The transition to `APPROACHED` need not mean that the drone has reached its closest point of approach, rather that it has made an unusually significant

approach and that this, in combination with its previous activity, now warrants labelling as an attack. A drone may continue to approach further in this state, but an alarm will already have been raised. Once the drone settles to a hover, the **SURVEILLANCE** state is entered. The drone may approach more to improve the surveillance quality, captured in state **APPROACHMORE**, but returns to surveillance again once a stable hover resumes. When the drone finally moves away, the system indicates the attack is ending by moving to the state **ESCAPE**. Shortly after reaching this stage, the system resets back to the **INIT.** state.

Table 3.3 details the specific tests on combinations of metrics that are required to transition between states. The system is parameterised in four values: message rate r_{msg} , total mean change Δ_{total} , long time window duration w_{long} and short time window duration w_{short} .

As shown in Section 3.8.2 above, transmitters that are not drones may display some characteristics of drone flight or approach behaviour. As such it is expected that non-drone transmitters may advance the FSM to early states, but that they will not exhibit the correct pattern of behaviours and properties to reach the alarm states.

Timeouts are used to handle various inactivity conditions:

- If a transmitter is still active, but has not progressed through the state machine for a long period the state is reset back to **INIT.**. This is primarily to ensure that transmitters which trigger early transitions can be monitored from a fresh start if their behaviour changes. It also provides a recovery mechanism if a monitoring error in the attack tracking causes e.g., the **ESCAPE** transition to be missed, ensuring that the effects of the error are limited and that future visits can still be detected. Long-term metrics, such as the maximum RSSI change, are not reset in these circumstances. Thus, as the conditions for a potential approach (**APPROACHING**) are met quickly by a drone and the conditions for a confirmed approach (**APPROACHED**) are based on the persistent, total RSSI change, the attacker cannot exploit the timeout behaviour to go unnoticed.

- A much shorter timeout is applied specifically for the APPROACHMORE state, ensuring the model returns to SURVEILLANCE even in the presence of a missed transition and can then continue operating normally.
- If a transmitter is not seen at all for a long period, it is eventually removed from the monitoring, along with the FSM and all metrics. The timeout for this is very long, such that it exceeds the flight time of any drone to prevent exploitation, without incurring substantial overheads from stored transmitters.

State	Next State	Condition	Tests
INIT.	POTENTIAL	High message rate	Rate > 15 msg./s
POTENTIAL	APPROACHING	LOS	Kurt. of diffs. < 1.0
		Approaching	Positive mean trend for w_{shortS}
		Accelerating	Positive variance trend for w_{shortS}
		Planar movement	Positive autocorrelation for w_{shortS}
APPROACHING	APPROACHED	LOS	Kurt. of diffs. < 1.0
		Approaching	Positive mean trend for w_{shortS}
		Decelerating	Negative variance trend for w_{shortS}
		Planar movement	Positive autocorrelation for w_{shortS}
		Far closer	Mean has increased > 15dB total
APPROACHED	SURVEILLANCE	Hovering	No mean trend for w_{shortS}
		Hovering	No variance trend for w_{shortS}
		Hovering	Significant autocorrelation drop
SURVEILLANCE	APPROACHMORE	LOS	Kurt. of diffs. < 1.0
		Approaching	Positive mean trend for w_{shortS}
		Accelerating	Positive variance trend for w_{shortS}
		Planar movement	Positive autocorrelation for w_{shortS}
APPROACHMORE	SURVEILLANCE	Hovering	No mean trend for w_{shortS}
		Hovering	No variance trend for w_{shortS}
		Hovering	Significant autocorrelation drop
SURVEILLANCE	ESCAPE	LOS	Kurt. of diffs. < 1.0
		Retreating	Negative mean trend for w_{shortS}
		Accelerating	Positive variance trend for w_{shortS}
		Planar movement	Positive autocorrelation for w_{shortS}
		Far further	Mean has decreased > 15dB total
ESCAPE	INIT.	—	Mean trend not negative

Table 3.3: Conditions for transition between states in the detection FSM.

3.9 Evaluation

For a system with an unconstrained, everyday use-case, the appropriate evaluation seemed to be a deployment in a real-world setting, with domestic activity mimicked and repeated drone privacy-invasion attacks mounted.

We sought to discover how accurate the system was in separating drone visits from observed benign traffic, how quickly the system could detect a drone attack and what effect changes to the major parameters had upon the behaviour.

3.9.1 Experimental Design

In designing the experimental setup, the following factors were of primary interest:

Multiple receiver locations A detection system following our system model would be expected to perform best situated immediately at the target window. Testing should compare this case with receivers in other locations; both those in other windows of the building and those in free-form deployments closer to the normal distribution of devices in a building.

Static & mobile transmitters A drone detection system will likely encounter a large number of transmitters, both those situated statically in a build and those being carried by occupants. Testing should include both classes of device, in a variety of deployed locations and movement patterns.

Multiple drones The drone detection system should be reliable against a variety of drone models, with commensurate variations in transmission power, speed, manoeuvrability and communication rate.

Variations in drone flight The attacker is free to fly their drone in any manner they wish, so the detection system should be tested against a variety of approach patterns to ensure it is robust to variations in flight style and an attacker's attempts to go unnoticed. We consider a set of example flight patterns; visualised in Figure 3.9. Whilst these patterns are not exhaustive, they demonstrate some extremes of the behaviour of an attacker. A *Straight* approach is simply a direct path from the launch point to the window, it is the simplest pattern. A *Zig-zag* pattern incorporates sideways motion as well as forward motion, in order to vary the short-term approach speed. A more extreme case is the *Back-and-forth* approach,

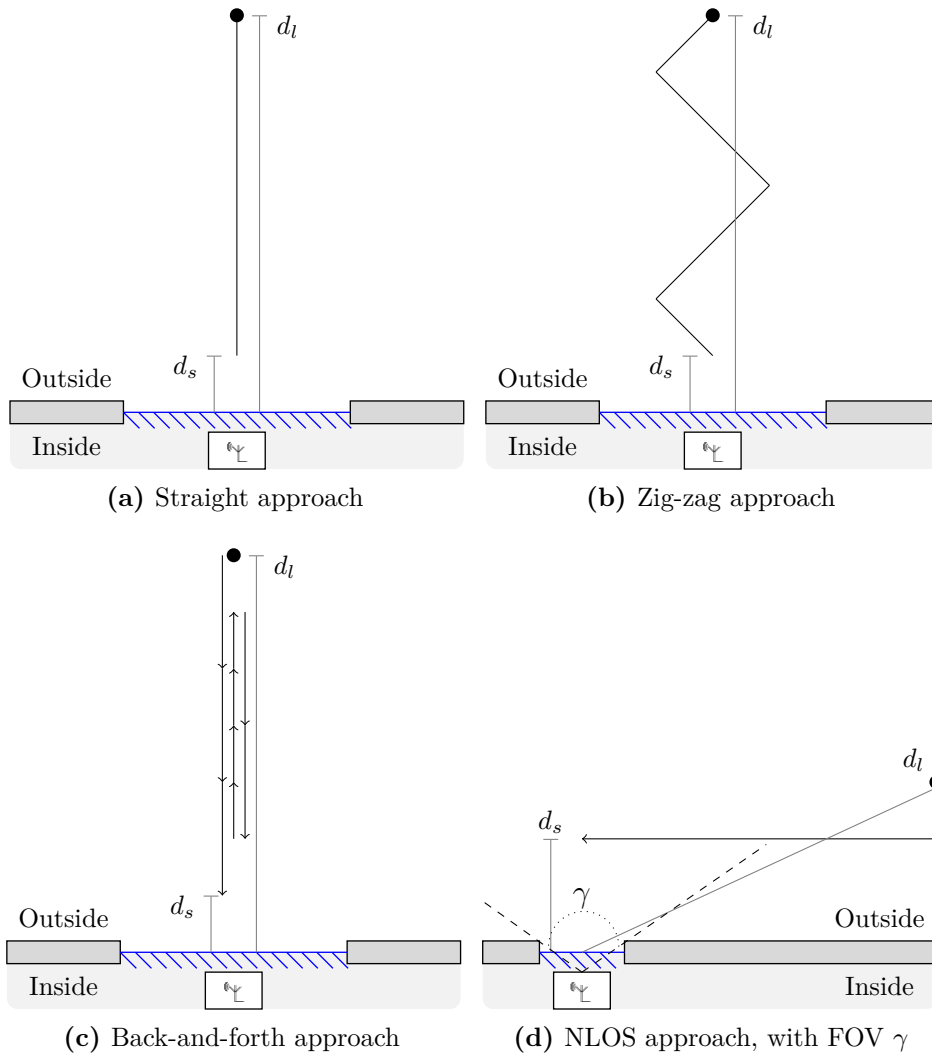


Figure 3.9: Various example approach patterns

in which some amount of the forward progress towards the detector is reversed by retreating, before approaching further again; potentially breaking the notion of a consistent approach. These three patterns vary properties of the movement alone, but LOS remains consistent throughout. The NLOS pattern also break this assumption by causing a LOS condition to emerge only shortly before the surveillance period starts. By approaching from the sides, above or below like this, the attacker can restrict the effectiveness of tests that assume LOS, whilst also giving any detection approach, whether human or machine, less opportunity to detect the drone prior to surveillance beginning.

3.9.2 Real-World Tests

The system model was realised using commercial off-the-shelf (COTS) hardware, in the form of Raspberry Pi Model A¹³ units equipped with Wi-Pi USB Wi-Fi adaptors. The Raspberry Pis ran the dumpcap utility¹⁴ to capture Wi-Fi traffic, with the Wi-Pi adaptors in monitor mode to allow traffic capture without having to associate to any network. RSSI values were provided with a 1dB resolution.

We undertook real-world experiments at a secluded¹⁵ domestic property in Devon, United Kingdom. The building was an old farmhouse with thick, stone walls that were expected to attenuate Wi-Fi signals heavily. The property was surrounded by open land, allowing a variety of approach flight paths. Figure 3.10 shows the building that was used.



(a) Rear view of the farmhouse used in the experiments. The red circle denotes the target land. Drones were launched from a point approximately 20m behind the viewpoint shown here. (b) Top view of the farmhouse and surrounding approaches, approx. 60m from the target. The \times shows the NLOS launch site, approx. 25m away.

Figure 3.10: The real-world experiment site

Receivers were deployed throughout the house, the majority were positioned in windows on the front and rear of the property over two floors. Additionally, a handful were placed away from windows, in locations that did not conform to

¹³www.raspberrypi.org

¹⁴Part of the Wireshark network protocol analyser (www.wireshark.org)

¹⁵A secluded property was selected in order to comply with CAA regulations as detailed in Section 3.5.2 and avoid disruption to others. The nearest property that was not under our control was over 350m away.

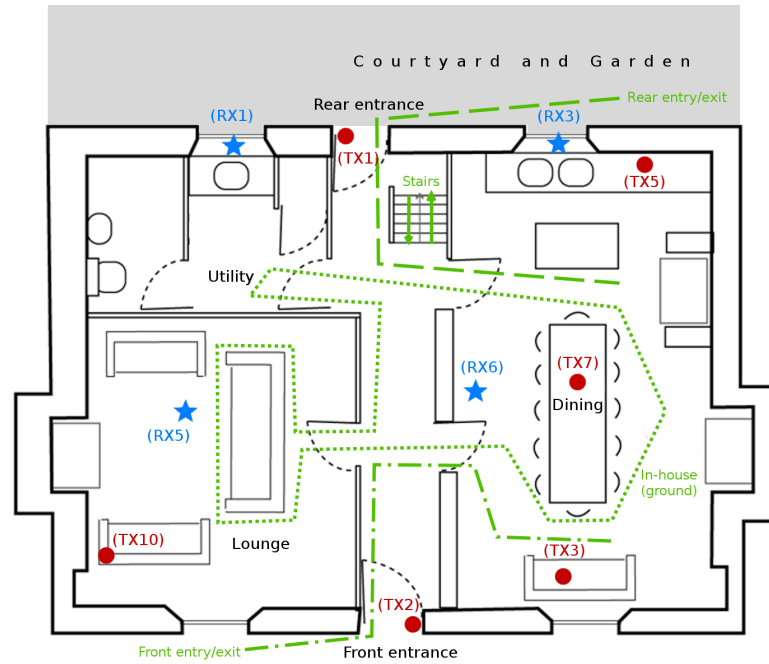
the system model, to test robustness to misplacement. In total 4 receivers were placed in windows and 4 in arbitrary locations. Maps of the locations are given in Figure 3.11. Example placement in a window is shown in Figure 3.12. A set of 10 static transmitters were placed throughout the building to emulate domestic devices. The maps in Figure 3.11 also show their locations.

Privacy attacks were conducted using two popular consumer drone models: a DJI Phantom 3 Standard and a Parrot Bebop. Both are pictured in Figure 3.13. The privacy invasion attack targeted a first-floor bedroom, marked in Figure 3.10a. Attack runs were performed for each approach pattern described in Section 3.9.1. Two versions of NLOS approach were made; one in which the drone passed over the roof of the building and descended to the window, another in which the drone panned around the side of the building at first-floor height. Five trials were made of each approach. Examples of flight traces are shown Google Earth¹⁶ plots in Figure 3.14. Notably, the GPS data became less precise close the window, as the drone no longer had a clear view of the GPS satellites.

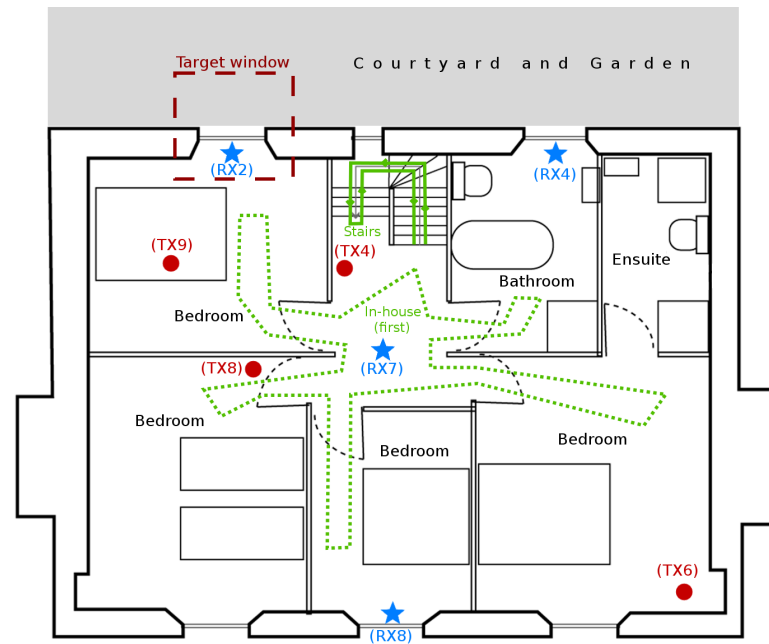
The launch distance d_l was approximately 60m for LOS approaches. For the NLOS approaches the launch distance was much shorter at around 25m as these started from the front of the house. Due to the added challenge of close-proximity flight, the operator moved to keep the drone in view during the approach. While this may be unrealistic for an attacker that cannot enter the premises, it was a prudent safety measure and was deemed only to make the situation more difficult for the detection system.

Extensive data were collected from benign transmitters throughout the test as well. The static transmitters distributed throughout the building broadcast messages at a rate of 1.5MB/s (approx. 90msgs/s) for the duration of the experiments. The message rate was similar to that of a live-streaming drone. A mobile device was also taken on a series of walking routes within and around the house, while broadcasting at the same rate. Seven routes, visiting different areas of the building were followed, shown in Figure 3.11. Routes covered entry and exit at the front and rear doors

¹⁶earth.google.com



(a) Ground floor. Receivers RX1 and RX3 were located in windows. Receivers RX5 and RX6 were resting on tables in the middle of rooms. Transmitters were distributed around the rooms. Six walking routes are shown (inc. one spanning both floors).



(b) First floor. The target window was a bedroom on the rear side of the house, with receiver RX2 located in that window. Receivers RX4 and RX8 were also positioned in windows. Receiver RX7 was ceiling-mounted in the hallway. Transmitters were distributed around the rooms. Two walking routes are shown (inc. one spanning both floors).

Figure 3.11: Receiver and transmitter placement in building. Walking routes are also shown.



(a) Receiver RX2, mounted in the target window. Pictured here from a drone. (b) Receiver RX7, ceiling-mounted inside.

Figure 3.12: Examples of receiver deployment.



(a) DJI Phantom 3 Standard

(b) Parrot Bebop

Figure 3.13: Drones used in the real-world experiments

(4 routes), roaming on the ground floor (1 route), roaming on the first floor (1 route) and walking up and down the stairs (1 route). The routes attempted to mimic normal domestic behaviour patterns that might show signs similar to drone approaches, thereby risking false positives in the detection.

The collected data were separated into trials. Each trial represented the observations of messages from one source (i.e., transmitter MAC address) by one receiver, as a time-ascending list: $T = [m_0, m_1, \dots, m_n]$ where $m_k = (\text{Time}, \text{RSSI})$. Trials were selected based on the recorded start and end timestamps of either a

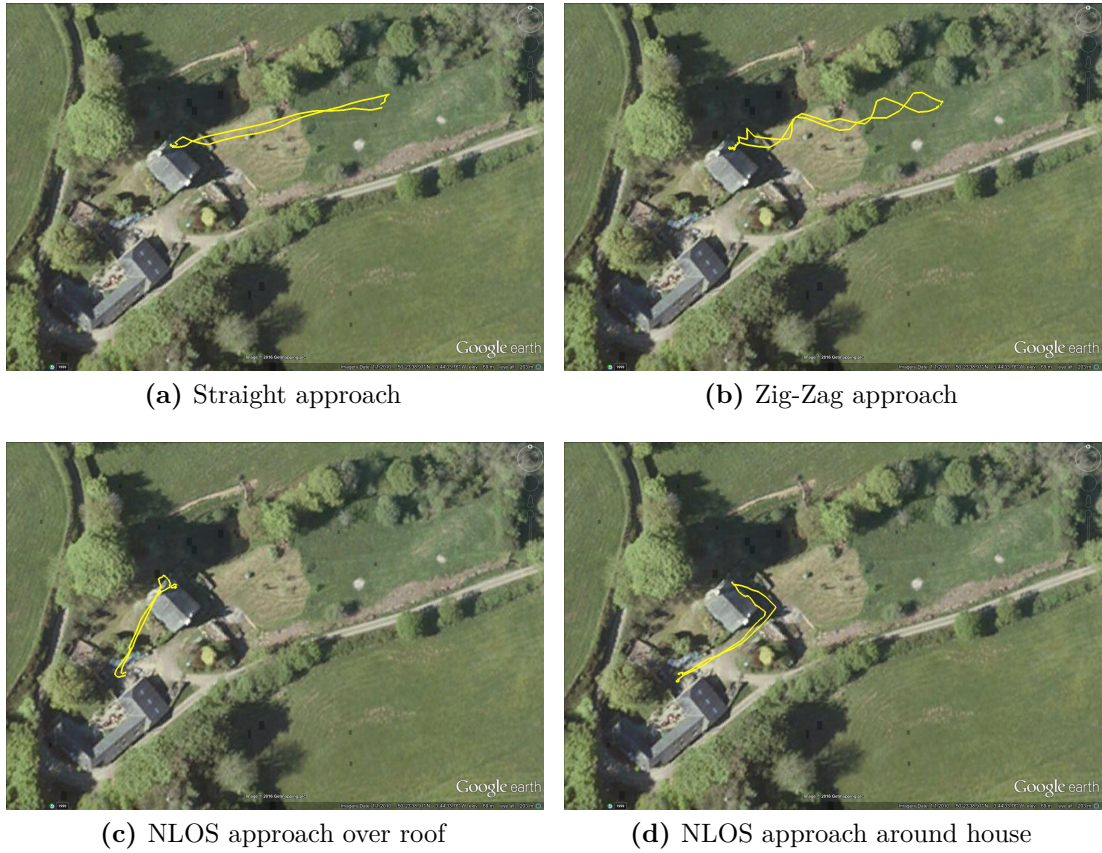


Figure 3.14: Different flight patterns during the approach

drone visit or a mobile device visit and included the drone or mobile device involved, as well as every static transmitter. As the static transmitters broadcast continually, there were long periods in which no activity was taking place. These periods were excluded on that basis that with no activity in the environment, these recordings would be the very unlikely to look like a drone visit, as opposed to periods in which there was activity in or near the building.

The system was implemented in Python, following the structure described in Section 3.8, with the exception of timeout behaviour which was omitted as trials had been individually separated. System values were set as follows: $r_{msg} = 15$ msgs./s, $\Delta_{total} = 15$ dB (approx. $12 \times$ distance decrease), $w_{long} = 10.0$ s, $w_{short} = 1.0$ s. Values for r_{msg} , w_{long} and w_{short} were selected empirically based on our observations of commonplace piloting capabilities and drone performance levels, while Δ_{total} was set based on a short survey of the size and shape of the test-site.

	N	TP	FP	TN	FN	Prec.	Rec.	Acc.
<i>RX1</i>	798	53	5	740	0	0.914	1.000	0.994
<i>RX2</i>	798	53	0	745	0	1.000	1.000	1.000
<i>RX3</i>	790	47	2	735	6	0.959	0.887	0.990
<i>RX4</i>	647	53	4	590	0	0.930	1.000	0.994
<i>RX5</i>	795	39	8	734	14	0.830	0.736	0.972
<i>RX6</i>	787	35	4	730	18	0.897	0.660	0.972
<i>RX7</i>	798	46	5	740	7	0.902	0.868	0.985
<i>RX8</i>	798	38	6	739	15	0.864	0.717	0.974
<i>Overall</i>	6211	364	34	5753	60	0.915	0.858	0.985

Table 3.4: Overall detection results, for each receiver and aggregated across all receivers. Figures here represent tests over all classes of transmitter and reflect performance in plausible in-home device deployments. Receiver RX2 was situated in the target window.

3.9.3 Results

In total 8360 trials were collected: 208 from the Bebop drone, 216 from the Phantom drone, 336 from the mobile device and 7600 from static transmitters. Not all receivers collected messages for transmitters during each trial. In 2149 cases there were no messages at all for a given transmitter-receiver pair during a trial, these were discarded to leave 6211 for testing. We considered a drone detection to have taken place iff the FSM reached a state APPROACHED. That is, if this state occurred during a drone trial then a true positive was declared, if it did not then a false negative was recorded instead. If this state occurred during a mobile device or static transmitter trial, a false positive was declared, otherwise a true negative was recorded.

The number of static-transmitter trials is far larger as there were only two drones and a single mobile device, but ten transmitters. While skewed, this distribution of sources seems consistent with common household life, in which there are many deployed devices operating continuously as well as a smaller number of mobile devices. For this reason we discuss overall results including all trials, but later dissect the results further for static versus mobile devices.

The results of processing all the captured data with the drone detection system are presented in Table 3.4. Across all trials and sensors, accuracy¹⁷ was 98.5%. For

¹⁷Accuracy is simply a measure of the number of observations classified correctly, in this case drones identified as ‘drones’ and other things identified as ‘not drones’. This is a useful overall metric, but does not provide insight into what sort of errors occur, for which precision and recall

individual receivers, peak accuracy was 100.0% for **RX2** (the receiver in the target window), dropping to 97.2% for **RX6** (on the other floor and side of the building). Indeed, the four receivers located in windows facing the drone visits (**RX1–RX4**) demonstrated higher performance than those in other locations.

Dissecting further, the recall¹⁸ was perfect for the receiver in the target window (**RX2**) and for those in the windows immediately adjacent (**RX4**) or below (**RX1**), while it fell by 11.3% (6 drone visits missed) for the receiver in the kitchen window (**RX3**) that was both below and across. This reduced performance would appear to be due to a combination of greater distance and less time with LOS, but no single factor is clearly the dominant cause. Indeed, the high performance of **RX1**, despite also being on the ground-floor, suggests that the potentially-greater multipath caused by ground reflections need not impact performance too greatly. The precision¹⁹ demonstrated the same trend of degradation, but had a lower peak value of 96.4% even in the target window. Together, these figures would suggest that in good conditions the system is very effective at detecting a drone attack but is still susceptible to reporting false alarms.

Receivers **RX5 – RX7**, that were not deployed in windows and instead placed in plausible locations for devices around the house, performed worse in both recall and precision. There did not appear to be a predictable pattern in the performance. **RX7**, on the ceiling of the first-floor landing, displayed the best recall among the freeform receivers, however **RX6**, on the ground floor in the kitchen, had the best precision (indeed yielding the second-fewest false positives overall). The most notable trend however, was that the relationship between recall and precision reversed here compared to the in-window receivers. For all four in-house receivers, the precision was higher than the recall, suggesting that the placement degraded the ability to detect drones more than the ability to reject non-drone devices. This

are useful in addition.

¹⁸Recall is a measure that just considers the performance at detecting the target, namely “what percentage of the drones were spotted”; leaving aside whether or not the system declares other things as drones as well.

¹⁹Precision is a measure that considers the performance at detecting *only* the target; namely “what percentage of things that were classified as drones, were actually drones”.

seems unsurprising given that a key assumption — the presence of line of sight — was broken in these cases and yet no assumption for non-drone device classification was. Nonetheless, by being in the middle of open areas these receivers were visited more closely by the mobile device than the in-window receivers and it is notable that the false positive counts increased very little.

Performance for Different Drones

Individual drone models differ from one another in many ways; from obvious differences such as speed, range and manoeuvrability to more subtle variations in transmission power or antenna location. From these experiments, we cannot infer anything about the impact of many of those properties. We do note a profound impact from differences in message rate.

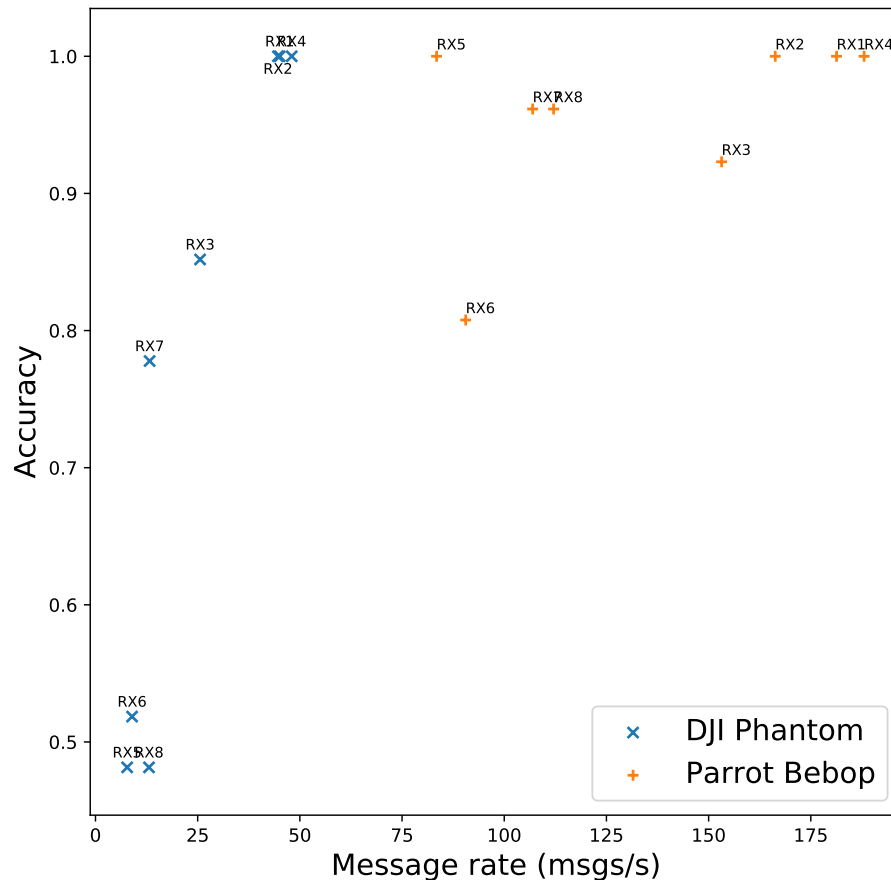


Figure 3.15: Plot of detection accuracy against message rate, for both drones. Only detection of the drones themselves is considered here, other traffic is excluded.

The Parrot Bebop drone, with an average observed message rate of 135.23 messages/s across all receivers, was detected more accurately by every receiver than the DJI Phantom 3 drone, with an average message rate of 25.79 messages/s. Indeed, the difference is stark. No receiver had less than 80.7% accuracy for the Bebop, while worst case for the Phantom was 37.0%. Both were detected perfectly by RX1, RX2 and RX4, but the Bebop was also detected every time by RX5 and RX8, despite their non-ideal, freeform positions.

The relationship between message rate and detection accuracy can be seen in Figure 3.15. The ordering of receivers by accuracy remains very similar between results for the two drones. However the accuracy does appear to be impacted strongly as the observed message rate falls. While there could be an impact from transmission power that also influences the rate of successful message reception, separate observation confirmed that the Parrot Bebop maintains a far higher message rate than the DJI Phantom and, as such, is far easier to detect with out methods.

Effect of Approach Patterns

For the three receivers that demonstrated perfect recall, there was of course no effect upon their performance from any variation in approach pattern. The other five receivers demonstrated some false negatives and the breakdown across different approach patterns is given in Figure 3.16. Three things are notable about the results.

Firstly, the straight approach accounts for the majority of the false negatives for RX3. As RX3 was verified to have clear LOS from drone footage during the attacks, it seems unlikely that this behaviour was due to LOS problems. Indeed it is unclear what the ultimate cause of was, although it is possible that the straight approach happened to cause the drone to stay in an area for which RX3 experienced substantial multipath effects, or saw the drone too briefly to identify consistent movement. Other approaches, with more complex manoeuvring and covering more space, could have allowed the receiver to measure more accurately and make a better detection.

Secondly, by contrast RX3 demonstrated no false negatives for the NLOS around pattern. This is reassuring as the drone passed close by the kitchen window (and




Figure 3.16: Normalised false negative rates for each approach pattern. Only receivers that experienced false negatives are shown.

thus RX3) on its way to the target, during this approach. This performance would support the idea that RX3 was still able to function correctly when a drone passed within its field-of-view more clearly and a closer range.

Thirdly, receivers RX6 and RX8 were predominantly affected by the NLOS over pattern. This is interesting as they were the receivers that were closest to the drone during the beginning of the NLOS over approach (one on each floor), primarily observing the initial ascent to high above the roof, while being far from the target window. While they ultimately did not detect that the drone would later approach a window, other receivers closer to the window did. If anything, given the view these receivers had, the results appear to indicate that the detector performed well in avoiding an alarm for drone behaviour that had not yet begun to look like an approach. Aside from these three observations, there appeared to be no pattern in the false negatives for other receivers, suggesting that they were the result of noise and random factors rather than systematic failings.


Detection Examples

Figure 3.17 shows example drone approaches following each flight pattern, as observed by RX2. In each subfigure, the RSSI is shown in red, after smoothing by


a $w_{long} = 10s$ moving average. The vertical lines indicate changes in state of the detection finite-state-machine. The gray line indicates the trial becoming eligible for consideration due to sufficient message rate. A cyan line marks entry into the detection process, when sustained, if brief, movement towards the receiver has been detected and the kurtosis-of-differences is not consistent with an indoor environment. This indicates a potential drone but is not considered a detection. A dark blue line, however, indicates when the approaching target has closed sufficiently to be considered ‘close’ and still does not display an indoor-environment marker (causing an FSM transition to state APPROACHED). At this point the attack is declared. A purple line indicates the start of a surveillance period, which can either be sustained or broken to approach further. A green line indicates an escape, before the system resets shortly afterwards to permit further detections.

Straight Approach Figure 3.17a shows a drone approach following a straight pattern. The pattern is clear to the eye, with a smooth initial rise in signal strength due to the closing distance, a plateau as the drone hovers in position for surveillance and a smooth escape afterwards. This, quite simple case, was initially spotted early in the approach and confirmed approximately $\frac{3}{4}$ of the way through. The surveillance phase was detected immediately that the drone settled into a hover and maintained for about 25s. As the mirror image of the approach confirmation, the escape was detected approximately $\frac{1}{4}$ of the way through its duration.


Zig-Zag Approach The zig-zag approach does show visible patterns in the RSSI levels over time, both indicating to the human eye the ‘zigging-and-zagging’, but also having the effect of smoothing the mean further. No doubt this is a function of the time window for the mean and the speed of the pattern, but the result is a less abrupt change in RSSI on approach. The system did nevertheless successfully detect the drone, indeed at an earlier point of the approach than for the straight pattern, but the gradual increase led to a small amount of bouncing between states just as surveillance was established.




(a) Straight approach, with a DJI Phantom drone.




(b) Zig-Zag approach, with a Parrot Bebop drone.



(c) Back-and-forth approach, with a DJI Phantom drone.



(d) NLOS approach around house, with a DJI Phantom drone.



(e) NLOS approach over house, with a DJI Phantom drone.

Figure 3.17: Examples of detection during various approaches; all as observed by receiver RX2. The 10s-mean RSSI is shown in red. The vertical lines indicate state changes in the detection FSM: Gray = INIT, Cyan = APPROACHING, Blue = APPROACHED, Purple = SURVEILLANCE, Green = ESCAPE).

Back-and-Forth Approach The oscillating behaviour seen in the zig-zag approach is far more pronounced in the back-and-forth approach. Indeed, the system quickly detects the drone, but does identify one of the later back-and-forth peaks as a drone attack itself, so that the system completes a full cycle before synchronising again on the correct surveillance time.

NLOS Around-House Approach In terms of the mean value, little difference is apparent in the NLOS approach versus the straight one. As the drone moves along the rear wall of the building, towards the target window, the normal detection process succeeds despite the inhibited LOS. It may be that in this case, the thick walls of the building meant that the signal in fact was behaving predominantly like a LOS signal, due to it being weakened by attenuation so much that multipath reflections were of minimal power. This is speculation, however. The drone did, nonetheless, pass very near the house on its way to RX2, which began far earlier than the detection was made, but it is reasonable to think that a detector placed on that side instead would identify the drone earlier.

NLOS Over-Roof Approach The over-roof approach was in some ways very easy to detect as the route through wide-open space above the building meant that smooth, planar movement was easy to spot; with no hesitations on the part of the pilot, or deviations to avoid obstacles. As with the around-house case, it may again be that initially the drone signal was being received with a dominant LOS component from its position high in the air. This would account for the fast early detection. However, a combination of the rapid change in LOS conditions passing over different sections of the roof, along with the slow movement of the drone as it was piloted through the fiddly blind-descent into position led to many state transitions in a short period. The system jumps rapidly between SURVEILLANCE and APPROACHMORE for periods during the final parts of the approach, with the conditions becoming more like one or the other nearly every sample. The system ultimately tracks the attack correctly and the fast transitions would not cause a problem for

	N	TP	FP	TN	FN	Prec.	Rec.	Acc.
<i>RX1</i>	756	53	0	703	0	1.000	1.000	1.000
<i>RX2</i>	756	53	0	703	0	1.000	1.000	1.000
<i>RX3</i>	748	47	0	695	6	1.000	0.887	0.992
<i>RX4</i>	622	53	0	569	0	1.000	1.000	1.000
<i>RX5</i>	753	39	0	700	14	1.000	0.736	0.981
<i>RX6</i>	745	35	0	692	18	1.000	0.660	0.976
<i>RX7</i>	756	46	1	702	7	0.979	0.868	0.989
<i>RX8</i>	756	38	0	703	15	1.000	0.717	0.980
<i>Overall</i>	5892	364	1	5467	60	0.997	0.858	0.990

Table 3.5: Detection results excluding mobile devices. Figures here represent tests of only drones and static transmitters. These two classes were the easiest to separate. Receiver RX2 was situated in the target window.

a user, who would have been warned of the approach and initial surveillance only. Nevertheless, some additional debouncing mechanism may be warranted here.

Static Devices

The good performance shown in the overall results is only accentuated further by considering static transmitters as the sole source of false positives. The full results are given in Table 3.5. Unsurprisingly, the recall was unaffected, as the exclusion of other sources did nothing to influence whether the system was successfully detecting drones. The precision results are drastically improved however, with only one false positive declared across the entire test period of 6539s (approx. 1h 48m). Static transmitters, as expected, are the easiest benign sources to exclude.

Mobile Device

As expected, mobile transmitters were far harder to separate from drone activity than static ones. Table 3.6 presents results that exclude the large number of static transmitter trials. The recall results are again unaffected, while the precision numbers now (very nearly) match those of the overall results. With the exception of the single static-transmitter false positive at RX7, the false positives were entirely contributed by the mobile devices. Nevertheless, receiver RX2 still performed flawlessly.

	N	TP	FP	TN	FN	Prec.	Rec.	Acc.
<i>RX1</i>	95	53	5	37	0	0.914	1.000	0.947
<i>RX2</i>	95	53	0	42	0	1.000	1.000	1.000
<i>RX3</i>	95	47	2	40	6	0.959	0.887	0.916
<i>RX4</i>	78	53	4	21	0	0.930	1.000	0.949
<i>RX5</i>	95	39	8	34	14	0.830	0.736	0.768
<i>RX6</i>	95	35	4	38	18	0.897	0.660	0.768
<i>RX7</i>	95	46	4	38	7	0.920	0.868	0.884
<i>RX8</i>	95	38	6	36	15	0.864	0.717	0.779
<i>Overall</i>	743	364	33	286	60	0.917	0.858	0.875

Table 3.6: Detection results considering only mobile transmitters. Figures here represent tests of only drones and walking with a mobile device, as these two classes were the hardest to separate. Receiver RX2 was situated in the target window.

Detection Speed

Across all receivers, the system detected approaching drones after a mean time of 41.2s since the trial began (thus approx. 30s after takeoff). Figure 3.18 illustrates the timings per visit, per receiver.




Figure 3.18: Detection timings by receiver (timings indicate the first transition to state APPROACHED). Individual detections are illustrated with a black dot, while the mean for each receiver is denoted by a red cross.

It is interesting that the receiver in the target window (RX2) was not the first

to detect the drone’s approach. Receivers **RX1** and **RX3**, both on the ground floor, detected visits first and were shortly followed by **RX4**. Given the arrangement of the receivers and the approach route of the drone, the performance of the **RX3** is not surprising, as it was the closest to the launch site and so would be expected to detect a total distance change first. It was surprising, however, that **RX1** made earlier detections than **RX4** as **RX4** was closer to the launch site from a top-down perspective. This could be explained by a slow ascent of the drone on approach, leading it to be closer to the ground-floor windows. Although the relative performance remained largely unchanged for the NLOS approaches as well, which would suggest that receiver arrangement may instead be the cause.


The transition to state **APPROACHED** happened comparatively late, generally halfway to three-quarters of the way through the drone’s approach (as 3.17 illustrates). This shows the drawback of the use of a multi-stage approach model to mitigate against false positives. The state **APPROACHING** was instead reached much earlier, with a mean 19.8s across all receivers, but was reached in far more false-positive trails than the **APPROACHED** state. We discuss how this limitation may be addressed in Section 3.10.1.

Effects of Window Sizes

Figure 3.19 illustrates the effect of varying the size of the long- and short-term windows w_{long} and w_{short} .

In Figure 3.19a, w_{long} is varied, while w_{short} is kept at a constant value of 1.0s. The effect on accuracy is surprisingly minor, almost negligible once the window size is above 2.5s. However, this consistent behaviour belies the trade-off of precision and recall that is occurring. At low values of w_{long} , the system is less susceptible to false positives, but at the expense of being *far* less likely to detect a drone. Between 2.5s and 10s, the system balances these two extremes, before they again diverge for larger window sizes. A value of w_{long} in the balanced region appears the best choice.

Figure 3.19b shows the effects of varying w_{short} , while w_{long} is held constant at 10.0s. There is a greater, albeit still small, change in overall accuracy here



(a) Long window (w_{long}) varied, with short window held at 1.0s. (b) Short window (w_{short}) varies, with long window held at 10.0s.

Figure 3.19: Detection performance as window sizes are varied.

compared to varying the long window, with a plateau between 1.0s and 1.5s. There is also a brief period of jointly-rising behaviour for both precision and recall as window sizes increase towards 0.75s, in contrast to the uniform opposition seen for the long window. At values of 1.0s and above, the precision begins to dominate; with diminishing returns but a linear falloff of recall. A value of w_{short} between 0.5s and 1.5s appears best here.

3.10 Discussion

We discuss here the impact of the observed results for the proposed system. We then describe below some of our observations in undertaking the work.

3.10.1 Results

The results from our experiments were encouraging. It was difficult for an attacker to avoid demonstrating the telltale hallmarks of drone flight, irrespective of approach pattern or drone model. Rejecting false positives was a greater problem than detecting drones themselves, although a sufficient message rate was important to maintain accuracy. Approach tracking was sometimes affected by rapidly-changing conditions, but this did not impact the effectiveness of the system.

The late timing of the confirmed APPROACHED state limited the warning that the system could provide to users of a drone approach (approx. 5-10s). While this may be sufficient time for an automated mechanism, such as an electronic window-blind controller, it is unlikely to be enough for an individual to choose and then enact a suitable countermeasure. However this timing is an effect of the construction of the FSM, rather than a fundamental problem of detecting physical markers. We believe there is still plenty of scope to hasten drone detection, without increasing false-positive rates, by improving the test criteria or adding intermediate states. With more substantial changes to the system model, a probabilistic, soft-decision approach could also be used and a more granular warning provided.

It was surprising to note that ground-floor receivers could still perform well, such as RX1 below the target window. While this receiver was more affected by false positives than the others in the windows, the presence of ground reflections did not cause it to miss any drone visits.

With an in-window deployment, only a little refinement to improve false positive rejection would be necessary to make the system usable in practice with acceptable usability trade-offs, at least in the tested environment. While the in-house, freeform placement clearly suffered by moving receivers away from the windows, the fact that performance was still reasonable for each receiver is impressive given the absence of LOS conditions. Further, the fact that performance was affected in different ways for different sensors suggests that, if an appropriate aggregation and weighting scheme could be devised, results could be combined with performance approaching in-window deployments.

Indeed, receivers were considered separately throughout this work, but it seems very likely that their use in concert could either improve performance, or bring additional features. Receivers in nearby windows could perhaps be tested for agreement on a drone approach before declaring an attack, with the aim of reducing false positives. Alternatively, with some knowledge of the receiver positions, a coarse-grained localisation approach could be developed.

As the detection system uses only RSS measurements, the equipment used here is but one of a number of possibilities. The majority of Wi-Fi hardware provides RSS data and many drivers make this information available. Purely software-based implementations are feasible for manufacturers and often for end-users as well. Normal user PCs, laptops, access points and IoT devices could all be expected to support the required measurement capabilities easily. User space drivers even exist for some mobile operating systems²⁰, allowing the system to be implemented as an app to turn a user’s mobile device into a drone detector. The low modification requirements of this approach greatly benefit its applicability, especially if they mean it can be incorporated into commonplace devices instead of requiring additional hardware. The fact that receivers placed in realistic, freeform locations around the building still performed reasonably suggests that the system could work well with the devices users already have, if a suitable aggregation or weighting scheme could be devised.

3.10.2 Observations

Our experimental experiences reinforced the expectation of a reliance by the pilot on the drone’s FPV video stream. At close range it is possible (and sometimes preferable) to watch the drone directly when flying, however the benefits of doing this quickly disappear as the drone operates further away and the pilot is less able to reason about its position relative to other objects purely by eye. In this case, and certainly where there is no direct line of sight at all, the attacker depends on the streamed video to pilot successfully. Even in the minimal case, conducting the privacy attack itself, the operator needs immediate feedback to ensure that the drone obtains a good view of the interior of the building — it proved easy to capture detailed footage of a window-frame by mistake. An attacker could, if their hardware permitted, disable the video stream for a period in order to avoid providing a suitable communication stream for the detection system. However, doing so near the target or when attempting surveillance, would likely jeopardise

²⁰For example, the Android user-space driver for the Alfa One NIC, with an RTL 8187 chipset

the attack. Indeed, a short surveillance distance is crucial in mounting a successful attack; at large distances a high-resolution camera or telescopic lens is required to capture a detailed image of the building interior and even then the visible area is heavily constrained by the window. Our observations suggest that values of d_s at one or two metres are realistic.

Taken together, the need for constant, live video and the requirement to approach close to a window for observation, suggest that the threat model considered here is realistic, at least for conventional house designs, and thus that the system could adequately protect the occupants from real drone privacy-invasion attacks.

3.11 Future Work

As we note, there appears to be potential for the system to work well using receivers in existing devices deployed throughout a building. Work to achieve this would be an interesting future direction, as indeed would study of how the system is affected by having mobile receivers. It would also be fascinating to investigate what could be achieved by combining information from multiple receivers; whether to improve performance or to integrate location-tracking. The possibility of using probabilistic decision criteria, instead of defined states, could make aggregation across receivers easiest, as well as helping to maximise warning time.

Further study of the behaviour in other conditions is also warranted. A long-term deployment would no doubt find more unusual false positives that could refine the detection behaviour. Testing in other environments would be particularly important, especially urban ones. Insight into whether performance is affected by enhancements of the RSSI measurement (e.g, resolution, accuracy, number of measurements per packet) would indicate whether deeper changes to hardware could bring drastic improvements. We also believe that the system could be easily applied using technologies other than Wi-Fi, but study to confirm this in practice would be helpful.

3.12 Conclusion

In this work, we developed a system to detect privacy-invasion attacks by drones based on their communication of a live video stream back to a controller. The approach uses the commonly-available RSSI measurement and a bank of statistical metrics to infer physical behaviour from patterns in the observations. Even with such limited measurement ability, appropriately-designed metrics are able to indicate an array of physical features; such as speed, movement direction, smoothness of progress and radio propagation environment.

By deriving physical features like this, a set of tests can then drive a finite-state machine that models the behaviour of a drone engaging in a privacy-invasion attack. With the system, it is possible to separate attacks from benign traffic with an accuracy of 98.5% overall, rising to 100.0% for a receiver deployed in the target window. The system is also able to track the progress of a privacy-invasion attack, providing a short warning period before surveillance is established.

We conducted real-world tests with two popular drone models and extensive sources of potential false positives. All drone visits were detected by suitably-placed receivers, irrespective of flight pattern. We found performance to be particularly sensitive to deployment location and message rate. Nevertheless, good performance could still be achieved with receivers in a natural deployment mimicking common devices in a building.

Our observations during the experiments suggest that the threat model and underlying assumptions are reasonable for opportunistic attackers. Furthermore, the results and the reliance only upon RSSI measurements suggest that the system could be incorporated, with minor refinement, into existing devices and provide drone detection without changes to user behaviour.

4

Malicious PLC Network Detection

Contents

4.1	Introduction	106
4.2	Motivation	108
4.3	Security Risk	109
4.4	Threat Model	112
4.5	Related Work	113
4.6	Background	114
4.6.1	HomePlug AV	116
4.6.2	EM Leakage	117
4.7	Designing EMPower	119
4.7.1	Detecting Radiated Emissions	120
4.8	Detector Design	123
4.8.1	Frequency Domain	124
4.8.2	Time Domain	125
4.9	Evaluation	126
4.9.1	Experimental Setup	126
4.9.2	Detection Accuracy	128
4.10	Discussion	129
4.11	Conclusion	134

4.1 Introduction

This chapter considers the use of physical-layer information from a new perspective. Unlike the work of previous chapters, in which the physical information was obtained

from an obvious effect of normal operations, in this chapter an unintended side-channel is used instead. The effect in question is the wireless leakage of a conducted signal sent over a powerline communication network.

While this unintended emission ultimately provides a signal that is very similar to the original communication, the change of transmission mode is important. Fundamentally it has different propagation properties: emitted out into three-dimensional space, instead of conveyed only along cable runs. Security systems can exploit this behaviour properties to derive additional functionality. From the perspective of security study, the security models associated with each transmission mode are different too. Wireless communication is analysed under the model of a shared medium, where only weak constraints can be placed on an attacker's capabilities. The challenges of this model have been paramount in directing the development of physical-layer security [22]. Wireline communication can also require this model of analysis (and in-home powerline communication is a prime example), but is often considered under the model of a private medium that an attacker cannot interfere with (e.g., powerline communication in vehicular deployments [96, 97]). By changing this perspective, new security considerations in the design of the system become apparent.

We consider the defensive 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 detection of power line networks being operated in the vicinity by detecting the radiated emissions. Using the radiated signal, instead of that carried within cabling, confers a number of benefits. It allows the detection system to avoid localised noise sources and disjoints in wiring, that can limit the coverage of a conductive system. It permits a more natural use of ranging and localisation techniques to locate malicious network devices in 3D space, rather than along 'cable distance'. It also, from a practical perspective, aligns well with commonplace wireless intrusion detection system deployments.

In particular we make the following contributions:

- Highlight the security threats posed by unmonitored power networks — with an example attack implementation
- Demonstrate the observability of power line communications based on their EM emissions
- Describe exemplary 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

4.2 Motivation

Power line communications (PLC) technologies have been used for over 70 years. Originally employed for long-distance measurement and control over high-voltage distribution lines, over time improved techniques have been developed that permit the use of fragmented local wiring to transmit data. Systems have since grown both more resilient to the difficult networking conditions that electrical wiring poses and also more efficient at transmitting data at high rates over these adopted media. Advances in technology, increased demand from consumers and successful standardisation initiatives such as the HomePlug Powerline Alliance, have permitted 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. The HomePlug Powerline Alliance claimed in 2016 that 220 million devices were in use worldwide [98]. 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 [99] and access from beyond a secure perimeter [33], little if any consideration is given to the openness of power networks. They present an easy target for a potential attacker.

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. The former of these has been considered in security literature such as [100][101][102], while the latter has been considered only from the perspective of electromagnetic compatibility. In particular, the radiated emissions have given rise to concerns from users of the HF band for radio communication, who strongly object to the growing broadband noise that the proliferation of power line adaptors brings [103].

4.3 Security Risk


Numerous studies of local-area power line communication 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. Stories of devices communicating between domestic houses or apartments are not unheard of [100][101] but there are clear mechanisms to resolve these problems by maintaining separate logical networks and securing them individually.

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:



(a) A prototype covert attack device.



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

Figure 4.1: A mock-up of a covert traffic capture and bulk data exfiltration attack using a maliciously-installed power line network.

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

Target	Attacker	Rate (Mbps)
L1 Private Office	L1 Social Area	33.6
	L1 Corridor	36.1
	L0 Social Area	N/A
L1 Shared Meeting Room	L1 Social Area	54.8
	L1 Corridor	45.4
	L0 Social Area	N/A
L1 Private Meeting Room	L1 Social Area	35.6
	L1 Corridor	51.7
	L0 Social Area	N/A
L0 Shared Meeting Room	L1 Social Area	N/A
	L1 Corridor	N/A
	L0 Social Area	42.1

Table 4.1: Staged attack examples. The data rate given is for application data, the PHY rate will be higher.

buildings, even if a casual observer spots the adaptor, they may well not conclude that it was put there with malicious intent.

We were able to demonstrate the effectiveness of such a straightforward attack easily; by connecting a power line adaptor at various locations within our building (a busy, recently-renovated, modern office) and attempting to establish connectivity from elsewhere. The building houses a large quantity of electronic equipment and is in constant use. We noted that the noise levels were high throughout the building and in all but the simplest deployments the adaptors indicated poor signalling conditions. Despite this we were able to communicate with the target adaptor from the majority of other power outlets on the same floor; barring only those that were near significant sources of electrical noise. There was no connectivity 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. The results of brief testing can be seen in Table 4.1.

Covert attack An attacker who is capable of only the simplest electrical tasks can conduct a more dangerous and more covert attack. Figure 4.1a shows a prototype

attack device constructed by the authors. A single-board power line networking implementation was taken from a Technomate TM-200 HP adaptor and the two power lines connected via short leads to insulation-piercing crimp connectors rated for mains voltages and wire sizes. A tiny section of CAT5 cable was connected from the adaptor's RJ45 port and into a CAT5 punchdown jack.

Figure 4.1b 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 punching down the wires 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 perpetuity¹.

4.4 Threat Model

The attacker wishes to establish connectivity to a target host or network, either for data exfiltration, traffic monitoring or as a platform for further attacks. The attacker has access to the premises, for example as an insider or a brief visitor (a courier or cleaner as in the 'evil maid' attack). In any case, the attacker wishes to establish persistent access to devices or networks in 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 4.3 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

¹The authors did not connect the prototype directly to live 230V cables, for obvious reasons, although the design was checked with a qualified electrician and tested at a nominal 5V to demonstrate continuity. A would-be attacker would be wise to add some insulation to the many exposed connections before attempting to deploy it.

office-building café or a nearby house. They can do this without establishing a rogue wireless network for which detection technologies exist 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. We further assume that the attacker is restricted to commercially-available power line hardware and is not in a position to develop their own or modify existing hardware. The HomePlug standard is complex and licensed, the authors are aware of no implementation of any kind outside member organisations of the HomePlug Powerline Alliance.

4.5 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 [26] and the ability for users to administer their devices securely [100]. More recently, practical attacks have been noted against weak implementations such that they can be co-opted into a network of the attacker's choosing [104][101]. Little consideration is given to power line networking as a vector for establishing a malicious network bridge to a target host or network, only oblique references to the possibility have been made [105].

Work on the *radiated emissions* from power line communication has been largely absent from the security literature also, and for understandable reasons; the power line channel presents such adverse signalling conditions that even legitimate nodes must manage their transmission parameters carefully in order to communicate with high data rates successfully. Data recovery even from conducted emissions is considered difficult (to say nothing of the packet encryption that would still be present afterwards) [100], to do so at distance via radiated emissions would only be more challenging. However, there is considerable work on the effects of radiated

emissions from power line communication on electromagnetic compatibility [106] and successful usage of spectrum for existing allocations [107][103]. Yet the security implications of unintentional emissions have been a rich field of study throughout the 20th century, with early military work under the codename of TEMPEST [108] entering the academic domain in the mid 1980s [109] and continuing to the present day. Side-channel attacks on displays, cabling, IO devices and embedded systems or even whole consumer devices are regularly being conducted [110][111][112].

Considerable attention has also been paid in the security community to the problem of rogue wireless access points; a topic that again has a long history [113] and is still receiving attention today both in academic circles [99] and for cyber security practitioners [114]. 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.

4.6 Background

Power line communication (PLC) systems have existed, at least in principle, since 1838, with commissioned systems used for long-distance signalling over electrical distribution networks since the 1950s [115]. These early systems used high signalling powers and operated at low data rates; being used primarily for remote equipment diagnostics and electric meter reading[116].

Work on a number of PLC systems for local-area communication began around the turn of the millennium. Some, such as X10, Universal Powerline Bus and latterly HomePlug GreenPHY were designed for robust, low-bandwidth communication for home automation, IoT and electric vehicle applications. However, the most well-known and commercially-successful emergence of the technology has been for high-bandwidth, 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 (often dubbed 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 (and indeed still are not) designed for high-frequency signalling. They are unshielded (permitting radiated emissions and susceptibility thereto) and filled with impedance mismatches, impedance variation and devices that were also never built to avoid them interfering with transmissions. As such the power lines are a very challenging environment for communication; exhibiting frequency-selective fading, plentiful multipath interference and non-linear distortion — more akin to indoor, urban wireless communication than to purpose-built, wired data networks. Equipment manufacturers [117], reviewers [118] and large-scale field tests [119] all point out that, for local-area PLC, maximum theoretical speeds are never reached in practice.

The dominant, standardised, broadband LAN PLC technologies are the HomePlug and G.hn families, which were ratified in the IEEE1901 and ITU G.9960 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 [120]. Contemporary devices at time of writing advertise operating distances in domestic settings of up to 300m [121][122] and are often deployed as Wi-Fi extenders to mitigate problems of poor coverage from the wireless network.

We concentrate in this paper on devices implementing the HomePlug family of standards in this work, in particular HomePlug AV. HomePlug AV was the second major development of the HomePlug specifications. It was introduced in 2005 to provide multimedia content-distribution capabilities beyond the capacity of the HomePlug 1.0 specification (2001, 14Mbps), which it superseded. HomePlug AV used the same 28MHz signalling bandwidth as its predecessor but achieved PHY-layer rates up to 500Mbps. It was itself superseded by the HomePlug AV2 specification in 2012 to introduce Gigabit rates using a far larger bandwidth of 84MHz, higher-order modulation modes and multiple-input multiple-output (MIMO)

transmission over earth wiring as well as live-circuit wiring [123]. Our selection of the HomePlug AV standard is due to it introducing the vast majority of functionality that persists in later specifications. However our findings are generalisable to the newer HomePlug AV2 standard and we discuss this in Section 4.10.

HomePlug AV adaptors are available in various configurations; as host NICs [124], Ethernet bridges [125] and wireless access points [126], sometimes as terminating devices and sometimes with power pass-through to permit normal use of the power outlet as well as networking [127].

4.6.1 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 [128]. The choice of OFDM in the standard’s design was to mitigate the effects of the noisy power line medium, particularly the unpredictable frequency-selective fading present in any given deployment. Each subcarrier will be affected differently by the channel noise and is managed separately. In addition, each HomePlug adaptor will experience different noise depending upon the devices and wiring near its connection point. As such, communication between HomePlug adaptors requires managing the distribution of data among the 1,155 subcarriers in either direction of each link between devices. Long-term field testing of the devices has demonstrated that noise profiles vary enormously between deployments and exhibit variation over time [119]. To successfully communicate at high bandwidths, HomePlug adaptors periodically exchange *Tone Maps* with each other, that indicate their signal-to-noise ratio (SNR) on each subcarrier of a test message sent using all the subcarriers. These Tone Maps indicate the noise profile for one direction of the link and are used to determine how many bits of data (from 0 to 10) can be modulated onto each subcarrier, in a process of adaptive bit loading. It is from this adaptive behaviour and the utilisation of a wide bandwidth that high data rates can be achieved over such a noisy channel. Where broadcast messages are required and cannot make such per-link optimisations of subcarrier usage, the standard also defines a ‘ROBO’ mode that is made resistant

to the different noise at each receiver by using time and frequency redundancy and longer error correction codes — with a corresponding reduction in data rate.

The HomePlug AV standard considers that individual logical networks may not be isolated from other devices where, for instance, there are multiple networks operating on the same power lines or where networks from other, connected parts of the power network can be overheard. 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 on the medium if a HomePlug device is connected and powered.

The lowest-level transmission structure defined in the standard is the *PHY-layer protocol data unit* (PPDU) and this is the only one that can be directly observed. The PPDU is the concatenation of a preamble, frame control data and a payload consisting of a series of OFDM symbols encapsulating the higher layers of the protocol and the ultimate data being communicated. During reception, a HomePlug AV adaptor will identify the predictable preamble, synchronise to it and then pass the frame control and payload on to the rest of the receiver processing chain.

4.6.2 EM Leakage

We have discussed above that power line wiring is not designed for high-frequency signalling and how this leads to substantial noise being present on the PLC channel. The dual of this effect is that PLC signalling is also prone to leak out of the

transmission medium, 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. In practice radiated emissions are commonly present across much of the utilised spectrum.

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 [107][106][103]. As per our threat model, we assume the attacker is restricted to commercially-available equipment and so is subject to any systems implemented therein to mitigate electromagnetic compatibility problems. In an attempt to mitigate electromagnetic compatibility issues the use of spectral masks is mandated in the HomePlug AV specification. A standard *Tone Mask*² is defined and used worldwide, while additional masks can be added to meet the needs of a particular regulatory environment or requirements of a deployment [129]. 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 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. Table 4.2 details the bands³ and Figure 4.2 shows the expected spectrum usage for a standards-compliant power line network. Usage of this spectral mask is hardcoded into power line adaptors and experimental results from emissions testing in [107] and [106] have shown that under lab conditions, emissions are consistent with these expectations. Some adaptors expose the ability to add additional spectral masks at a user-configurable level, while removing notches is not possible without substantial modification to the hardware implementation of the device.

²A blanket mask, not to be confused with the Tone Maps distributed between individual adaptors.

³Allocations are given as the widest band coverage across all three IARU international regions.

Band	IARU		HomePlug / IEEE1901	
	From	To	From	To
AM*	—	—	—	1.71
AM-160 metres*	—	—	1.71	1.8
160 metres	1.8	2	1.8	2
80 metres	3.5	4	3.5	4
60 metres	5.3515	5.3665	5.33	5.407
40 metres	7	7.3	7	7.3
30 metres	10.1	10.15	10.1	10.15
20 metres	14	14.35	14	14.35
17 metres	18.068	18.168	18.068	18.168
15 metres	21	21.45	21	21.45
12 metres	24.89	24.99	24.89	24.99
10 metres	28	29.7	28	—

Table 4.2: IARU amateur band allocations and HomePlug/IEEE1901 spectral mask frequencies. Entries marked by * are blanket masks in the relevant standards not relating directly to a single IARU allocation. Values in MHz.

While the spectral usage of HomePlug AV devices has some degree of configurability, other element of the standard exhibit far less flexibility; being necessary for correct operation of the system. The primary example is the preamble waveform that precedes every data packet. As is common in random-access networks, a preamble is transmitted to allow a receiver to detection an incoming message, synchronise in time and estimate the channel distortion conditions. This preamble is consistent across all HomePlug AV networks and must be reproduced correctly in order for compliant devices to operate correctly. In line with our threat model, an attacker’s equipment must transmit this preamble waveform before any message and thus exposes a reliably-detectable trait.

4.7 Designing EMPower

In this section we document our observations of the radiated emissions from power line communication and then describe our methods for detecting networks using the phenomenon.





Figure 4.2: Spectral usage with standard mask defined in HomePlug and IEEE1901 standards.

4.7.1 Detecting Radiated Emissions

Figure 4.3a shows the distinctive EM emissions of a power line 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 antenna⁴, 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. Contrasting Figure 4.3a against Figure 4.2 shows the comparative weakness of lower-frequency signals and particular drops around 11MHz, from 14-18MHz and around 22.5MHz.

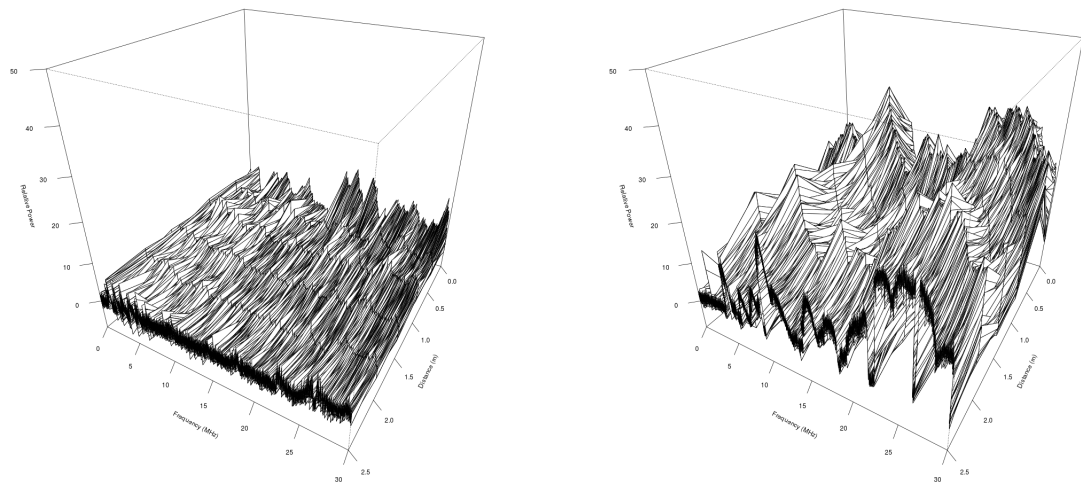
Indeed the wiring to which the power line adaptors are connected greatly affects the radiated emissions. To demonstrate this we conducted a test with two power line adaptors communicating across a 4-plug power strip. The power strip was connected

⁴At the comparatively low frequencies used by power line adaptors, resonant antennas have large dimensions. The popular G5RV dipole antenna used for HF voice communication is 30m in length. Any convenient antenna will be electrically small and as such we selected a unit that was on-hand.



(a) Observed emissions from TP-Link TL-PA511 power line adaptor. (b) Smoothed comparative emissions spectra from four power line adaptors.

Figure 4.3: Observed radiated emissions from power line adaptors.



(a) Adaptors in power strip connected to distant socket. (b) Adaptors in power strip connected to wall RCD-protected socket.

Figure 4.4: Surface plots of relative signal power against the ambient background.

to a distant socket that was fitted with an RCD protector, which severely impedes the power line signal from propagating across it due to the coils used to detect earth leakage⁵, leaving only the power strip and its cable as radiating medium. The power

⁵Precisely what legitimate power line network users are advised not to do due to the negative

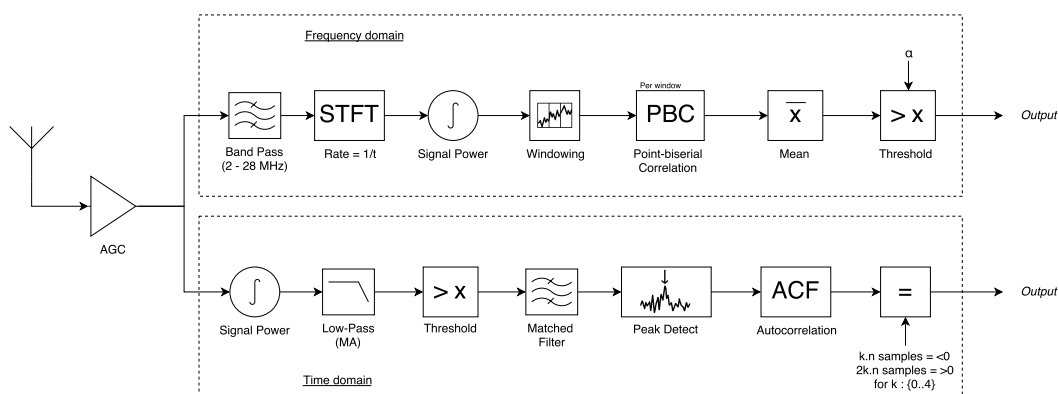


Figure 4.5: Block diagram of the system, showing the full processing chain for frequency- and time-domain analyses.

line adaptors were connected across the power strip with a Raspberry Pi and laptop running the `iperf` benchmark utility across the connection. The receiving antenna was arranged on the opposite side of the power strip, furthest from the cable. The same arrangement was then connected to a wall socket with no RCD protection, such that the signals could propagate into the surrounding ring main sections.

Figure 4.4 shows the maximal spectral power of the radiated emissions in each arrangement, with the background levels removed. In the first scenario, the radiated emissions from the power strip are all but absent. In Figure 4.4a, some peaks can be seen appearing at high frequencies when at zero range, but these are small and quickly disappear. By comparison Figure 4.4b shows the far stronger emissions that occur when the adaptors are connected to a larger circuit. Even a single room presents enough lighting, power distribution and switch wiring to present many convenient antennas, not to mention its connections to the rest of the building. The presence of certain types of common devices in, 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 [132][107]. Indeed, if the emissions are measured with a receiver situated inside the building, the unintentional antennas can conceivably surround it on all sides.

effects of the RCD upon the transmission [130][131]

While the emissions from an adaptor vary by environment, they are broadly consistent between adaptors. Figure 4.3b shows the emissions compared between four different adaptors. The TP-Link TL-PA511, TP-Link TL-WPA281 and Technomate TM-200 HP are all HomePlug AV adaptors, while the Sunvision SVW1000 is a HomePlug AV2 adaptor. To make any distinctions clearer, the spectra have been smoothed. While some amplitude level variations are visible, most notably from the TM-200 HP, there is very little difference in the pattern of emissions across adaptors⁶.

4.8 Detector Design

EMPower performs analyses of the received signal in the frequency and time domains. With the radiation properties of a given building hard to predict, examining spectral content can provide useful information even with considerable signal loss. Where the signal is received clearly enough, time domain analysis permits insight into the protocol taking place, as well as just the signalling.

Figure 4.5 shows the structure of the system. A received signal is passed through an automatic gain control circuit with a long time-constant, to provide a consistent level for processing irrespective of the received signal amplitude. The values are then passed to each processing chain for analysis; these are described in detail below in Sections 4.8.1 and 4.8.2.

The frequency-domain analysis operates on a short time period T , using values collected across that period, while the time-domain analysis runs continuously on the received signal. The time-domain results can be used directly, or combined over the same T period to obtain synchronous results. We use the latter method here for conceptual consistency. Each analysis computes a score indicating the presence or absence of a power line network. The processing chains shown in

⁶It is interesting that the TM-200 HP shows a greater signal power in a number of places as we believe it implements the optional Amplitude Map feature from the HomePlug AV specification, and so in some circumstances it demonstrates quieter emissions instead. The Amplitude Map feature allows an adaptor to alter its transmission power on any or all the subcarriers, either when mandated or simply when conditions permit. It was the only adaptor in our test set that varied its transmission power to suit conditions.

Figure 4.5 use those scores internally to output a straightforward Yes/No detection, although the raw scores can also be provided.

4.8.1 Frequency Domain

The frequency-domain method detects the presence of the distinctive Tone Mask present in HomePlug AV adaptors. The approach filters to the HomePlug AV band, computes the short-term Fourier transform (STFT) of the signal at regular intervals and then calculates the signal power. For 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 frame rate in Hz. The detector maintains a set of the maximum observed power values in each frequency bin across the observation period T . For each STFT output, the maximum value is updated if a larger one has been measured. After T has elapsed, these maxima are passed onwards through the detection chain.

The full bandwidth is first split into windows, to combat the effects of highly-variable radiation across the full bandwidth. The task of the detector is to ascertain the presence of the Tone Mask, so local variation is of far more importance than the total variation across the band. The window size is taken as the smallest size for which the Tone Mask has a change of amplitude in every window. The Tone Mask is simplified to 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 a correlation measure specifically designed for comparing continuous values against binary classifications; here the signal powers against the two expected signal levels in the Tone 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.

4.8.2 Time Domain

The time-domain method makes use of the preamble section at the start of every PPDU transmitted over a HomePlug AV network. The preamble is designed to have a reliable structure and as such, is likely to be detectable even in adverse conditions. EMPower thus performs a similar process to that deployed in a normal receiver, adapted to the conditions of radiated emissions. The power of the time-domain signal is calculated and then passed through a moving average low-pass filter to reduce noise from the amplifier and receiving radio. The moving average is short relative to the preamble length, such that it does not have a large impact on the structure of the signal. The signal is then thresholded such that baseline noise is excluded. Where the signal exceeds the threshold it is passed to a finite impulse response (FIR) filter. The FIR filter has an impulse response that matches a time-reversed copy of the preamble, from a pre-loaded template, and as such acts as a matched filter. A matched filter is optimal in separating a known signal from white noise, and serves to make preambles far more distinct from the background, and therefore easier to classify. 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 for each peak a section of the signal the length of a preamble is passed through the autocorrelation function. The preamble structure has either nine or ten⁷ repetitions of equal, known length. The result of autocorrelation is tested at these known intervals. 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.

⁷Depending on PPDU type

4.9 Evaluation

We evaluate our approaches for their capability of detecting power line networks in a range of deployments within a real-world scenario.

4.9.1 Experimental Setup

A staged attack was conducted by placing a power line adaptor in a series of locations within a normal, shared office building. The locations were selected as all those described in Section 4.3 for which a power line connection could be established. The locations are listed in Table 4.3 and can be seen in Figure 4.6, as can the static location of the EM monitoring system. The adaptor (a TP-Link TL-PA511) was placed in a power socket and a Raspberry Pi computer connected via the Ethernet port. The attacker then inserted another power line adaptor (a TP-Link TL-WPA281) at publicly-accessible locations on the same floor of the building. Tests of EM emissions were run with the power line adaptors switched off to leave only *Background*, then 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 more adequately 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, with a UBX daughterboard permitting it to tune to the low frequencies required. As in Section 4.7.1, we made use of an electrically-short, unmatched antenna, with signals being pre-amplified before entering the USRP by means of a low-noise amplifier powered by a bench supply. 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 4.8.

	#	Description
Target	1	L1 Private Office
	2	L1 Shared Meeting Room
	3	L1 Private Meeting Room
	4	L0 Shared Meeting Room
Attacker	A	L1 Social Area
	B	L1 Corridor
	C	L0 Social Area

Table 4.3: Target and attacker locations. These correspond to those given in Table 4.1 above.

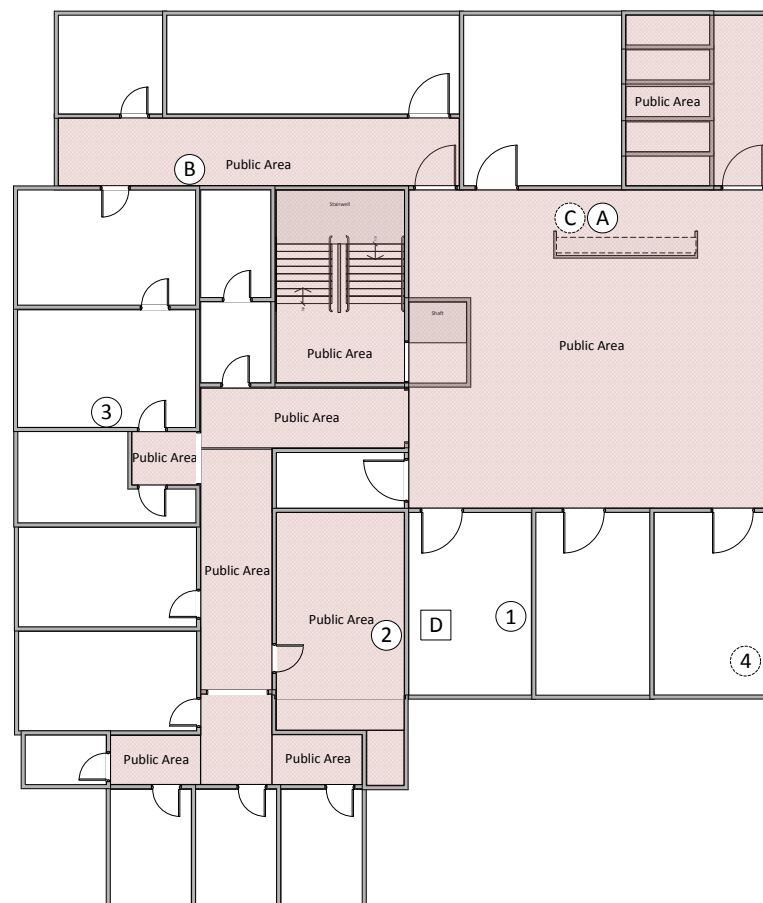


Figure 4.6: Floorplan of the target building, showing the public locations (shaded in red) and the private locations (in white). The detector is marked with a D symbol, the attack locations with numbers corresponding to entries in Table 4.3. The markers with dashed lines are on the floor below.

For the frequency-domain approach, the STFT rate was 120Hz, corresponding to $t = \frac{1}{120}$ with peaks 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.

4.9.2 Detection Accuracy

Figures for detection accuracy with each method can be seen in Table 4.4, along with precision and recall metrics.

EMPower performed well across the tested locations, although differently for each detection approach. Peak accuracies were 97.8% using frequency-domain detection and 100% using time-domain detection, whilst minimal accuracies were 74.6% and 50.2% respectively. The wide variation in accuracy is determined by the complex factors discussed in Section 4.7.1 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 irrespective of distance, although its recall fell as distance increased. In other words, this 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; for communication on the floor below. By contrast, the time-domain approach performed near-perfectly when operating at close range, correctly identifying the presence or absence of a power line network in every case except one at these distances. However, this performance degraded far more quickly, as conditions deteriorated, than with the frequency-domain method. For the further targets, even on the same floor, the recall of the time-domain approach had fallen below 32% and for the attack on the floor below, the detection was effectively no better than random. It appears that the frequency-domain and time-domain approaches

provide complementary properties that can contribute to better combined detection than either method achieves individually.

Higher data rates over the malicious network contributed to better performance in every case; with the frequency-domain approach detecting the network in every case under these circumstances. Even the minimal management traffic on an idle network was enough in most cases, though. This means that a power line adaptor within range *would be detected mere seconds after being powered on*. Furthermore, it is worth reiterating at this juncture, that each sample used to calculate the performance metrics above and in Table 4.4 represents a single T period — only 1s of elapsed time. Over even a brief observation period, the classification accuracy can be improved further by simply repeating the measurement and determining the ratio of positive vs. negative results.

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 4.7 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 (TPR - FPR = 0.822, F-Score = 0.905) is achieved with α set at -0.038. The values in Table 4.4 are with that threshold value.

4.10 Discussion

We have seen that radiated emissions from power line networks can be used to detect them. The radiating behaviour is complex and difficult to predict for a given building, as indeed is the communication quality between individual adaptors. Some frequencies radiate better than others and some deployments are more detectable than others.

An attacker clearly would like to avoid their malicious deployment being detected. Given the properties we have established, they have two broad approaches. The attacker can reduce emissions by reducing either the transmission power of the

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

Table 4.4: Detection results. Distances are taken from target to detector. Accuracy metrics are shown for each network state and aggregated across all three.

adaptors or the utilised bandwidth. There are limitations to both approaches, however. Reducing the signalling power has a notable effect on the level of EM emissions produced by the adaptors, but also on their signalling range. The further an attacker must be from their target, the less scope they have to minimise the signal power. In our testing the Technomate TM-200HP adaptors, which implement a variable power level, are undetectable when plugged into adjacent power sockets. However over even a short distance they are unable to communicate at this low power level, periodically switching up to a detectable level when separated by a 5m power strip and consistently operating at detectable power between rooms. Alternatively, reducing the bandwidth means disabling further subcarriers than




Figure 4.7: ROC curve for the detector, computed across all test locations and network states.

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. The emissions already display sizeable variation across the full band, small additional changes will not affect the correlation a great deal. A large amount of the bandwidth must be excluded to reduce the correlation noticeably. Not only does this have a profound effect on the communication rate available for the attacker’s purposes, it also risks denying them access entirely if the remaining subcarriers are affected heavily by noise. Both adaptors must be configured to exclude the bandwidth and doing this in advance of an attack is unreliable because, as [119] notes from a large scale field test, “[n]o line is like the other”. Individual deployments display substantial variation in base noise levels across the spectrum, noise levels change over short periods of time and the lowest-noise subcarriers are those at the upper end of the frequency range, where the EM emissions are also most pronounced. The attacker must observe the noise

profile after deployment and then remotely configure both adaptors to communicate successfully whilst still minimising their spectral occupancy.

The attacker can instead increase the EM emissions in the same band used by their power line adaptors in the hope of drowning the telltale spectral pattern in noise. 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. The detector simply observes this powerful, suspicious signal and raises the alarm. Assuming the attacker does not know the location of the detector they would have to jam with enough power to overcome the emissions from power line adaptors everywhere in the building, whilst simultaneously attempting to remain under the threshold for triggering the jamming alarm.

Detection of power line networks using radiated emissions should naturally be compared to detection using conducted emissions, i.e., a monitoring 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 shown that power line network connectivity, while sometimes very far-reaching, can also be severely limited if noise conditions are poor. Furthermore, some electrical devices effect powerful attenuation upon a conducted signal (such as RCD devices, or transformers in distribution boards). Indeed, we demonstrated in Section 4.3, conditions under which no power line communication is achievable between floors. A detector that relies upon conducted emissions must be very carefully placed to avoid it monitoring only a small segment of the network, if it is even possible to achieve full coverage with a single adaptor. 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 low-cost hardware and appropriate software. Nothing need be attached to the power network, a user

need only walk around their premises with a detector, or distribute a handful of them. No HomePlug-compatible hardware need be developed and even a home-made detector can be built for less cost than purchasing a commercial power line adaptor.

Indeed, while our experimental setup made use of bulky equipment and desktop computing, the methods could of course be implemented on more lightweight equipment. This in turn could realise some important use cases:

Distributed detection A number of detectors could be distributed around a premises in order to provide coverage of each area with high detection accuracy. They could monitor constantly and inform the user of the appearance of a malicious power line network shortly after it is established.

Mobile detection A portable detector could be used to periodically survey a wide area easily. Additionally, given the distance component of the emissions that we have observed in Section 4.9.2, there appears to be some scope for implementing coarse-grained localisation of communicating adaptors to assist in removing the devices.

A notable limitation of EMPower is that it does not differentiate between legitimate and malicious network deployments. In a large building this is unlikely to be a problem, although the operators of small premises that are located close to those operating legitimate networks, must carefully manage the system properties to exclude weak, far-away signals and assiduously check their premises using a short-range detector. By the same token, the detection methods presented herein are incapable of differentiating between multiple cohabiting power line networks. If a user is operating their own legitimate network, the detectors cannot inform them of the presence of a second, malicious one.

We have focused on the HomePlug AV standard throughout this paper; a standard introduced in 2005 and since superseded. A large number of contemporary devices still implement only this standard, but HomePlug AV2 and G.hn compliant

adaptors are also in widespread existence⁸. Both of these standards share an enormous amount in common with the older HomePlug AV design; including the OFDM signalling, utilised (albeit extended) frequency band, filtering notches and preamble structures. As seen in Figure 4.3b, the Sumvision SVW1000 adaptors (implementating HomePlug AV2 and advertising gigabit speeds) demonstrate the same spectrum usage as with the standard HomePlug AV adaptors. They also exhibited additional emissions at various points all the way to the maximum 86MHz limit, while the additional notches specified in the HomePlug AV2 standard were visible through the spectrum [123]. As such we are confident that such devices are equally detectable by our frequency-domain method as earlier implementations. The time-domain method would require reworking to handle the increased clock rate (200MHz, up from 75MHz) and incorporate the multiple preamble types that exist in that standard, but should otherwise apply in most cases. There are also two PPDU formats introduced in HomePlug AV2 that contain a ‘short delimiter’ (which omits a separate preamble and includes it in the frame control data) and our time-domain method would not work in these cases. However, legitimate receivers must synchronise in time even without a preamble and we believe we could adopt that mechanism also to achieve the same ends. For the similar reasons we believe that devices implementing the G.hn standard, again with manifold similarities, will also be detectable and to the same extents [133][134], however no devices are currently available for sale in our region and so we were unable to test this in practice.

4.11 Conclusion

We have identified how an attacker can easily make use of ubiquitous power networks to establish an unmonitored covert communications channel for eavesdropping and bulk data exfiltration or as a platform for further network attacks.

We have demonstrated that unintentional electromagnetic emissions, previously only studied in lab settings, are abundant in real-world settings. We have argued

⁸Broadly, all HomePlug devices advertising speeds ‘up to 200MBit/s’ are HomePlug AV devices, those offering ‘up to 500MBit/s’ implement the derivative HomePlug AV500 and devices offering a PHY rate greater than that use the newer HomePlug AV2 standard.

for the usefulness of these emissions in detecting the presence of malicious networks, given that the wireless propagation confers advantages beyond the wireline channel.

We have demonstrated simple frequency- and time-domain detection methods and shown that even these can accurately identify the presence of a network in a real office environment; with perfect accuracy within the same room and still 74.6% accuracy two rooms away and on a different floor. The methods have been shown to detect an attacker at a maximum distance of 12.9m. More complex signal processing could undoubtedly enhance this performance further.

Crucially, it was also shown that detection using wireless emissions was possible in locations that were physically close to the source, but from which the conducted signal in the wireline channel could not be received. This provides an obvious immediate security benefit by detecting networks more effectively, but also suggests the potential for developing further to allow source localisation as well. The appearance of these benefits supports the idea that defensive systems built atop physical phenomena can also take advantage of the wealth of unintended side-effects that are known. The growing accessibility of these side-effects with commodity systems can benefit security systems just as it benefits attackers.

5

Electric-Vehicle Charging

Contents

5.1	Introduction	137
5.2	Motivation	138
5.3	Background	140
5.3.1	Combined Charging System (CCS)	143
5.3.2	CCS Security	144
5.4	Related Work	145
5.5	A Near-Ideal Side-Channel	147
5.6	Threat Model	151
5.7	PLC Eavesdropping Tool	151
5.8	Real-World Measurement Campaign	155
5.9	Results	157
5.9.1	Eavesdropped Communications	157
5.9.2	Effects of Location	160
5.9.3	Message Recovery	160
5.10	Security Analysis	162
5.10.1	Unencrypted Communications	162
5.10.2	Private Data	163
5.10.3	Charging Attacks	164
5.11	Lessons Learnt	166
5.11.1	Wireless Threats	166
5.11.2	Reliance on a Non-Existent PKI	166
5.11.3	Available PHY Security Disabled	167
5.12	Potential for Active Attacks	168
5.12.1	Jamming in detail	169
5.12.2	Message injection in detail	171
5.12.3	Relay in detail	171
5.13	Countermeasures	173
5.13.1	Protocol Changes	173

5.13.2 Equipment Changes	174
5.13.3 Physical-Layer Security	175
5.14 Conclusion	175

Disclosure Statement

We disclosed the findings of this chapter to the tested vehicle and charger manufacturers, along with AutoCharge operators.

5.1 Introduction

Through the previous chapters, the effective use of physical-layer features for security has been shown across a range of systems. In this chapter, the focus turns to examining a modern system to understand whether concepts of physical-layer security and the growing accessibility of physical-layer attacks are being incorporated into designs.

The system in question is one for charging electric vehicles. As electric mobility expands, a new wave of charging infrastructure is fast becoming critical infrastructure. It is an ideal example of new cyber-physical systems proliferation; integrating a core physical process of power delivery with communications and logical processing for control, centralised monitoring, billing and as a platform for envisaged third-party services.

The system is widely recognised in the transport industry and codified in international standards. Considerable attention is paid, in its standardised description, to the security model and widespread use is made of traditional cryptographic primitives to secure data and limit the trusted computing base.

Nevertheless, the system has come into being at a time when physical-layer security is already an established field of study. Attacks, at the physical-layer, on both publicly-accessible hardware and communications systems are widely documented. In particular, many communications attacks [135–138] and serious physical-layer vulnerabilities [20] have been noted in automotive systems and are

now a burden upon users worldwide. The presumed use of this system as a key part of individuals' lives, along with the potential for it to be an underpinning platform for other systems, makes it imperative to consider carefully all known attack vectors when designing it.

The system makes use of powerline communication to allow communication between a vehicle and a charger. The potential for leakage in powerline communication has been known for years, even if its security relevance has only recently been articulated¹. Its selection raises questions about the potential vulnerability of the charging system to physical-layer communication attacks.

This chapter examines whether the problems of Chapter 4 have been controlled in the design of the system. It focuses on the possibility of wirelessly recovering communications from vehicle-charging sessions in deployed production systems. As well as yielding direct observations about the system's resistance to a passive attacker, this method is also used to infer the likelihood of vulnerability to active physical-layer attacks as well.

The results are discussed in the context of the design methodology used to create the system.

5.2 Motivation

The rise of electric vehicles (EVs) as a contemporary and future transport mechanism has been swift in recent years and continues to accelerate, helped by prevailing attitudes, technological advances and notable personalities contributing in the area. There are already widespread government plans to eradicate fossil-fuel vehicles in cities [139], states [140] and countries [141] in the coming years.

As EV technology advances rapidly, the availability of charging infrastructure has become a challenge for users, who require access both to private charging points at home and public ones on longer journeys. The lack of sufficient charging points is noted as a slowing influence on adoption of electric mobility [142] and this has prompted endeavours to expand the infrastructure, both from governments

¹Herein and in associated publications

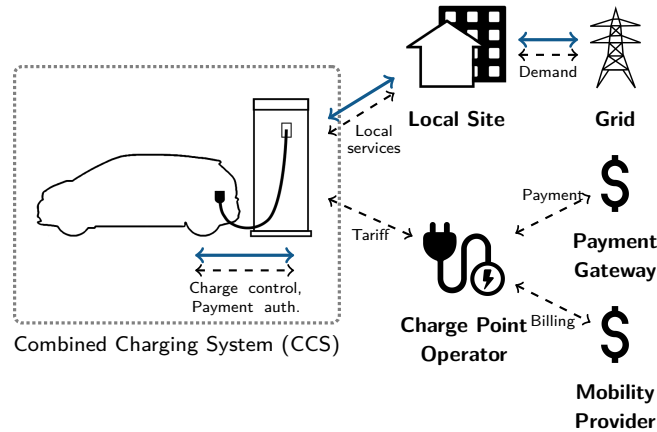


Figure 5.1: Overview of EV charging with V2G and payment options shown. Solid, blue lines indicate power flow whilst dashed, black lines indicate communication.

recognising the potential public good and from competing EV manufacturers who understand that having the best infrastructure makes their vehicles more appealing to purchasers. There are already multi-billion dollar public deployment plans in progress [143] and predictions of worldwide numbers exceed 50 million chargers by 2025 if private systems are included [144].

With several major charging standards in existence, the race to become the dominant one has reached a fervour in recent years and a new generation of high-power charging systems has emerged. But the pressure to achieve rapid expansion has so often been seen to inhibit secure implementation. Users demand charging systems that are consistent and convenient, but with such drive for the adoption of electric mobility, it is critical that they are also secure. The security community has raised concerns in the past that standards do not fully address security and privacy issues [145–147], as well as noting vulnerabilities in back-end and payment systems of earlier charging system deployments [148, 149].

Meanwhile the complexity of developing all the infrastructure required for a secured charging network is enormous. As Figure 5.1 shows, vehicle charging involves interaction between the vehicle, the owner, the charger operator, a payment gateway and the grid regulator. This requires establishing communication links capable of supporting the higher-level protocols for this interaction, within a dynamic and untrusted environment, where many thousands of users come and

go. It also necessitates trust relationships between all the participants to ensure each is acting legitimately.

In light of the challenges this infrastructure development faces and the acknowledged side-channel vulnerabilities that exposed cabling presents, we undertook to investigate the security of the charging cable communication.

We make the following specific contributions:

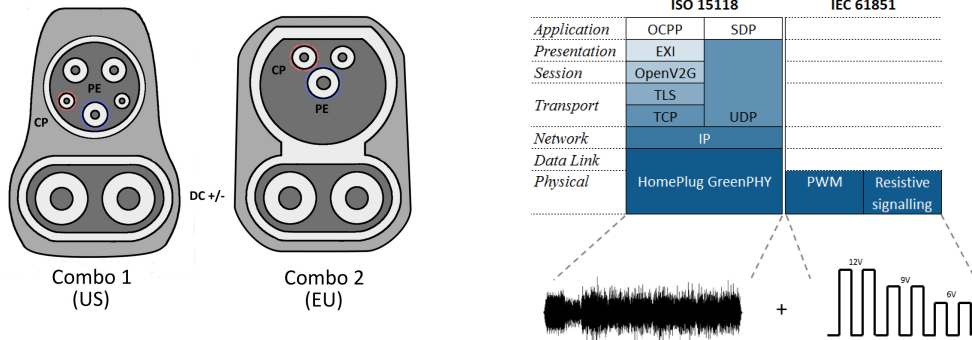
- *Demonstrate that the use of powerline communication, and its specific configuration in CCS, makes systems particularly vulnerable to EM eavesdropping*
- *Develop an eavesdropping system for HomePlug GP and the DIN70121/ISO15118 PHY-layer*
- *Conduct a real-world measurement campaign, demonstrating the widespread nature of the problem*
- *Highlight the potential for privacy violation and user tracking with existing systems*
- *Propose potential active attacks on existing and future charging systems that permit denial of service and billing fraud*
- *Propose countermeasures to mitigate the capabilities of an eavesdropper*

Our findings are relevant to thousands of chargers across Europe and North America [150, 151], along with having implications for ongoing deployments both in public locations and private homes.

5.3 Background

The availability of EV charging infrastructure is growing enormously. Early, simple alternating-current chargers are being superseded by a new generation of charging technologies that provide greater charging power and advanced functionality. The greater power is provided by the use of direct-current (DC) charging, allowing an enormous increase in current delivery over previous alternating-current designs. Public DC charging stations currently well exceed the 3kW power levels commonly

available in a home, with 50kW supplies plentiful and those providing up to 350kW soon to appear [152][153]. But the improvements in *power* are only part of the benefit of this new generation of technologies. The *communication* capabilities are also vital to enable a host of new uses:



(a) Two charging cables are used by CCS. The Combo 1 and Combo 2 plugs are dominant in the US and Europe respectively, while other locations adopt one or the other. DC power is delivered by the large conductors at the bottom of the plug, meanwhile communication happens over the Control Pilot and Protective Earth lines (red and blue, respectively).

(b) CCS high-level and low-level signalling share the same communication lines. The corresponding ISO 15118 PLC and IEC 61851 systems have their signals superposed at the physical layer. The PLC provides a standard IP stack for use by charging traffic and other services.

Figure 5.2: Illustrations of the physical connectors for CCS charging, along with the network stack used for communication.

Reactive charging allows a vehicle to vary its charging process based on electricity price or expected time of departure.

Automatic billing or “plug-and-charge” allows a vehicle to authorise billing of its owner for charging, without the owner explicitly interacting with it. Aside from the obvious convenience benefit, the same capability also allows the user to ‘roam’ between charging providers with a seamless experience as cross-provider billing is handled automatically as well.

Vehicle-to-Grid (V2G) makes use of bidirectional power flow to allow the vehicle to deliver energy as well as consume it. As energy prices fluctuate with demand, the vehicle can either act as a storage battery for a user’s home or sell energy

back to the grid on demand. This can bring economic benefits for the user and stability improvement for the grid operator.

External payment is commonly provided by RFID cards [149], apps that communicate with the charger separately or card payment terminals. Additional systems exist though, for payment through separate providers or via a blockchain network [154–157].

Additional services that operate in conjunction with charging are proposed [158]. In a private environment this might comprise access to the local network to communicate with smart-home devices or make use of domestic internet service and avoid mobile network charges. At public charging stations site-specific services such as loyalty schemes, to-vehicle delivery, parking charges or ‘where-have-I-parked’ reminders can operate, with middleware layers to support an app ecosystem in commercial development [159]. Internet access can also be made available for connected vehicles in areas without mobile network coverage, such as underground parking complexes.

Examples of each are in production use and deployment is becoming more widespread. The underpinning communication mechanisms go beyond indicating presence and readiness to charge, also providing a general-purpose channel for software operating in the vehicle and charger. Figure 5.1 shows the potential extent of communication during charging. The vehicle can demand current flow, the charger can provide tariff information for reactive charging or reverse current demands for vehicle-to-grid, and the two can interact with external parties for automatic billing or to provide additional services.

Four major next-generation charging systems exist: CHAdeMO², Supercharger³, GB/T 20234⁴ and the Combined Charging System (CCS)⁵. Each uses the charging

²An open standard developed by Nissan and dominant in Japan

³A proprietary standard developed by Tesla Motors

⁴A nationwide standard in China

⁵An open standard backed by the European Union

cable for primary communication: CHAdeMO, Supercharger and GB/T 20234 make use of CAN-Bus, whilst CCS makes use of powerline communication (PLC).

We examine the CCS standard as it has the most extensive, current functionality (supporting reactive charging, automatic billing and additional services) and has been adopted by seven out of the ten largest automobile manufacturers by production numbers [160]. In addition it is being integrated by competing manufacturers, such as Tesla [161].

5.3.1 Combined Charging System (CCS)

The Combined Charging System (CCS) is an amalgamation of standards governing all physical and logical elements of the charging infrastructure; from the physical connector to the protocols for automated billing. Figure 5.2a shows the charging plug, while Figure 5.2b illustrates the communications undertaken. The communication between vehicle and charger is standardised as DIN 70121 and ISO 15118. These use powerline communication (PLC) over the Control Pilot (CP) and Protective Earth (PE) lines of the charging cable. The PLC shares the lines with the older IEC 61851 signalling system for backwards-compatibility reasons, with the signals superposed at the physical layer. The specific PLC implementation is HomePlug GreenPHY (HPGP) [162], a derivative of the commonplace broadband LAN technologies sold to consumers, that has been modified to support pairing between devices with no pre-shared key, and to be more robust to noise. Atop the PLC, a full IP stack is provided to act as the general-purpose channel. The same standards also define interactions for identification, authorisation and control, with additional features in ISO 15118. Communication persists throughout the duration of charging and allows charge parameters to be varied quickly.

CCS provides reactive charging by allowing a charger to present current and future tariff information to the vehicle, which can then make charging requests based on a user's settings. The user may have a price preference or timing constraints for when the vehicle should be charged. Contract-based automated billing is implemented by having a user's contract with a charging provider represented by a

public-key certificate stored on the vehicle. A complex public-key infrastructure (PKI) then allows the vehicle to authenticate the charger, the charger to validate the charging contract and the provider to produce verifiable metering receipts. The same PKI is used to underpin the TLS tunnel for protecting traffic.

Competing automated billing approaches do exist however, that do not use the contract-based approach, nor rely on the PKI. Blockchain-based payment systems, seeking to protect the user’s privacy from charging operators, simply use the communication channel as a building block for their own service [156, 157]. A system named “AutoCharge” [163] is also used in some networks [164, 165] to enable automated billing for even those users whose vehicles do not support the required certificates. The AutoCharge system is based on a simplified ISO 15118 use-case [166] that uses only vehicle-provided identifiers to match the vehicle to a customer record at the provider.

As there is a general-purpose channel, any IP communication is supported for additional functionality. Fast internet access is suggested in the ISO 15118 standard and a selection of data collection, targeted marketing, on-demand entertainment and third-party app platforms are emerging to take advantage of this [156, 159].

5.3.2 CCS Security

Communication security is considered in many of the systems making up the CCS standard; with traffic encryption available at the PHY layer and possible TLS and XML Security at higher layers [167, 168].

At the PHY layer, the HPGP PLC network maintains a shared secret key called the Network Membership Key (NMK), with ephemeral Network Encryption Keys (NEKs) rotated periodically. All MAC-layer traffic is encrypted via AES-128 using the NEK. However, HPGP security is based upon a private-network model, while EV charging is fundamentally a public-network model. To adapt the technology to the use case, additions were made to HPGP to incorporate an initial association protocol⁶, during which the vehicle and charger verify that they are connected to

⁶The comprehensively-named GreenPPEA, or “GreenPHY Plug-in-electric-vehicle Electric-vehicle-supply-equipment Association”

each other and are not communicating with the wrong party due to crosstalk. The determination is known as Signal-Level Attenuation Characterisation (SLAC) and is illustrated in Figure 5.3. The protocol involves the vehicle sending a series of sounding messages, for which the charger reports the measured attenuation. If multiple chargers respond, the one reporting the least attenuation is selected and communication commences. Once a charger is selected, a Network Membership Key is created by the charger and used to establish a private network. The key is then sent to the vehicle in the final `CM_SLAC_MATCH.CNF` message of the protocol. The SLAC protocol can operate in a secure mode, with mutual authentication and encrypted communication, but this capability is optional if supported by both parties. Indeed, despite the availability of this mechanism, the DIN 70121 and ISO 15118 standards specify that SLAC only operates in its plaintext mode, leaving message security to TLS.

Once a network is established, a TLS connection is only created under certain conditions. Within DIN 70121 this is never done [169]. Within ISO 15118, if contract-based automated billing is used then TLS is required, similarly for the discovery of additional services, but only when they are ones that are defined in the ISO 15118 use cases. When charging is externally authorised, no TLS is required for the control traffic [168]. The options for external authorisation are open; including RFID cards, mobile-app networks, manual authorisation by a charger operator or some other service operating on the charger. The method need not be external to the charger, only external to the ISO 15118 scope. For all other traffic not managed through a standard use case, security is left to the implementer. In the alternative payment example of [157], independent IP communication is undertaken completely outside the scope of ISO 15118 (although secured in an SSH tunnel in that case).

5.4 Related Work

The privacy and security issues surrounding EV charging are the subject of ongoing work; with attempts to devise architectures that protect each stakeholder [170, 171] and analyses of the security of upcoming standards [145, 146, 172]. These

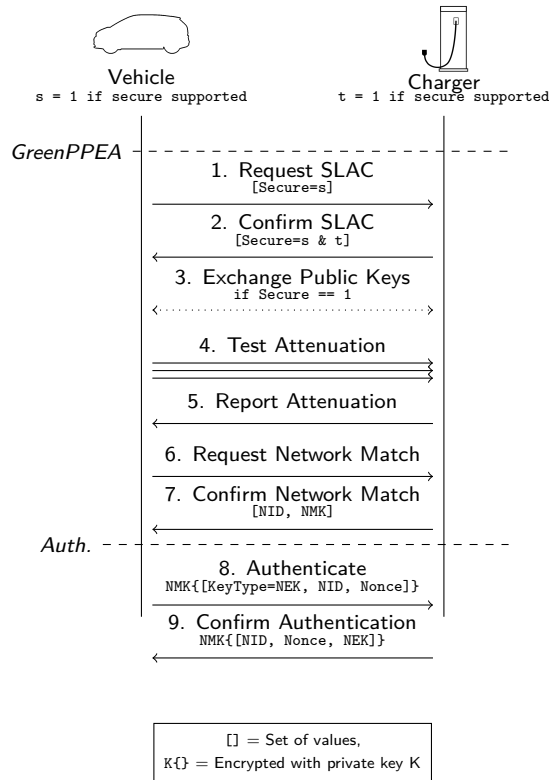


Figure 5.3: An overview of vehicle-to-charger network establishment in HomePlug GreenPHY. If the secure mode is supported by both parties and enabled in initialisation then step 3 occurs, allowing the messages in steps 5–7 to be signed and the one in step 7 also encrypted.

works are theoretical in nature, however, and leave aside implementation issues. They also assume a wireline threat model for attacks on the vehicle-to-charger communication, discussing where an attacker must use “a modified cable or an adapter plug installed on the [charger]” [146]. By contrast, we consider a wireless threat model that permits deniability on the part of the attacker.

Practical attacks have been demonstrated on previous-generation infrastructure, particularly against RFID authorisation [148, 149], but require the attacker to clone a user’s physical token or access debug ports on an unlocked charger. General, system-level attacks against common RFID authorisation implementations have been known for nearly a decade too [173].

Since electromagnetic emissions security was brought from a military discipline into academic study by van Eck’s work on eavesdropping video displays [109], efforts have been devoted to studying a wide range of systems [174]. Recent

work has focused primarily on extracting secrets from operating devices [175, 176], although the emissions security of digital communication systems have been studied in the context of eavesdropping on RS232 serial devices [177] and 100BaseT ethernet [178], along with use as a covert channel for USB [179]. While radiated emissions from powerline communication have been studied from an electromagnetic compatibility perspective [107], we demonstrate the first practical wiretap attack using these emissions.

Vehicle tracking using unique identifiers has been studied in the context of electronic license plates [180], tire-pressure monitoring systems [135] and vehicular ad-hoc networks [181], highlighting the impact upon individuals' location privacy and inspiring this work on new charging technologies. Practical attacks have also been demonstrated to wirelessly compromise in-vehicle systems [136], to unlock vehicles for theft via remote keys [137] or passive entry [20, 138] and to misdirect drivers to unwanted locations [182]. These attacks consider an active attacker with different goals to those studied here and as such could be considered orthogonal to our work.

Energy monitoring has been shown to enable the tracking of individuals [183] and this has prompted proposals to mask energy signatures, such as by using rechargeable vehicle batteries [184], which assumes that data about vehicle power flow cannot be monitored.

5.5 A Near-Ideal Side-Channel

The underlying principles of electromagnetic (EM) side-channels are very well-explored and their study has informed modern security design [174]. Despite this, we describe here how the use of PLC and its specific arrangement within CCS exacerbates the vulnerability to EM attacks.

The design of PLC technologies assumes differential signalling; wherein two identical transmission lines that are located in close proximity are driven with equal but opposite signals, such that those fields largely cancel and no residual electric field exists. Practical challenges often break these underlying assumptions for in-home PLC deployments, leading to EM interference and susceptibility thereto [107].

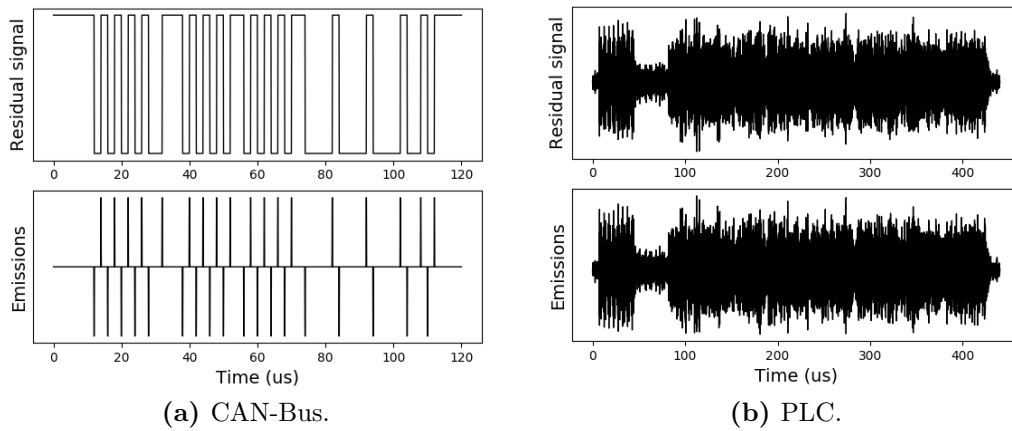


Figure 5.4: Example single-ended signals, with the radiated emissions that result. As the emissions are the gradient of the signal, the square wave produces only impulses while the OFDM waveform is all but unchanged.

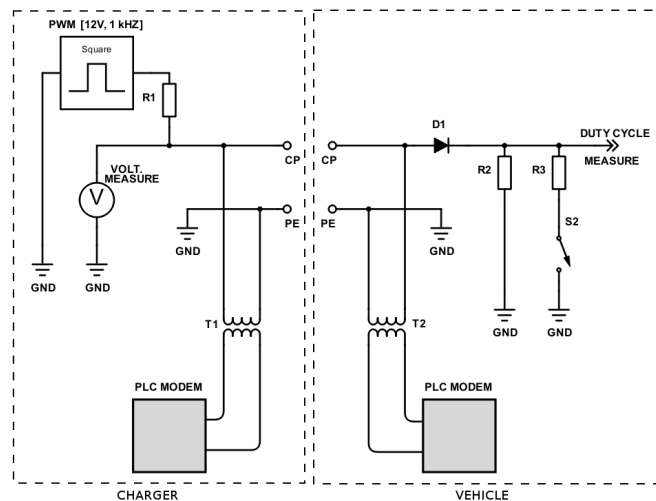


Figure 5.5: A diagram of the CCS communication circuit. The loads on each line connected to the PLC modem are not balanced. Resistors R2 & R3 alter the voltage in the low-level communication, but also vary the imbalance further.


Despite EV charging only requiring simpler and more constrained wiring than domestic electrics, these assumptions are still broken in CCS.

Figure 5.5 shows the communication circuit for PLC in CCS charging systems, including the connection of the circuit to the Control Pilot and Protective Earth lines, along with the additional components affecting the Control Pilot line due to the need for backwards-compatibility with the IEC 61851 communication that shares the lines. The design choice to incorporate backwards compatibility with an

earlier low-power charging standard led to a PLC circuit design that connects one transmission line to ground (see Figures 5.2b and 5.5). This renders the signalling *single-ended instead of differential*. With no inverse field, the charging circuit functions as a suitable antenna for emissions or interference.

The nature of the PLC waveform itself, however, makes it ideal for wireless observation and interaction. It can be seen in Figure 5.4, operating as a single-ended system alongside single-ended CAN-Bus communications for comparison. The radiated signal represents the gradient of the original signal: only the changes in voltage. This can introduce problems for an attacker whenever they wish to observe or inject signals with constant voltage levels (respectively a minor problem and a major one), most notably the square waves used ubiquitously in digital communication (and in other EV charging communication based on CAN-Bus). For observation, a static voltage produces no useful emission, so only the state transitions are detectable. The attacker uses these where they can or hopes for the signal to leak elsewhere in the circuit and be modulated onto a more easily-observable one [174]. For injection the attacker cannot directly induce the desired static voltage level without enormous bandwidth usage and instead must exploit nonlinearities in components or undersampling effects in order to synthesize the desired signal at the victim [16]. In the absence of any such components in the target circuit that can be subverted, or in the presence of additional filtering, the attacker's opportunities to inject an arbitrary waveform are limited.

Broadband PLC technologies predominantly use orthogonal frequency division multiplexing (OFDM); in which the data are modulated in the frequency domain before constructing a time-domain waveform using an inverse Fourier transform. The resulting, transmitted waveform is a finite sum of sinusoids and does not exhibit any non-zero static voltage levels. *The observed emissions simply form a phase-shifted replica of the original signal.* The attacker therefore does not need to make inferences to determine the original signal from eavesdropped observations, nor predict what transformations an injected signal will undergo in the receiver. They need only contend with the characteristics of the channel itself.



(a) Signal given as input to waveform generator. (b) Signal observed on oscilloscope connected to PLC circuit.

Figure 5.6: Observations from PLC signal injection in lab conditions. The gray signal is the emitted signal, the black signal is that observed on the PLC line when communication was idle.

It follows naturally from the principle of radio reciprocity, and from the identity transformation of the PLC signal under radiation, that the PLC system would not only exhibit emissions from its transmissions lines, but that they would also be susceptible to signal intrusion.

This is illustrated in Figure 5.6, from confirmatory tests in a lab. The gray signal in Fig. 5.6a when given as input to a Rigol DG4062 waveform generator⁷ driving a long wire antenna positioned near a PLC circuit injects the black signal in Fig. 5.6b onto the PLC transmissions lines, as was observed on an oscilloscope connected to the PLC circuit directly. Some degradation of the signal was evident, as the channel properties affect the injected signal just as they affect a radiated one. Nevertheless, the channel is bidirectional.

While we focus on emissions throughout our real-world testing, as could be tested on public equipment without fear of causing unintentional damage, we discuss

⁷Due to limitations in the memory depth of the waveform generator, the signal here is substantially downsampled from the full HomePlug rate and so has lost fidelity against the example in Fig. 5.4b. No SDR was available at the time that could transmit at such a low frequency, so the generator was required. Nevertheless, with appropriate hardware a full-rate signal could be injected easily.

the consequences of signal injection and active attacks in Section 5.12.

5.6 Threat Model

While we discuss the channel properties in a bidirectional sense above, we focus our further investigation and practical attacks on passive eavesdropping. Testing on deployed infrastructure restricts us to only passive operation.

The attacker listens to the unintended electromagnetic radiation of the EV charging communication. Their goal is to eavesdrop on the general-purpose channel established between the vehicle and the charger; such that they obtain access to private data it carries. The attacker can approach close to the target vehicle and charger but cannot modify or interfere with the equipment. They perform their attack either *in-person* from a nearby location, or by *situating a device* at the site and leaving it unattended.

We justify this model on the basis of deniability and access. Interfering with a vehicle or charger is an immediately suspicious activity that would draw attention from the owner, people nearby and operators reviewing CCTV footage. The charging equipment is also handled regularly by drivers, so a cable modification or plug insert is more likely to be noticed. By contrast parking near another vehicle at a public station or briefly visiting a private property appear to be benign actions.

5.7 PLC Eavesdropping Tool

Given the properties described in Section 5.5, the passive attacker's task is the same as that of a legitimate receiver; to maximise the signal-to-noise ratio (SNR) and bandwidth (BW) of the received signal. In a real setting, additional complicating factors exist. While the exposed components are the most obvious targets, any element of the communication circuit (i.e., charging plug, cabling, vehicle, charger), or indeed multiple elements, could act as an unintentional antenna(s). The size of the equipment makes potential antennas physically distant from one another, so it can be difficult to predict the location that optimises the SNR and BW for

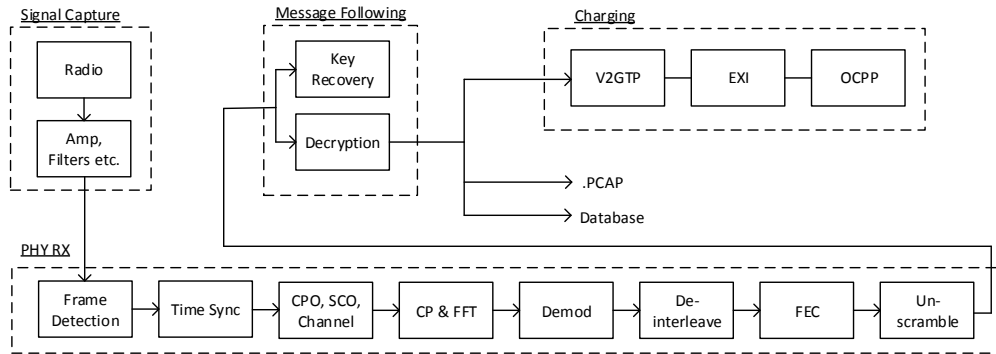


Figure 5.7: Architecture of PLC monitoring tool. The signal is captured and prefiltered, before moving through a software receiver chain to recover messages. The message following behaviour extracts security-relevant data and stores all messages. Charging traffic can be further processed, while traffic using other protocols will need separate onwards processing.

each target. Similarly, electric vehicles and chargers operate at high electrical powers and can easily cause significant interference levels, which must be suitably mitigated by careful positioning or filtering.

Exploiting the properties and design choices of CCS, we developed a tool for wireless eavesdropping of the underlying physical layer; a HomePlug GreenPHY (HPGP) network. The tool is applicable to monitoring any HPGP network as well as network management traffic in HomePlug AV and AV2 networks, although the vehicle charging scenario is particularly beneficial for the reasons discussed above. The tool is open-source software⁸.

The eavesdropping tool broadly resembles a normal HPGP receiver. While the HPGP standard is public, all compatible implementations are proprietary and implemented as integrated circuits. Our pure-software implementation allowed far greater insight and flexibility during captures however, particularly for experimenting with different preprocessing steps to improve reception and collecting partial data that would be discarded by a black-box implementation. The receiver architecture can be seen in Figure 5.7. Given that Wi-Fi shares the same OFDM underpinnings, the overall structure bears many similarities to a Wi-Fi receiver.

⁸<https://gitlab.com/rbaker/hpgp-emis-rx>, (MIT licence)

Similar stages of time synchronisation, channel correction, domain-transformation, demodulation and error-correction occur, albeit distinct in details to match the HPGP protocol specification.

As the signal processing chain is complicated we describe it briefly here but elide full details from the main text, providing them in Appendix A instead. The signal is captured and digitally filtered to suppress local interference. Messages, known as PHY-layer Protocol Data Units (PPDUs), are identified using a power detector and correlation of the signal preamble against the known preamble structure. As an OFDM technology, data are represented in individual *symbols* throughout the Frame Control and Payload sections of the PPDU. Once the receiver is time synchronised to the PPDU, each symbol is processed in turn; with channel estimation and frequency offset correction applied before demodulation. With complete messages the Turbo Code error correction is processed to reduce errors and the Cyclic-Redundancy Check checksums are calculated (a CRC24 for the Frame Control and a CRC32 for the Payload). The application of the Turbo Code decoder is limited in our tool, owing primarily to the computational cost of the process. A Turbo Code is intended to be decoded by iterating a probabilistic decoder over various rearrangements of the received bits. We use only a single pass of the decoder and its application already dominates the message reception time; exceeding the rest of the software processing chain. As such we suffer from reduced error-correction performance compared with an arrangement using multiple repetitions. Such an arrangement could be expected to receive more messages correctly in all circumstances.

Site	Location	Type	Charger (Operator)	i3	Vehicle		Charge Sessions
					I-PACE	e-Golf	
A	Oxford Belfry, Oxon.	Hotel	DBT Dual DC [185] (Polar [186])	✓			1
B	Abingdon, Oxon.	Superstore	DBT Dual DC [185] (Polar [186])	✓			1
C	Maldon, Essex	Superstore	ABB Terra 53 C/JG [187] (POD Point Open [188])	✓			1
D	South Mimms, Herts.	Road services	DBT Dual DC [185] (Ecotricity [189])	✓			1
E	Bishops Stortford, Herts.	Road services	DBT Dual DC [185] (Ecotricity [189])	✓			1
F	Hythe, Kent	Road services	DBT Dual DC [185] (Ecotricity [189])	✓	✓	✓	9
G	Dover, Kent	Superstore	ABB Terra 53 C/JG [187] (POD Point Open [188])	✓			10
H	Marden, Kent	Local garage	Chargepoint CPE200 [190] (InstaVolt [191])	✓	✓	✓	15
I	Chatham, Kent	Racetack	Chargemaster Ultracharge 500S [192] (Polar [186])			✓	1
J	Ticehurst, Kent	Golf club	Chargemaster Ultracharge 500S [192] (Polar [186])		✓	✓	4
K	Hawkhurst, Kent	Local garage	EVTronic QUICKCHARGER [193] (GeniePoint [194])		✓		2
L	Tunbridge Wells, Kent	Local garage	Efacec QC45 [195] (Shell Recharge [196])			✓	2
M	Hastings, Sussex	Local garage	EVTronic QUICKCHARGER [193] (GeniePoint [194])			✓	1
N	Milton Keynes, Bucks.	Public car park	Efacec QC45 [195] (Polar [186])			✓	5

Table 5.1: Details of all tested charging locations, across the southern United Kingdom. There were a total of 54 unique charging sessions. Multiple signal captures were taken during each session; at initialisation, during charging and at shutdown. At sites **F** and **H**, two vehicles were charged and monitored simultaneously.




Figure 5.8: A composite diagram showing the experiment layout. The five antenna locations are denoted with a dashed \times symbol.

5.8 Real-World Measurement Campaign

To explore the accessibility of the wireless side channel, we undertook a data collection campaign with three fully-electric vehicles: a BMW i3, a Jaguar I-PACE and a Volkswagen e-Golf. The campaign comprised over 800 miles of driving and spanned six major administrative regions of the UK. A total of 54 unique charging sessions were conducted, at locations including garages⁹, motorway services, supermarkets and hotels.

The three vehicles we observed all implemented the DIN 70121 standard, a derived subset of ISO 15118 that does not require the use of TLS and does not provide additional features beyond charging control. At a physical communication level the standards are the same and our claims apply equally to both.

During charging sessions, we monitored radiated emissions to measure the extent of signal leakage and the ability of an attacker to eavesdrop it. Where we were able

⁹A fun aside that the author had time to contemplate, while sat waiting for a car to charge, is that the prohibition on mobile phone use in forecourts seems hollow when a 50kW charger with the error state ‘welding detected’ can be deployed there without concern.

to receive sufficient emissions we used the tool detailed in Section 5.7 to recover the original transmissions and examine the communication itself. For the majority of our testing we monitored one vehicle at a time, although we did conduct testing with multiple vehicles to examine the effects of cross-traffic. Further details of the locations and installed hardware are given in Table 5.1, while examples can be seen in Figures 5.10 and 5.11. All of the chargers are state-of-the-art at the time of writing. We tested only public chargers due to their availability, but equivalent chargers for private use are also on sale [197]. As the chargers were public, we did not modify or interfere with the equipment in any way. The vehicle, charger and associated cabling remained entirely untouched. While this prevented us from injecting messages or capturing ground-truth via a directly-connected receiver, it was necessary to conduct a widespread survey of existing infrastructure.

At each site, the vehicle was parked and connected to the charger for a series of charging sessions¹⁰. The receiving antenna was placed at various locations to investigate the reception capabilities. As noted in Section 5.7, deriving an optimal attack location beforehand is challenging, so this placement was exploratory. The locations are illustrated with a dashed \times symbol in Figure 5.8. Locations near the cable itself, on the outside of the vehicle, within the vehicle, hidden in a nearby hedge and on a nearby car were all tested. As each site had a different layout, Figure 5.8 is a composite to show the arrangements, rather than a meticulous depiction of any one site.

The data were collected using a bladeRF software-defined radio, an RF Explorer Upconverter and a GNU Radio flowgraph running on a Lenovo Thinkpad X1 Carbon laptop. We made use of an electrically-short monopole antenna to collect the signal. Owing to the long wavelengths involved, testing with a suitably-tuned directed antenna was not possible. The equipment for our experiments cost approximately \$800, although equivalent setups are available for less than \$300. The collected signal was passed through 25dB amplification and upconversion (+530MHz) to

¹⁰Care was taken to ensure we only observed signals from our own vehicles. Upon arrival we waited for any other users to leave before capturing traffic and aborted immediately if another arrived.

bring it into the tunable range of the bladeRF. Initial filtering and packet detection was performed with further GNURadio flowgraphs, while subsequent processing was implemented using Python and numPy libraries. We tuned the receiver's interference-rejection filter by observation at each site, but left all other reception parameters constant throughout.

5.9 Results

In this section we examine the results of our testing in real environments, both in terms of raw observable signal and message recovery.

5.9.1 Eavesdropped Communications

Table 5.2 details the observations for each site. It indicates the peak signal-to-noise ratio (SNR) over all the sessions, along with the widest bandwidth (BW) with a positive SNR. It then lists the count of all PPDU's detected, the number of data PPDU's, the rate at which messages were well-formed and the rate at which messages had a correct CRC32 checksum.

Every site displayed some form of unintentional wireless channel from the PLC communication, with properties that exceeded our expectations. The weakest signal showed 9dB from the peak to the background and spanned a bandwidth of 4.5MHz. In the best case 25MHz could be seen, up to a peak of 35dB. This was true irrespective of charger manufacturer, indeed varying notably between sites with the same charger hardware antenna location. This would seem to confirm the expectation that the site layout and variations in parking have a substantial impact upon reception.

Figure 5.9 shows spectrograms of the captured signal at a selection of sites, covering each tested antenna location. Overlaid on each subfigure is the utilised HPGP spectrum, showing the regions of the band in which transmission occurs. A transmission will originally have a frequency-domain representation that matches the spectral mask, with a peak power of -50dBm in utilised regions. Apparent power levels up to approximately -70dBm were observed, although the receiver was not calibrated against a reference scale so this value is uncertain. The degradation

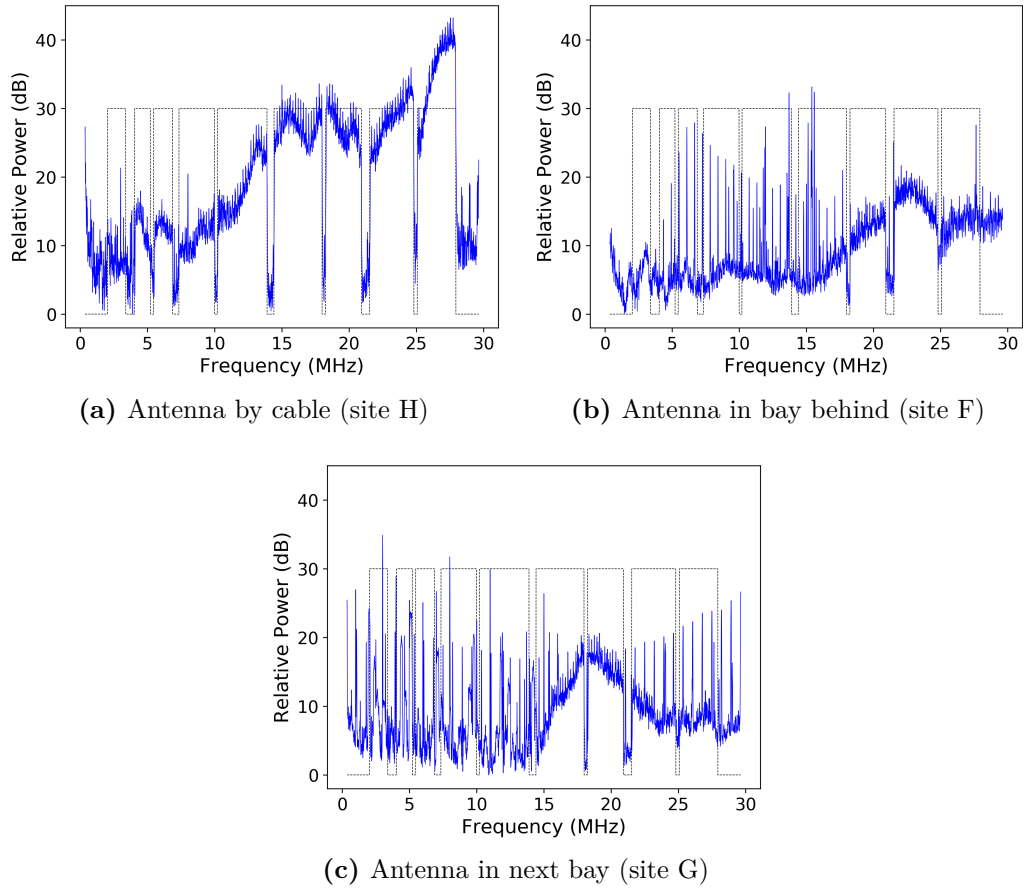


Figure 5.9: Observed signal across the HPGP bandwidth, at each antenna location. The HPGP spectral mask is overlaid to indicate the regions in which transmission occurs, although no valid comparison can be made with its power value as the measurement was not calibrated. Signal degradation and noise ingress is visible in every case, although far more prominently in (b) and (c).

of signal across the band is clear in every case; the flat-topped spectral usage of the transmission is observable as a jagged range with many subcarriers severely attenuated, particularly at lower frequencies. This correlates well with studies of the wireline channel that legitimate receivers (with a conductive connection) experience, albeit with a different noise profile [119].

Site	Antenna	Peak SNR (dB)	BW (MHz)	Total PPDUs	Data PPDUs	Bi-direc.?	Start?	RX%		CRC32%	
								Mean	Min	Mean	Max
A	In car	15	6	526	272	✓		99.3	1.1	1.8	3.3
B	In car	18	12	1063	567	✓		29.8	0.5	3.3	5.3
C	In car	25	14	2976	1819	✓		99.9	46.6	48.1	50.3
D	In car	10	12	556	293	✓		88.2	1.4	2.3	3.0
E	In car	9	4.5	569	306			100	11.0	11.1	11.2
F	In car	21	12	3660	2009	✓	✓	99.3	27.8	36.8	45.8
	Bay behind	15	8	1434	1430	✓		99.3	43.5	43.5	43.5
	Outside car	10	10	12987	8255	✓		76.2	34.9	46.6	89.5
	Two cars	14	11	2449	2274			99.1	24.3	47.5	70.8
G	In car	19	12	5837	3670	✓	✓	99.0	51.1	60.3	71.4
	Next bay	15	13	4157	2749	✓		99.7	91.8	91.8	91.8
	By cable	29	23	23984	17246	✓	✓	80.2	52.9	74.0	99.8
H	In car	16	12.5	15052	9362	✓		99.2	69.9	71.0	72.8
	Outside car	20	11	16243	10407	✓		99.5	27.7	61.6	80.6
	By cable	35	25	19535	14717	✓	✓	92.1	34.2	70.0	92.8
	Two cars	15	12	24121	21006			99.6	42.2	71.9	94.8
I	In car	20	12	1501	1193	✓	✓	98.0	94.8	97.4	100.0
J	In car	20	7	14231	10291	✓	✓	81.0	1.0	33.6	67.9
	Outside car	23	7	1084	935	✓	✓	96.0	49.2	49.2	49.2
K	In car	8	5	1971	1278	✓		92.5	0.0 †	22.0	38.3
L	Outside car	8	7	3004	1849		✓	25.8	0.0	0.0	0.0
M	In car	20	12	13631	9743	✓	✓	98.8	42.4	64.9	82.5
N	In car	24	14	4317	3364	✓	✓	68.3	0.0 †	44.5	72.6

Table 5.2: Eavesdropping results, from all sites and antenna locations. Raw signal properties are quantified as Peak SNR and Bandwidth. PPDU counts are given and the observance of bidirectional traffic and session startup is indicated. The rates of well-formed messages are then shown, along with the rates of CRC32 checksum validations. The worst and best performance for each antenna location is highlighted in **bold** († indicates joint-worst).

5.9.2 Effects of Location

While systematic examination of performance by location was not our goal, we were able to observe trends across tested antenna positions, with the fidelity of the wireless channel varying substantially. The closest representation of the transmitted signal is that shown in Figure 5.9a, obtained approximately 0.5m from the charging cable. At other antenna positions the signal loss was more pronounced, both inside and outside of the vehicle, and in isolated cases the signal was swamped by interference more than a short distance from the cabling. Making general predictions about the channel gain at specific distances is not feasible due to the low frequencies at which the PLC operates (2 – 28MHz). Even at 28MHz the wavelength is still 10.7m and so all observations were taken well within the near field of the transmitter. In this region, common path loss calculations like the Friis equation [198] are not defined and near-field effects can change the channel gain drastically from position to position. Nevertheless, Figures 5.9b and 5.9c show the results of tests at the greatest distances; 4.2m in the latter case when the antenna was positioned by a vehicle in an adjacent parking bay (shown in Figure 5.10). Interference is still substantial at these distances (e.g., everything below 15MHz in Figs. 5.9b and 5.9c), but in the higher reaches of the band signal still easily visible.

The consistency of observed leakage across different charger hardware indicates that the issue is not isolated to a single implementation; supporting the claim that the design choices in CCS make a wireless side-channel for the PLC communication a systemic problem.

5.9.3 Message Recovery

With such a clear channel, message recovery proved highly successful, with hundreds of complete messages captured even in short sessions. In the best case, at site I, 100.0% of received messages had correct CRC32 checksums, more surprisingly 91.8% were still received when the antenna was located in the next parking bay. Reception rates were broadly correlated with raw SNR and BW, with improvements to either benefiting the performance. However this was not universal, as the very




Figure 5.10: Eavesdropping from the next parking bay (site G), more than 4 metres away on the other side to the charging cable. In this arrangement 91.8% of messages were received successfully.

poor performance at sites **B** and **K** shows. Site **B** showed poor results despite far higher SNR and BW than Site **K**. Reception performance is broken down by location in Table 5.2, with the lowest minimum and highest maximum for each location highlighted in bold. Without ground-truth for the number of messages sent by each party, we cannot determine the number of messages missed entirely (only those received with errors), although the only unreported messages would be those that did not even trigger the packet detection algorithm (see Appendix A). Examining Frame Control headers showed that traffic was observed bidirectionally between vehicle and charger in all but two cases.

As charging stations, at least in public, are busy venues, we tested whether multiple simultaneous charging sessions caused interference that affected the wireless channel quality. Two vehicles (a Jaguar I-PACE and a VW e-Golf) charged simultaneously in 5 charging sessions at 2 locations, one of which is shown in



Figure 5.11: Two vehicles charging simultaneously. With the eavesdropper between the two vehicles 42.5% of messages were received successfully, including the NMK key establishment for both vehicles.

Figure 5.11. In each case, one vehicle initiated charging first and then the second did so. The eavesdropper's antenna was located between the two vehicles and attempted to listen to both. In all cases, the eavesdropper was able to listen to traffic from both vehicles, albeit with varying success. At worst, 24.3% of messages were received with correct CRC32, at best 94.8% (mean 59.7%).

5.10 Security Analysis

5.10.1 Unencrypted Communications

Where our testing campaign captured session initialisation, we were able to examine the NMK exchange to form a network. In line with the DIN 70121 and ISO 15118 standards, every SLAC interaction we observed operated in insecure form. As such, the NMK was delivered in plaintext and the only barrier to acquiring it was receiving the message intact. We intercepted the `CM_SLAC_MATCH.CNF` message in 31

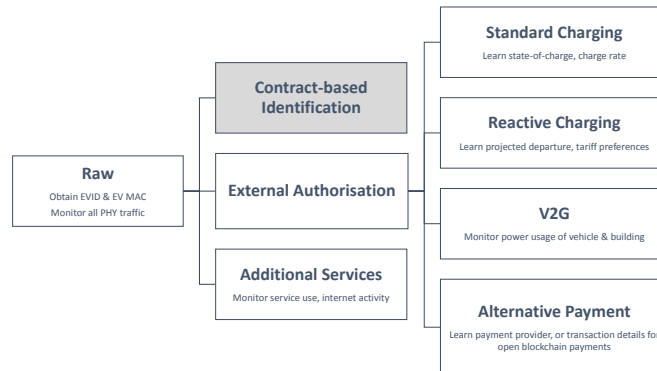


Figure 5.12: Tree diagram indicating the potential data available under a range of communication scenarios.

cases and acquired the NMK. When testing two vehicles side-by-side, in 4 sessions the NMK was extracted for one vehicle, while in one session both NMK values were extracted. In 9 cases, the subsequent `CM_GET_KEY.CNF` message was also recovered to obtain the ephemeral NEK and permit *passive decryption of physical-layer traffic*.

Examining compromised sessions, we saw the expected behaviour as the vehicle and charger established a network, the vehicle undertook the discovery protocol to find a charge controller and the two established a TCP connection. *No TLS tunnel was established in any charging session we observed*, leaving the high-level protocols exposed. In DIN 70121, or in ISO 15118 where external authorisation is used, the use of TLS is optional. Yet its complete absence from any vehicle or charger came as a surprise, especially given the charging locations were all public.

As a result we confirm that a passive attacker can wirelessly monitor all traffic at the PHY layer and that this ability *results from standards-compliant behaviour*, suggesting it is persistent. Likewise, the absence of TLS means charging data is also left in the clear. We discuss this situation and its implications in Section 5.11.

5.10.2 Private Data

Figure 5.12 provides a breakdown of potential data available when eavesdropping, under various charging conditions or in the presence of different services. The PHY-layer traffic is always available and permits access to any higher-level communication, such as charging or internet access, that does not take additional steps to secure itself.

Two unique identifiers for the vehicle are also available: its EV ID and its MAC address. These identifiers are persistent for the entire lifetime of the vehicle, including between owners, and are globally unique. They have been noted as personal data in previous privacy studies [171] and are covered by the European Union’s GDPR as data that can be easily combined with other sources to identify an individual.

With contract-based billing, we do not expect charging traffic to be available, as TLS is always required in this case. Nevertheless, this currently leaves the majority of charging traffic in the clear at public locations, although these are likely to be the earliest adopters of contract-based billing (or some alternative). The long-term omission of TLS at private locations is of greater concern. Indeed it is in this case that there is more potential for behavioural profiling, due to the vehicle staying far longer at the user’s home or workplace and with the emerging Reactive Charging and V2G systems far more beneficial to them there. The introduction of ‘Vehicle-to-Home’ capabilities, for instance, is prioritised for introduction as early as 2020 by the CCS standards body [199]. Resulting indicators of the user’s day-to-day behaviour such as the vehicle’s state-of-charge and projected departure time are contained within normal charging traffic, while reverse power flow data in a V2G system yields insights into the power usage of the building.

In addition to internet access for in-vehicle entertainment systems, third-party apps and alternative payment networks, the traffic of any local services would also be available at public locations, as would smart home integration traffic in private ones.

5.10.3 Charging Attacks

A reliable eavesdropping capability presents a range of opportunities for an attacker, both immediate and longer-term in their impact. We consider here a selection of potential attacks using these techniques. Although we did not perform the attacks against public chargers, we describe how they would be conducted.

AutoCharge Extant AutoCharge systems, such as one operating in production across a 60-location network in three European countries [164] are at particular risk from wireless eavesdropping. The use of the vehicle’s charge-controller MAC address for billing identification [163, 165], while highly questionable from a purely-security standpoint, was undertaken for compatibility and convenience benefits (and has been lauded as such by customers). What may be an acceptable trade-off when physical interference is required to extract the values, is far less so when this can be done from another vehicle without any observable signs. The identifiers of the vehicles are shown partially-masked below (none is a customer of an AutoCharge system):

Vehicle	MAC
BMW i3	f0:7f:0c:02:●●:●●
VW e-Golf	00:7d:fa:01:●●:●●
Jaguar I-PACE	00:1a:37:70:●●:●●

We were able to obtain the identifiers in 41 cases (76%)¹¹ from a variety of locations including the two-car arrangement shown in Figure 5.11. Here the identifiers for both vehicles were acquired from the same antenna position, suggesting that an attacker could simply park next to a charging station and collect identifiers as other users arrive subsequently providing them¹² in order to obtain free charging on another user’s account. As the charging spots are operated by a single provider, the attacker can be confident of targeting valid customers.

User Tracking In the simplest attack, charging sessions are linked by monitoring a number of busy public chargers for the appearance of vehicle identifiers. From time-of-day, charge duration and location information, behavioural profiles can be inferred. The invasiveness of the attack increases where the attacker is able to match a vehicle identity to other data. Popular charger-sharing schemes [202] allow anyone to register their home or business charger as a public site; any user booking to charge can then be associated with their vehicle identifier and tracked at any monitored station. Monitoring a charger near a sensitive event such as a

¹¹31 cases from SLAC initialisation messages and 10 more from network management messages

¹²Typically updating the MAC setting using `open-plc-utils` [200] and a serial debug port over UART or SPI [201]

union meeting, protest gathering or compromising night-spot would reveal more personal information about an individual's habits.

With a wireless attack, a wardriving approach also allows an attacker to associate a vehicle with a street address. This could easily be conducted by a delivery driver or postal worker as they visit properties regularly. Known MAC allocations to manufacturers also provide a coarse-grained indication of the vehicle, such as identifying expensive vehicles and then determining when they have been left in a public car park, or indeed when their owner is out of the house.

5.11 Lessons Learnt

The refinement of EV charging systems is still ongoing. In light of our observations, we have distilled a set of security lessons that can improve existing and future designs.

5.11.1 Wireless Threats

The most notable finding here is that the design of CCS communication allows a wireless attacker to observe it at a distance without prior interaction or tampering. In this case the attack was entirely passive, but has similar implications for the potential of active attacks that would currently be far more invasive. As in-vehicle wireless systems have been plagued by attacks in recent years, our results indicate that a testing model which considers emissions security as well as unwanted interference is crucial in future development.

5.11.2 Reliance on a Non-Existent PKI

While DIN 70121 does not attempt to provide security, the ISO 15118 security model relies on the existence of a complex PKI to underpin TLS at the Transport Layer and XML Security for external message values at the Application Layer. The merits of that infrastructure are an ongoing topic of academic study [146, 147, 166], but its complexity also presents a more practical problem. At the time of

writing, no widespread ISO 15118 PKI is deployed. While small-scale pilots have been attempted, there is still open debate about provision of the infrastructure and the authors are aware of public proposals from three different commercial entities to provide transaction brokerage and act as the Root Certification Authority [203]. There is even disagreement about the model the PKI will take; whether it will derive from a single root of trust, a consortium of trusted entities or some more open model [204]. Meanwhile the competing pressures to provide new functionality remain, spurring alternative solutions such as AutoCharge and encouraging service development without underlying security provision.

Even once a PKI is operating for public chargers in large charging networks, it remains unclear to what extent private units in individual homes or offices will benefit. A capacity for self-signed contract certificates to be manually installed into vehicles by users does exist, but unless contract-based billing is used ISO 15118 exempts charging installations from any security requirements; instead relying on the physical security of the location and cabling — which we have demonstrated to be insufficient. Manufacturer choices (and indeed user willingness) will determine whether private chargers can enjoy these security benefits.

It is important therefore to provide at least some security implementation that is decoupled from the need for access to a PKI. We discuss such an approach in Section 5.13.

5.11.3 Available PHY Security Disabled

The HomePlug GreenPHY (HPGP) PLC technology supports a *Secure SLAC* mode that protects the pairing and NMK distribution process, but this is disabled by specification in the ISO 15118 standard, relying instead on TLS for all security properties. While this can meet the charging use cases outlined in that standard, it ignores an opportunity for a pervasive security baseline, despite proposing the communication channel for general use. History has often shown that leaving security to individual developers atop insecure platforms produces widespread security problems, even more so when the channel is considered physically private.

5.12 Potential for Active Attacks

While the work in this chapter has focused on a passive attacker, we identified in Section 5.5 that the propensity for wireless signal emission was paired with a susceptibility to signal intrusion. Having observed the emissions phenomenon in real-world deployments and comprehensively demonstrated its use to a passive attacker, we now consider the potential for active attacks that exploit signal injection.

We stress that we have only conducted signal intrusion under lab conditions and have neither tested it in real settings, nor extensively documented the quality of the channel. Nevertheless we believe it represents a realistic attack vector for deployed systems. This is based in part on our observations, but also on elements of the ISO 15118 and CCS design. We recollect that the SLAC protocol exists specifically to avoid vehicles and chargers pairing incorrectly due to crosstalk and that it is expected that several PLC stations may participate in a given protocol run, with only the least-attenuated matching against one another. Crosstalk is not necessarily due to radiation, as it can occur conductively as well, but combined with our observations we consider it likely to occur radiatively.

An active attacker can mount more serious attacks against the charging system. By combining passive and active efforts, the attacker obtains a Dolev-Yao position on the network, as well as the ability to operate below the network layer. Attacks yield the potential for denial-of-service, network intrusion or charging fraud. We identify the following:

Jamming The attacker jams the PLC communication, thereby disrupting the charging communication atop it. The outcome is that charging stops after a short period and does not recommence until jamming ceases. The attacker can thereby deny one or more users charging over a period of their choosing. The attacker can optionally implement a more insidious selective-jamming strategy in which they both observe messages and reactively apply jamming when desired, such that only certain messages are jammed (e.g., based on source, destination or metadata values) [wilhelm2011wifire].

Message Injection The attacker injects PHY-layer PLC messages of their choosing. They can always interact using control messages that are unencrypted and unauthenticated, such as SLAC messages. If they obtain keys for encrypted communication then they can operate as a full network participant.

Relay The attacker redirects flows of traffic between participants. In particular they redirect the traffic from a victim vehicle to a charger that they have access to, such that the victim authorises billing for a charge session that the attacker then receives.

5.12.1 Jamming in detail

The attacker emits a jamming waveform that disrupts the PLC communication carrying the charging protocol. The power required to achieve this will depend heavily on the distance, channel and coupling quality. However, as the signal-to-noise ratio of the genuine signal is finite, for some level of jammer power the attacker becomes successful¹³.

Once jamming is successful, charging control is no longer possible as no messages can be exchanged. According to the ISO 15118 specification [168], under these circumstances either a Message Timeout (at the vehicle side) or a Sequence Timeout (at the charger side) will occur and the V2G Session will be interrupted. This will in turn terminate the communication and the TLS/TCP connection used for it. Behaviour in this situation can vary depending on the use case, meaning that successful interruption depends primarily on whether the charging is AC or DC and whether automated billing is being used.

For DC charging with automated billing (the heavily promoted method for public charging), charging cannot continue without further communication. The communication is required for charging control, charge schedule/tariff updates and continued signing of energy receipts by the vehicle. As such, the interruption of communication will initiate a state transition in the low-level IEC 61851 control

¹³In our lab testing, it was possible to disrupt communication at distances of a few metres by emitting white noise using an off-the-shelf SDR.




Figure 5.13: Theorised PHY-layer relay attack on ISO 15118 Plug & Charge.

system – thus terminating the power delivery. The vehicle and charger will enter an ‘Unmatched’ state for ISO 15118 communication and an X1 state for IEC 61851 signalling, meaning that charge initialisation can take place again if communication can be restored.

It is thus then under the attacker’s control whether to cease jamming and allow charging to recommence, or to continue jamming and prevent future attempts at

communication. The attacker can use this ability to mount a denial-of-service on vehicle charging; either to strand a driver whose vehicle has a low battery, or to damage the revenue and reputation of a charge-point operator. Were the attack conducted at scale, it may be possible to disconnect a large number of vehicle-charging sessions at once. As charging sessions are substantial power draws¹⁴, the aggregate reduction may have impact on the power grid.

An attacker can further implement a selective jamming strategy if desired, in which only certain messages are jammed (e.g., based on source, destination or meta-data values). They can then subvert interactions by deleting or modifying messages.

5.12.2 Message injection in detail

The attacker can emit a suitable waveform to inject arbitrary messages at the physical layer. Low-level messages for network control and initialisation, particularly those for the SLAC protocol, are unencrypted and unauthenticated. If the attacker employs the eavesdropping approach detailed to recover the NMK and NEK then they can join the network as a full participant and are only restricted by any cryptographic encapsulation applied by other parties.

The attacker thereby establishes a presence on the network, as if they were connected directly. As well as any interference with charging communication, they are free to perform reconnaissance of other network services or conduct the range of standard IP network attacks that are known. Experimental work has suggested that commonplace services such as SSH, Telnet or web-based control panels are exposed on some real-world chargers [205], and thus accessible wirelessly.

5.12.3 Relay in detail

This complex attack is illustrated in Figure 5.13. It assumes that the victim vehicle is using automated billing, such as ISO 15118 ‘Plug & Charge’. The attacker has modified their vehicle to offer control over its connection to a charger, as well the

¹⁴Charging stations delivering 50kW and 150kW deployed in the UK at time of writing, stations with 350kW outputs in mainland Europe.

control they can exert over other vehicle and chargers wirelessly. They position their vehicle at a charger and connect the cable, but do not initiate charging immediately.

The attacker then jams the victim vehicle's communication to end the existing charging session (if any) and ensure a new initialisation can take place, before isolating the vehicle's network traffic using one of a number of techniques:

- Subverting SDP
- Subverting SLAC protocol
- Selective jamming

In subverting the Session Discovery Protocol (SDP), the attacker responds to a vehicle's initial discovery message faster than the genuine charge controller, informing it of a false address to which it will try to connect. In subverting the SLAC protocol, the attacker masquerades as a charging station and falsifies attenuation reports in Step 5 such that they are selected by the vehicle. In selective jamming, the attacker jams only messages from the charge controller to the vehicle by reading the packet metadata.

The first two methods are more appealing for an attacker, where possible, as they avoid the vehicle's traffic being recognised by the charger it is actually connected to. However, both are restricted if secure SLAC is enabled and the NMK is not known. The third method is the most difficult, but is immune to the security mechanisms.

With whatever method the attacker chooses, they then commence relaying. Messages observed from the victim vehicle are relayed to the attacker's charger. Messages received from the attacker's charger are injected to the victim vehicle. The vehicle dutifully conducts the initialisation process and authorises a charging session, however the power is delivered from the attacker's charger and replenishes their vehicle's battery instead.

This attack exploits the key problem of being unable to tie the logical state (charge authorisation) to the physical state (location of the vehicle). A variant that considered only the SLAC subversion approach was suggested in [146]. However, the use of Secure SLAC would defeat this variant and its widespread deployment

is expected to coincide with the introduction of automated billing. Our proposed physical-layer approach using selective jamming is a novel attack that we believe to be effective against even a fully-secured ISO 15118 implementation.

5.13 Countermeasures

5.13.1 Protocol Changes

At a network level, we have argued for the use of the available HPGP security mechanisms above, but note that in their present form they are still reliant on a PKI to function. In addition the HPGP key distribution behaviour itself introduces an unnecessary risk of interception. Whether the SLAC protocol operates in its secure mode or not, it is still unilateral: the charger generates a network key and then provides it to the vehicle. However, the SLAC process is typically implemented in software by the same devices that undertake the higher-level communication, including possible TLS sessions, and as such require the capabilities for an Elliptic Curve Diffie-Hellman key derivation for AES128 [168].

We propose additional steps in the SLAC initialisation, as a fallback to provide confidentiality from the MAC-layer upwards in the event that PKI access is unavailable. Figure 5.14 shows the modified protocol. Upon receiving a network match request, the charger generates an Elliptic-Curve key-pair (d_C, Q_C) and instructs the vehicle to commence a key exchange, along with Q_C . If the vehicle also supports the protocol then it generates (d_V, Q_V) and responds with Q_V . The derived key becomes the new NMK and the charger blanks the NMK field in the subsequent `CM_SLAC_MATCH.CNF` message. If the vehicle does not support the protocol then the unrecognised message will be dropped. The charger maintains a timeout counter after step 6.1 and, upon expiry, falls back to the existing protocol's step 7.

While it cannot provide authentication and therefore cannot mitigate man-in-the-middle attacks, the threat of passive eavesdropping is eliminated using this approach. By building only on existing functionality, the protocol is deployable in existing vehicles as well as new ones.

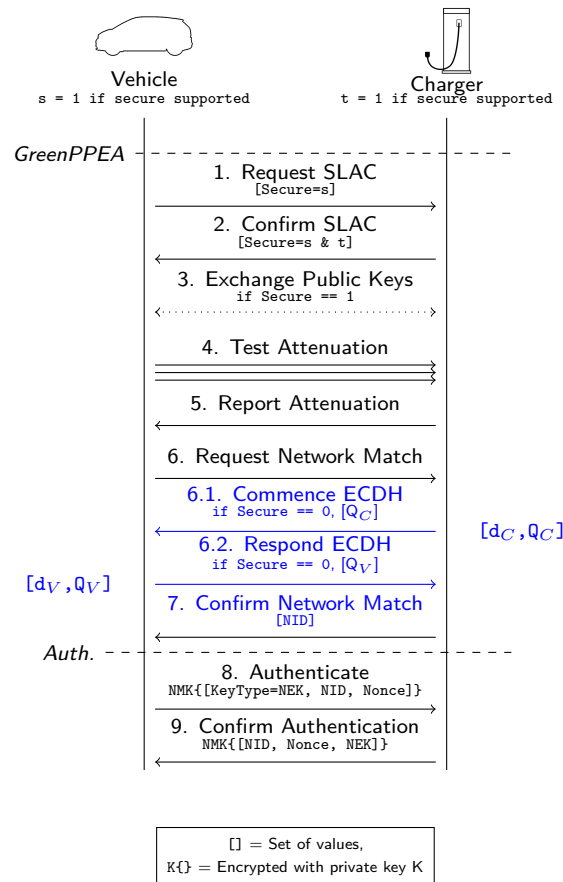


Figure 5.14: The modified SLAC network establishment. Steps 6.1 and 6.2 are new, while step 7 has been modified.

5.13.2 Equipment Changes

To mitigate the unintended wireless channel, familiar emissions security mechanisms such as chokes or shielding can be applied to reduce leakage [174]. Some proposals for future, high-power chargers include liquid-cooled charging cables and we would expect this to attenuate the signals if the cooling jacket wraps the communication lines as well as the power-delivery ones. This would not eliminate emissions from the vehicle or charger circuitry however, nor is it likely to exist in smaller, private chargers.

Although hardware modifications for existing systems are costly and time-consuming, it seems implausible that these issues were not foreseen during the design. The absence of physical mitigations is curious. For such short cable runs, shielding seems neither too costly nor too mechanically debilitating. Effective

differential signalling also does not seem beyond the realms of possibility, either using the proximity pilot as the second line, or introducing a new one. Even the use of PLC at all is puzzling, not in comparison to CAN-Bus (which has its own raft of security issues), but in comparison to other commonplace networking technologies with similar functionality and far-reduced emissions problems.

5.13.3 Physical-Layer Security

We note that physical-layer security techniques have the potential for direct application in this context. While they, perhaps ironically, cannot stop the purely-physical attack of jamming, they can frustrate network level attacks; particularly the relay attack.

Firstly, distance-bounding techniques could limit the attacker's range to be that of the victim vehicle — the length of the charging cable. Doing so would make an attack at least far less convenient, and likely easier to detect.

Secondly, techniques involving measurement of the communication channel itself seem promising. The attack is fundamentally one of a change of communication medium; allowing an attacker to eavesdrop and interact without having a direct connection. The behaviour of this wireless leakage channel is likely to be markedly different from the constrained wireline channel in the cable. Existing techniques to derive shared keys from the channel characteristics [25] or detection mechanisms based on the physical properties of received signals seem likely to have use respectively in the initialisation process and preventing message injection.

5.14 Conclusion

We have demonstrated that use of powerline communication in EV charging and the design of the CCS standard lead to an unusually high-quality, unintentional wireless channel. Although conditions vary substantially between sites, for eavesdropping we achieved a peak successful recovery rate of 100% in one case and could intercept traffic several metres from the target, in a different parking bay, with a rate of 91.8%. We showed how a series of further design choices allow recovery of network

keys and passive monitoring of all traffic in plaintext. We also describe cases in which even this fledgling system is already being misused to foster convenience at the expense of security — which validates concerns that the positioning of the system as a platform for other services only accentuates further the risks associated with any vulnerabilities.

Deployed systems reliably demonstrate wireless leakage on a scale similar to PLC deployments in domestic environments. This is despite a far more controlled wiring arrangement, suggesting that leakage is exacerbated by the coupling mechanisms used in the standardised design. The leakage effects are similar to those observed in lab conditions, where propensity for signal emission was found to also pair with susceptibility to signal intrusion. As such, we have described potential vulnerabilities to active attacks that pose more significant challenges for the system than passive attacks alone. Future work should test the effectiveness of signal intrusion, and the described active attacks, against real deployments.

Given that PLC was one option among many as a communication technology and that the risk of leakage was known, these observations suggest that physical-layer security was not well-considered in the design of the system. We suggest short-term countermeasures, but also note that the judicious use of physical properties still appears to offer the potential for securing the system better in the long term.

6

Conclusion

6.1 Collected Results

In the previous chapters, systems have been presented that exploit physical properties to enhance the security of a number of cyber-physical systems.

A secure location verification system (Ch. 2) was described that exploits the core physical-layer security property of radio propagation time and the practical advantages associated with moving receivers. The system was shown to display high ($>97\%$) accuracy in simulation with realistic parameters, and be amenable to deployment without modifying the original system.

A drone detection system (Ch. 3) was presented that derives information about movement and environment from only a single widely-available radio metric; using this to infer device type. An implementation using consumer-grade equipment was shown to be $>98\%$ accurate even when averaged across all receiver locations in a realistic in-home deployment.

A system for identifying malicious powerline networks (Ch. 4) was described that makes use of unintentional electromagnetic emissions instead of conducted signals, exploiting the advantages of wireless propagation for this purpose. Even with simple signal-processing methods, the system was shown to display perfect

accuracy within the same room and to still be effective (>74% accurate) for targets on a different floor, which could not be detected by conducted means.

An existing, large-scale vehicle-charging system (Ch. 5) was also tested with an eavesdropping attack exploiting the same unintentional emissions phenomenon seen with consumer powerline networks. The emissions were found to be an inherent feature of wireline OFDM and exacerbated by poor design choices. The eavesdropping attack was found to be effective in a range of real settings, with >91% of messages recovered correctly from a charging vehicle when the attack was conducted from a nearby parking bay. Use of the attack indicated immediate privacy and security concerns with deployed charging infrastructure.

In each case, the systems or attacks presented have been practically realised at low cost using only COTS equipment.

6.2 Conclusions

The challenges of securely modelling the physical world, or monitoring a security-relevant physical property, have been seen in a range of real systems; from legacy air-traffic monitoring designs developed before physical-layer attacks were widely-accessible, to brand-new vehicle-charging systems developed with knowledge of, but not regard for, the capabilities of a modern adversary.

The repeated successful development of security systems has confirmed the claim that it is possible to exploit physical-layer features for defensive means. Moreover, the use throughout of commodity equipment has demonstrated that the growing accessibility of complex sensing and actuation technologies is indeed providing a means to implement physical-layer security controls in cyber-physical systems, as well as attacks.

It has been demonstrated that physical-layer features can provide substantial defensive security benefits if used judiciously. Information about the physical world was shown to be useful either when it had been measured directly or when it had been inferred. Indeed, it was seen that a surprising wealth of information can be derived from a single metric when its behaviour is well-characterised. It was also

shown that creative use of unintentional effects and side channels can unlock new realms of system behaviour. The potential for physical factors to impose prodigious challenges for an attacker has been shown throughout; where drastic changes to the threat model and more resources have been required to defeat the system.

The defensive use of physical phenomena has only been seen insofar as detection is concerned. Primarily this has been a matter of applicability, as there was no scope to modify the systems involved and so detection of malicious activity was the only means to integrate additional security. While the results presented herein cannot therefore be used to make claims outside the realm of detection, there does not seem to be a clear reason to doubt that other systems are possible and the author remains confident that future work will bear this out.

It has also been confirmed that it is necessary to make use of physical-layer factors to secure cyber-physical systems¹. Ignorance of physical-layer attacks has been shown to leave even a modern system, with a sophisticated security regime, open to the compromise of core functionality. The substantial application of cryptography is insufficient when the trusted computing base includes assumptions upon the physical world that can be neither fully-verified nor fully-controlled. The ways in which the physical layer of the system could be abused must be considered in the development of new systems.

It remains challenging to develop good physical-layer security systems; ultimately there is always some threat model which can overcome any security system. There is also an enormous range of physical phenomena and myriad ways in which they can conspire to undo a system's design — hence the great efforts over the years to minimise the trusted computing base. But the dual perspective is that this also yields a great range of opportunities to turn those physical phenomena to a system's advantage. No implementation will be perfect, but even as theoretical work seeks to develop impenetrable protocols using unbreakable physical bounds, plenty of scope still exists to meaningfully enhance security with practical means. Indeed, codifying

¹As indeed the name rather suggests.

the physical dependencies of a system in order to perform formal modelling also yields the information with which to make engineering choices about how to defend them.

The many ways in which physical factors can be used to enhance security offer potential benefits that should not lightly be ignored.

6.3 Final Remarks

While the field referred to as physical-layer security focuses enormously on wireless communication, the use of physical features in securing systems has a long history in other fields distributed across security study. While the security capabilities that can be constructed by integrating particular physical factors varies enormously with the factors themselves, the approach to integrating them does not. As such, it seems that the field of physical-layer security may² come to encompass the general use of physical properties as primitives in security systems. In the author's opinion it should.

Should general methods by which to tie physical and logical states together be discovered, they would be powerful tools in securing systems of all kinds. Their pursuit seems an admirable goal, albeit one whose possibility is not clear.

There is also still far more practical work left to be done in protecting real systems from physical-layer attacks. This too is an important endeavour. The proliferation of cyber-physical systems and their involvement in daily life is only growing — it would be great if they remained secure.

²Exercising here the prerogative at the end of a long work of sober factual reflection, to engage in a bit of wild speculation.

Appendices

A

HomePlug GreenPHY Receiver

In this appendix we describe in detail the eavesdropping tool implementation used in Chapter 5. As noted in Section 5.7, the tool is effectively a modified receiver design, although newly-implemented entirely in software. Since HomePlug GreenPHY (HPGP) [162] is an orthogonal frequency-division multiplexing (OFDM) technology, many elements of the tool structure are similar to a Wi-Fi receiver.

Raw signals are first collected using a suitable capture device. A Rigol DSA-2302A oscilloscope was used in our testbed arrangement, as can be seen in Figure A.1. Even here the radiated emissions were easily observed; the yellow line in the figure represents the conducted signal, while the blue line is the radiated signal received by a short random-wire antenna. Although the distance shown here is very short, we were still able to observe the signal from the other side of the lab, several metres away. We later employed software-defined radios for signal capture, for their ability to receive and stream a captured signal in real time.

The captured signal is filtered in the frequency domain, benefiting from knowledge of the active regions of the HPGP band and the ability to survey initially an individual site's leakage before eavesdropping in earnest. A sharp-edged digital filter is used to remove regions with notable interference ingress or where channel gain is so low as to provide no useful information. The signal is then resampled into the HPGP native timebase of 75MHz.




Figure A.1: HomePlug AV adaptors communicating across a short wire. Conducted signals and radiated emissions can be seen on the oscilloscope (top in yellow and bottom in blue, respectively).

Frame Detection and Time Alignment With the signal suitably pre-filtered and digitised, the PPDU's are detected using a *Double Sliding-Window* power detector; a design that accurately identifies the rise in power that accompanies the start of a packet. The detector calculates the power of the incoming signal and maintains two windows A and B of equal length L that are arranged with a time lag such that calculated power levels are included in window A at time t , subsequently passing out of window A and into window B at time $t + L$ and out of the detector entirely at time $t + 2L$. At each sample, the power in each window is updated and the total power in A is divided by that in B. This configuration causes the output signal to spike quickly on increases in power levels, while remaining stable at equal power levels (i.e., prior to or during a frame). By selecting an appropriate value of L (based on the frame's structure), transient noise can be prevented from triggering a frame.

OFDM requires precise time synchronisation in order to demodulate correctly. We performed this by correlating the entire preamble against a template, which provided sample-accurate alignment.

CPO, SCO & Channel Estimation In practice, a transmitter and receiver in an OFDM system will have neither precisely-aligned oscillators nor synchronised sample clocks, leading to Carrier Phase Offset (CPO) and Sampling Clock Offset (SCO). CPO causes an apparent frequency offset for the entire received signal, meaning that the frequency-domain representation exhibits a phase rotation. Meanwhile, SCO leads to an apparent phase drift across subcarriers in the frequency domain. Both phenomena hamper demodulation and must be corrected beforehand. Channel estimation is also crucial to successful reception; assessing the gain and phase alterations that have been experienced by the signal due to the propagation environment.

Our receiver estimates the CPO using a method derived from Bloessel et. al.'s work; estimating the CPO using the seven full-amplitude SYNCP preamble symbols in place of the Wi-Fi short-training sequence [206] (omitting the initial 192 as they have been windowed in symbol shaping):

$$cpo_{est} = \frac{1}{384} Arg \left(\sum_{i=0}^{7 \cdot 384} x[i] \bar{x}[i + 384] \right)$$

where x is the received signal samples.

From the extracted section of the preamble, complex samples are multiplied with the conjugate of the same sample in the next SYNCP block. This produces an estimate of the phase progression introduced between those SYNCP blocks by the mismatch between transmitter and receiver (plus noise). Dividing through by the length of the SYNCP block gives an estimate of the phase offset per sample. The length of the sequence (2688) and the number of repetitions (7) permit an accurate CPO estimate. The per-sample CPO estimate can then be used to correct the remainder of the captured signal.

$$x[i] \leftarrow x[i] \cdot e^{-jcpo_{est}i}$$

As the estimated CPO will not precisely match the actual CPO, ongoing correction is applied to each received symbol by estimating the CPO between the cyclic prefix and the symbol tail, with a suitable correction being applied over that symbol.

$$c_{PO_{estcp}} = \frac{1}{3072} \text{Arg} \left(\sum_{i=0}^{GI} x[i] \bar{x}[i + 3072] \right)$$

where GI is the guard interval (with four values depending on the symbol and system settings).

The channel estimation is performed in the frequency domain, by comparing the received preamble symbols to a locally-computed template. HPGP provides no pilot symbols so all estimation must be performed from the preamble and maintained across the PPDU. The results for each preamble symbol are averaged and a channel estimate from the active preamble subcarriers computed. From this a channel estimate for the full channel is derived by interpolation, while the SCO is estimated from the slope of the phase differences in the channel estimate. As the CPO and SCO are due to hardware imperfections in the transmitter and receiver, rather than channel properties, estimates are maintained between received PPDUs by way of a moving average. The channel estimate, by contrast, is discarded after a PPDU has been received.

Demodulation Demodulation takes place in the frequency domain (via a 3072-point DFT), after the removal of the cyclic prefix for the symbol and correction for the channel effects at each subcarrier. As HPGP uses QPSK modulation, the receiver compares the measured value for the subcarrier in the in-phase and quadrature channels to the nominal values and estimates, under an additive white Gaussian noise assumption, the likelihood of the transmitted value having been a 0 or 1 bit. These probabilities are expressed as a ratio, the Log Likelihood Ratio (LLR) and then scaled according to the gain for the subcarrier in the channel estimate, such that the uncertainty inherent in weakly-received subcarriers is represented.

Post Processing Demodulated soft bits are combined by averaging to benefit from HPGP's redundancy schemes. They are then rearranged in read-by-row-write-by-column fashion to undo the channel interleaving process.

The FEC decoding is applied to produce hard decisions about the bit values. HPGP uses an unpunctured Turbo code with two systematic, rate $\frac{2}{3}$ constituent codes. Each pair of input bits (i, j) produces a codeword (i, j, p, q) , where p and q are parity bits, p from the in-order input and q from an interleaved input.

Finally, the bits are unscrambled by XORing with the same generator polynomial used in the transmitter to recover the original sequence.

The CRC checks are computed over the received bits to determine if the contents have been received successfully, however the PHY-layer bits are delivered to the higher layers irrespective as even messages containing errors may provide useful information.

Each stage of the receiver is configurable with a wide range of parameters. In particular, the power threshold to trigger PPDU capture, the frequency-domain filtering, the initial CPO estimate and the estimated noise variance for demodulation all permit tailoring the receiver to a given scenario.

Considering the emissions as a wireless channel, the simple modulation and redundancy in HomePlug GreenPHY's robust ("ROBO") transmission modes mean the attacker need not match the channel characteristics of any particular receiver; they need only to receive the transmissions with enough of the signal intact. Specifically, the attacker requires a positive signal-to-noise ratio (SNR) over some fraction B of the transmitted bandwidth. The selection of B depends upon the transmissions mode in use and the effectiveness of any error-correction mechanisms, however for a rough estimate the level of redundancy can be used. Thus for MINI_ROBO, STD_ROBO, HS_ROBO B can be taken as 5.2MHz, 6.5MHz, 13MHz ($\frac{1}{5}$, $\frac{1}{4}$ and $\frac{1}{2}$ of the 26 MHz HPGP bandwidth) respectively.

References

- [1] Theodore S Rappaport et al. *Wireless communications: principles and practice*. Vol. 2. prentice hall PTR New Jersey, 1996.
- [2] Steven W. Smith. *The Scientist and Engineer's Guide to Digital Signal Processing*. 2nd ed. Available online at: <https://www.dspguide.com/>. California Technical Publishing, 1999.
- [3] Richard G. Lyons. *Understanding Digital Signal Processing*. 3rd ed. Prentice Hall, 2010.
- [4] Harry Nyquist. "Certain topics in telegraph transmission theory". In: *Transactions of the American Institute of Electrical Engineers* 47.2 (1928), pp. 617–644.
- [5] David Declercq, Marc Fossorier, and Ezio Biglieri. *Channel Coding: Theory, Algorithms, and Applications: Academic Press Library in Mobile and Wireless Communications*. Academic Press, 2014.
- [6] James W Cooley and John W Tukey. "An algorithm for the machine calculation of complex Fourier series". In: *Mathematics of Computation* 19.90 (1965), pp. 297–301.
- [7] Luigi Federico Menabrea and Ada Lovelace. *Sketch of the analytical engine invented by Charles Babbage*. 1842.
- [8] Robert D. Farley. *NSA Oral History Interview: Dr. Howard Campaigne*. 1983. URL: <https://www.nsa.gov/Portals/70/documents/news-features/declassified-documents/oral-history-interviews/nsa-oh-14-83-campaigne.pdf>.
- [9] Paul Benioff. "The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines". In: *Journal of statistical physics* 22.5 (1980), pp. 563–591.
- [10] Paul C Kocher. "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems". In: *Annual International Cryptology Conference*. Springer. 1996, pp. 104–113.
- [11] Eli Biham and Adi Shamir. "Power analysis of the key scheduling of the AES candidates". In: *Proceedings of the second AES Candidate Conference*. 1999, pp. 115–121.
- [12] Daniel Genkin, Adi Shamir, and Eran Tromer. "RSA key extraction via low-bandwidth acoustic cryptanalysis". In: *Annual Cryptology Conference*. Springer. 2014, pp. 444–461.
- [13] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. "Electromagnetic analysis: Concrete results". In: *International workshop on cryptographic hardware and embedded systems*. Springer. 2001, pp. 251–261.

- [14] Alessandro Barenghi et al. “Fault injection attacks on cryptographic devices: Theory, practice, and countermeasures”. In: *Proceedings of the IEEE* 100.11 (2012), pp. 3056–3076.
- [15] Yoongu Kim et al. “Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors”. In: *ACM SIGARCH Computer Architecture News*. Vol. 42. 3. IEEE Press. 2014, pp. 361–372.
- [16] Denis Foo Kune et al. “Ghost talk: Mitigating EMI signal injection attacks against analog sensors”. In: *2013 IEEE Symposium on Security and Privacy*. IEEE. 2013, pp. 145–159.
- [17] Guoming Zhang et al. “Dolphinattack: Inaudible voice commands”. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2017, pp. 103–117.
- [18] Ralph Langner. “Stuxnet: Dissecting a cyberwarfare weapon”. In: *IEEE Security & Privacy* 9.3 (2011), pp. 49–51.
- [19] Saar Drimer, Steven J Murdoch, et al. “Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks.” In: *USENIX security symposium*. Vol. 312. 2007.
- [20] Aurélien Francillon, Boris Danev, and Srdjan Capkun. “Relay attacks on passive keyless entry and start systems in modern cars”. In: *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Eidgenössische Technische Hochschule Zürich, Department of Computer Science. 2011.
- [21] Stefan Brands and David Chaum. “Distance-bounding protocols”. In: *Workshop on the Theory and Application of Cryptographic Techniques*. Springer. 1993, pp. 344–359.
- [22] David Basin et al. “Let’s get physical: Models and methods for real-world security protocols”. In: *International Conference on Theorem Proving in Higher Order Logics*. Springer. 2009, pp. 1–22.
- [23] Kasper Bonne Rasmussen and Srdjan Capkun. “Realization of RF Distance Bounding.” In: *USENIX Security Symposium*. 2010, pp. 389–402.
- [24] Gerhard P Hancke and Markus G Kuhn. “Attacks on time-of-flight distance bounding channels”. In: *Proceedings of the first ACM conference on Wireless network security*. ACM. 2008, pp. 194–202.
- [25] Chunxuan Ye et al. “Information-theoretically secret key generation for fading wireless channels”. In: *IEEE Transactions on Information Forensics and Security* 5.2 (2010), pp. 240–254.
- [26] Richard Newman et al. “HomePlug AV security mechanisms”. In: *2007 IEEE International Symposium on Power Line Communications and Its Applications*. IEEE. 2007, pp. 366–371.
- [27] David Adamy. *EW 101: A first course in electronic warfare*. Vol. 101. Artech house, 2001.
- [28] Nikolaos Karapanos et al. “Sound-proof: usable two-factor authentication based on ambient sound”. In: *24th {USENIX} Security Symposium ({USENIX} Security 15)*. 2015, pp. 483–498.

- [29] Fabiola Colone et al. “WiFi-based passive bistatic radar: Data processing schemes and experimental results”. In: *IEEE Transactions on Aerospace and Electronic Systems* 48.2 (2012), pp. 1061–1079.
- [30] Fadel Adib and Dina Katabi. *See through walls with WiFi!* Vol. 43. 4. ACM, 2013.
- [31] Marco Rocchetto and Nils Ole Tippenhauer. “Towards formal security analysis of industrial control systems”. In: *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*. ACM. 2017, pp. 114–126.
- [32] Saandeep Depatla, Lucas Buckland, and Yasamin Mostofi. “X-ray vision with only wifi power measurements using rytov wave models”. In: *IEEE Transactions on Vehicular Technology* 64.4 (2015), pp. 1376–1387.
- [33] Deepak Vasisht, Swarun Kumar, and Dina Katabi. “Decimeter-level localization with a single WiFi access point”. In: *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*. 2016, pp. 165–178.
- [34] Srdjan Capkun et al. “Secure location verification with hidden and mobile base stations”. In: *Mobile Computing, IEEE Transactions on* 7.4 (2008), pp. 470–483.
- [35] Francesco Papi et al. “Radiolocation and tracking of automatic identification system signals for maritime situational awareness”. In: *IET Radar, Sonar & Navigation* 9.5 (2014), pp. 568–580.
- [36] Jun Luo, Hersh V Shukla, and Jean-Pierre Hubaux. “Non-Interactive Location Surveying for Sensor Networks with Mobility-Differentiated ToA”. In: *INFOCOM*. 2006.
- [37] Wai Chen. *Vehicular Communications and Networks: Architectures, Protocols, Operation and Deployment*. Woodhead Publishing, 2015.
- [38] Alan Bensky. *Wireless Positioning: Technologies and Applications*. Artech House, 2016.
- [39] Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. “Experimental analysis of attacks on next generation air traffic communication”. In: *Applied Cryptography and Network Security*. Springer. 2013, pp. 253–271.
- [40] Andrei Costin and Aurélien Francillon. “Ghost in the Air (Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices”. In: *Black Hat USA* (2012).
- [41] Paul Pounds et al. “Design of a four-rotor aerial robot”. In: *Proceedings of the 2002 Australasian Conference on Robotics and Automation (ACRA 2002)*. Australian Robotics & Automation Association. 2002, pp. 145–150.
- [42] British Broadcasting Corporation. *Didcot Power Station collapse: Major search for missing*. 2016. URL: <http://www.bbc.co.uk/news/uk-england-oxfordshire-35647397>.
- [43] Federal Aviation Authority Unmanned Aircraft Systems Task Force. *FAA UAS Task Force Recommendations*. Nov. 2015. URL: https://www.faa.gov/uas/publications/media/RTFARCFinalReport_11-21-15.pdf (visited on 11/25/2015).
- [44] Civil Aviation Authority. *Regulations for Small Unmanned Aircraft*. Nov. 2015. URL: <https://www.caa.co.uk/default.aspx?catid=1995&pageid=16012> (visited on 11/25/2015).

- [45] Tom Mendelsohn. *Amazon to test drone deliveries in UK after government clears runway*. 2016. URL: <http://arstechnica.co.uk/business/2016/07/amazon-drone-delivery-trial-uk-government-cao/>.
- [46] Robert Pearce. *Presentation to the UAS COE Public Meeting*. 2014. URL: http://www.faa.gov/about/office_org/headquarters_offices/ang/offices/management/coe/media/pdf/UAS_COE_Briefing_Robert.pdf.
- [47] Christopher D Wickens et al. *The future of air traffic control: Human operators and automation*. National Academies Press, 1998.
- [48] Federal Aviation Authority. *NextGen Update 2014*. Aug. 2014. URL: <http://www.faa.gov/nextgen/media/NextGenUpdate2014.pdf> (visited on 11/26/2015).
- [49] Martin Strohmeier and Ivan Martinovic. “On Passive Data Link Layer Fingerprinting of Aircraft Transponders”. In: *Proceedings of the First ACM Workshop on Cyber-Physical Systems-Security and/or Privacy*. ACM, 2015, pp. 1–9.
- [50] Pericle Perazzo et al. “The verifier bee: A path planner for drone-based secure location verification”. In: *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2015 IEEE 16th International Symposium on a*. IEEE, 2015, pp. 1–9.
- [51] Matthias Schäfer, Vincent Lenders, and Jens Schmitt. “Secure Track Verification”. In: *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. 2015.
- [52] Nils Ole Tippenhauer et al. “On the requirements for successful GPS spoofing attacks”. In: *Proceedings of the 18th ACM conference on Computer and communications security*. ACM, 2011, pp. 75–86.
- [53] GlobalTop Technology Inc. *FGPMMOPA6H GPS Standalone Module Data Sheet*. 2011. URL: <http://www.adafruit.com/datasheets/GlobalTop-FGPMOPA6H-Datasheet-VOA.pdf>.
- [54] Matthias Schäfer et al. “Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research”. In: *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*. IPSN '14. Berlin, Germany: IEEE Press, 2014, pp. 83–94.
- [55] Centre Tecnològic de Telecomunicacions de Catalunya. *GNSS-SDR operation with a Realtek RTL2832U USB dongle DVB-T receiver*. URL: <http://gnss-sdr.org/node/50>.
- [56] Martin Strohmeier et al. “Realities and challenges of nextgen air traffic management: the case of ADS-B”. In: *IEEE Communications Magazine* 52.5 (2014), pp. 111–118.
- [57] Jean-Louis Naudin. *ThermoPilot - a Thermal Hunter Glider Drone*. URL: <http://diydrones.com/profiles/blogs/thermopilot-project-a-thermal-hunter-glider-drone>.
- [58] Flightradar24. *Live flight tracker*. 2017. URL: <https://www.flightradar24.com/>.

- [59] BI Intelligence. *Drone market shows positive outlook with strong industry growth and trends*. Accessed: 08.10.2019. 2017. URL: <https://www.businessinsider.com/drone-industry-analysis-market-trends-growth-forecasts-2017-7>.
- [60] Teal Group. *Global drone market estimated to reach \$14 billion over next decade: study*. Accessed: 08.10.2019. 2019. URL: <https://uk.reuters.com/article/us-usa-security-drones/global-drone-market-estimated-to-reach-14-billion-over-next-decade-study-idUKKCN1UC2MU>.
- [61] Federal Aviation Administration. *FAA Drone Registry Tops One Million*. Accessed: 08.10.2019. 2018. URL: <https://www.transportation.gov/briefing-room/faa-drone-registry-tops-one-million>.
- [62] Financial Times. *China's DJI targets agriculture as consumer drone sales slow*. Accessed: 08.10.2019. 2019. URL: <https://www.ft.com/content/afa5e042-4c50-11e9-bbc9-6917dce3dc62>.
- [63] BBC. *Gang who flew drones carrying drugs into prisons jailed*. Accessed: 08.10.2019. 2018. URL: <https://www.bbc.co.uk/news/uk-england-45980560>.
- [64] BBC. *Japan radioactive drone: Tokyo police arrest man*. Accessed: 08.10.2019. 2015. URL: <https://www.bbc.co.uk/news/world-asia-32465624>.
- [65] Der Spiegel. *Merkel Buzzed by Mini-Drone at Campaign Event*. Accessed: 08.10.2019. 2013. URL: <https://www.spiegel.de/international/germany/merkel-campaign-event-visited-by-mini-drone-a-922495.html>.
- [66] User 'GWA88' et al. *Gatwick Airport drone incident*. Accessed: 08.10.2019. 2018. URL: https://en.wikipedia.org/wiki/Gatwick_Airport_drone_incident.
- [67] Linzi Sheldon. *Woman terrified by drone outside her window*. KIRO 7. Accessed: 29.06.2016. 2014. URL: <http://www.kiro7.com/news/woman-terrified-drone-outside-her-window/81721261>.
- [68] WDRB. *Hillview man arrested for shooting down drone; cites right to privacy*. Accessed: 29.06.2016. 2015. URL: <http://www.wdrb.com/story/29650818/hillview-man-arrested-for-shooting-down-drone-cites-right-to-privacy>.
- [69] T Eshel. "Mobile Radar Optimized to Detect UAVs, Precision Guided Weapons". In: *Defense Update* (2013).
- [70] Fraunhofer Institute. *Detection of Small Drones with Millimeter-Wave Radar*. Accessed: 08.10.2019. 2019. URL: <https://www.fhr.fraunhofer.de/en/businessunits/security/Detection-of-small-drones-with-millimeter-wave-radar.html>.
- [71] Artem Rozantsev, Vincent Lepetit, and Pascal Fua. "Flying objects detection from a single moving camera". In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE. 2015.
- [72] Joël Busset et al. "Detection and tracking of drones using advanced acoustic cameras". In: *SPIE Security+ Defence*. International Society for Optics and Photonics. 2015.

- [73] Gabriel C Birch and Bryana L Woo. “Counter Unmanned Aerial Systems Testing: Evaluation of VIS SWIR MWIR and LWIR Passive Imagers”. In: *SANDIA REPORT SAND2017-0921* (2017).
- [74] Philip Church et al. “Aerial and surface security applications using lidar”. In: *Laser Radar Technology and Applications XXIII*. Vol. 10636. International Society for Optics and Photonics. 2018, p. 1063604.
- [75] Ellen E Case, Anne M Zelnio, and Brian D Rigling. “Low-Cost Acoustic Array for Small UAV Detection and Tracking”. In: *2008 IEEE National Aerospace and Electronics Conference*. IEEE. 2008.
- [76] Juan R Vasquez et al. “Multisensor 3D tracking for counter small unmanned air vehicles (CSUAV)”. In: *SPIE Defense and Security Symposium*. International Society for Optics and Photonics. 2008.
- [77] Sanjay K Boddhu et al. “A collaborative smartphone sensing platform for detecting and tracking hostile drones”. In: *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics. 2013.
- [78] Stephan Sand, Armin Dammann, and Christian Mensing. *Positioning in Wireless Communications Systems*. John Wiley & Sons, 2014.
- [79] Matthew Peacock and Michael N Johnstone. “Towards detection and control of civilian unmanned aerial vehicles”. In: (2013).
- [80] Ben Nassi et al. “Drones’ Cryptanalysis-Smashing Cryptography with a Flicker”. In: *IEEE Symposium on Security and Privacy (SP)*, Vol. 00. 2019, pp. 833–850.
- [81] Michigan Tech. *Robotic Falconry*. Accessed: 10.07.2016. 2016. URL: <http://me.sites.mtu.edu/rastgaar/home/news/>.
- [82] Guard From Above. *Intercepting hostile drones*. Accessed: 10.07.2016. 2016. URL: <http://guardfromabove.com/>.
- [83] Sait Murat Giray. “Anatomy of unmanned aerial vehicle hijacking with signal spoofing”. In: *Recent Advances in Space Technologies (RAST), 2013 6th International Conference on*. IEEE. 2013, pp. 795–800.
- [84] Samy Kamkar. *SkyJack*. Accessed: 10.07.2016. 2013. URL: <http://samy.pl/skyjack/>.
- [85] Daniel P Shepard et al. “Evaluation of smart grid and civilian UAV vulnerability to GPS spoofing attacks”. In: *Proceedings of the ION GNSS Meeting*. Vol. 3. 2012.
- [86] Yunmok Son et al. “Rocking drones with intentional sound noise on gyroscopic sensors”. In: *24th USENIX Security Symposium (USENIX Security 15)*. 2015, pp. 881–896.
- [87] Skylogic Research. *DJI market share: here’s exactly how rapidly it has grown in just a few years*. Accessed: 08.10.2019. 2018. URL: <http://thedronegirl.com/2018/09/18/dji-market-share/>.
- [88] Civil Aviation Authority. *The Air Navigation Order 2016 And Regulations (CAP393)*. Accessed: 25.08.2016. 2016. URL: http://publicapps.caa.co.uk/docs/33/CAP%20393_AUG2016.pdf.
- [89] Paramvir Bahl et al. “RADAR: An in-building RF-based user location and tracking system”. In: (2000).

- [90] Zhuoling Xiao et al. “Identification and mitigation of non-line-of-sight conditions using received signal strength”. In: *2013 IEEE 9th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE. 2013, pp. 667–674.
- [91] Kavitha Muthukrishnan, Berend Jan van der Zwaag, and Paul Havinga. “Inferring motion and location using WLAN RSSI”. In: *Mobile Entity Localization and Tracking in GPS-less Environments*. Springer, 2009, pp. 163–182.
- [92] Shuyu Shi et al. “Accurate location tracking from CSI-based passive device-free probabilistic fingerprinting”. In: *IEEE Transactions on Vehicular Technology* 67.6 (2018), pp. 5217–5230.
- [93] Daniel Halperin et al. “Tool release: Gathering 802.11 n traces with channel state information”. In: *ACM SIGCOMM Computer Communication Review* 41.1 (2011), pp. 53–53.
- [94] Yongsen Ma, Gang Zhou, and Shuangquan Wang. “WiFi sensing with channel state information: A survey”. In: *ACM Computing Surveys (CSUR)* 52.3 (2019), p. 46.
- [95] Matthias Schulz, Daniel Wegemer, and Matthias Hollick. *Nexmon: The C-based Firmware Patching Framework*. 2017. URL: <https://nexmon.org>.
- [96] Wilfried Gouret, Fabienne Nouvel, and Ghais El-Zein. “Powerline communication on automotive network”. In: *2007 IEEE 65th Vehicular Technology Conference-VTC2007-Spring*. IEEE. 2007, pp. 2545–2549.
- [97] Massimo Antoniali et al. “Measurements and analysis of PLC channels in a cruise ship”. In: *2011 IEEE International Symposium on Power Line Communications and Its Applications*. IEEE. 2011, pp. 102–107.
- [98] HomePlug Powerline Alliance. *HomePlug Powerline Networking Technology Hits Maturation as Global Broadband Standard*. 2016. URL: <http://www.homeplug.org/news/member-pr/398/> (visited on 12/08/2017).
- [99] Chen Wang et al. “Locating Rogue Access Point using Fine-grained Channel Information”. In: *IEEE Transactions on Mobile Computing* 16.9 (2017), pp. 2560–2573.
- [100] Richard Newman et al. “Protecting domestic power-line communications”. In: *Proceedings of the second symposium on usable privacy and security*. ACM. 2006, pp. 122–132.
- [101] Sébastien Dudek. “HomePlugAV PLC: practical attacks and backdooring”. In: *NoSuchCon*. 2015.
- [102] Alberto Pittolo and Andrea M Tonello. “Physical Layer Security in Power Line Communication Networks”. In: *Physical and Data-Link Security Techniques for Future Communication Systems*. Springer, 2016, pp. 125–144.
- [103] PA Consulting Group. *The Likelihood and Extent of Radio Frequency Interference from In-Home PLT Devices*. Tech. rep. Ofcom, June 2010.

- [104] Ben Tasker. *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> (visited on 11/22/2017).
- [105] Todd W Colvin. “Ethernet over low-voltage power line communication networks a security analysis and audit of the HomePlug 1.0 standard an auditor’s perspective”. In: *SANS GSNA Practical v 2.1* (2003).
- [106] Tim Williams. “RF emissions of powerline ethernet adapters”. In: *EMC Journal* 82 (2009), pp. 15–18. URL: http://www.elmac.co.uk/RF_Emissions_of_Powerline_Ethernet_adapters.pdf.
- [107] Brad Zarikoff and David Malone. “Experiments with radiated interference from in-home power line communication networks”. In: *Communications (ICC), 2012 IEEE International Conference on*. IEEE. 2012, pp. 3414–3418.
- [108] U.S. National Security Agency. *TEMPEST: A Signal Problem*. 1972.
- [109] Wim Van Eck. “Electromagnetic radiation from video display units: An eavesdropping risk?” In: *Computers & Security* 4.4 (1985), pp. 269–286.
- [110] Markus G Kuhn and Ross J Anderson. “Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations.” In: *Information hiding*. Vol. 1525. Springer. 1998, pp. 124–142.
- [111] Jake Longo et al. “SoC it to EM: electromagnetic side-channel attacks on a complex system-on-chip”. In: *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer. 2015, pp. 620–640.
- [112] Daniel Genkin, Itamar Pipman, and Eran Tromer. “Get your hands off my laptop: Physical side-channel key-extraction attacks on PCs”. In: *Journal of Cryptographic Engineering* 5.2 (2015), pp. 95–112.
- [113] Raheem Beyah et al. “Rogue access point detection using temporal traffic characteristics”. In: *Global Telecommunications Conference, 2004. GLOBECOM'04. IEEE*. Vol. 4. IEEE. 2004, pp. 2271–2275.
- [114] 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 (visited on 12/08/2017).
- [115] Xavier Carcelle. *Power line communications in practice*. Artech House, 2009.
- [116] Klaus Dostert. “Telecommunications over the power distribution grid—possibilities and limitations”. In: *IIR-Powerline* 6 (1997), p. 97.
- [117] TP-Link Technologies Co. *Most frequent questions about TP-Link Powerline adapters*. 2017. URL: <http://uk.tp-link.com/faq-406.html> (visited on 12/14/2017).
- [118] Simon Jary. *The best powerline adapters for 2017*. 2017. URL: <https://www.techadvisor.co.uk/test-centre/network-wifi/best-powerline-adapters-for-2017-3490638/> (visited on 11/25/2017).

- [119] Michael Himmels. *Devolvo Real World Field Tests*. 2011. URL: http://www.homeplug.org/media/filer_public/25/4f/254f6adb-096a-4913-842b-91e3775da045/devolvo_presentation.pdf (visited on 11/22/2017).
- [120] Md Mustafizur Rahman et al. “Medium access control for power line communications: an overview of the IEEE 1901 and ITU-T G.hn standards”. In: *IEEE Communications Magazine* 49.6 (2011).
- [121] devolo AG. *dLAN@ 1200+ Starter Kit Powerline*. 2017. URL: <https://www.devolvo.co.uk/article/dlan-1200-starter-kit-powerline/> (visited on 12/05/2017).
- [122] TP-Link Technologies Co. *TL-PA8010*. 2017. URL: http://www.tp-link.com/us/products/details/cat-5509_TL-PA8010-KIT.html#specifications (visited on 12/05/2017).
- [123] HomePlug Powerline Alliance. “HomePlug AV specification version 2.0”. In: *Beaverton, OR, USA* (2012).
- [124] devolo AG. *devolo dLAN 200 AVminiPCI Datasheet*. 2009. URL: <https://www.devolvo.com/products/Integrationmodules/dLAN-200-AVminiPCI/data/Product-sheet-dLAN-200-AVminiPCI-com.pdf> (visited on 11/24/2017).
- [125] TP-Link Technologies Co. *AV500 Gigabit Powerline Adapter TL-PA511*. 2011. URL: <http://static.tp-link.com/resources/document/TL-PA511.zip> (visited on 11/24/2017).
- [126] TP-Link Technologies Co. *AV200 Wireless N Powerline*. 2011. URL: http://static.tp-link.com/resources/document/TL-WPA281_V1_Datasheet.zip (visited on 11/24/2017).
- [127] NETGEAR Inc. *PLP1200 Datasheet*. 2015. URL: <http://www.downloads.netgear.com/files/GDC/datasheet/en/PLP1200.pdf> (visited on 11/24/2017).
- [128] HomePlug Powerline Alliance. “HomePlug AV Specification Version 1.1”. In: (2007).
- [129] Haniph A Latchman et al. *HomePlug AV and IEEE 1901: a handbook for PLC designers and users*. John Wiley & Sons, 2013.
- [130] Mitch Mitchell. *Using a Powerline Adapter with an RCD*. In comment response. 2015. URL: <https://www.techanswers.org.uk/answers/powerline-adapters-and-rcd-protection.html> (visited on 11/25/2017).
- [131] Ectophile. *Powerline adapter does not work on some sockets*. In comment response. 2010. URL: <https://community.bt.com/t5/Connected-Devices-Other/Powerline-adapter-does-not-work-on-some-sockets/td-p/1131228> (visited on 11/25/2017).
- [132] Kaywan H Afkhamie et al. “An overview of the upcoming HomePlug AV standard”. In: *Power Line Communications and Its Applications, 2005 International Symposium on*. IEEE. 2005, pp. 400–404.
- [133] International Telecommunication Union (ITU). “G9960: Unified high-speed wireline-based home networking transceivers – System architecture and physical layer specification”. In: (2015).

- [134] Vladimir Oksman and Stefano Galli. “G. hn: The new ITU-T home networking standard”. In: *IEEE Communications Magazine* 47.10 (2009).
- [135] Rob Millerb Ishtiaq Roufa et al. “Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study”. In: *19th USENIX Security Symposium, Washington DC*. 2010, pp. 11–13.
- [136] Stephen Checkoway et al. “Comprehensive experimental analyses of automotive attack surfaces.” In: *USENIX Security Symposium*. Vol. 4. San Francisco. 2011, pp. 447–462.
- [137] Roel Verdult, Flavio D Garcia, and Josep Balasch. “Gone in 360 seconds: Hijacking with Hitag2”. In: *21st USENIX Security Symposium*. 2012, pp. 237–252.
- [138] Flavio D Garcia et al. “Lock it and still lose it—on the (in) security of automotive remote keyless entry systems”. In: *25th USENIX Security Symposium*. 2016.
- [139] Reuters. *Paris plans to banish all but electric cars by 2030*. <https://www.reuters.com/article/us-france-paris-autos/paris-plans-to-banish-all-but-electric-cars-by-2030-idUSKBN1CHOSI>. 2017.
- [140] Engadget. *California bill would ban new fossil fuel vehicles from 2040*. <https://www.engadget.com/2018/01/04/california-bill-would-ban-new-fossil-fuel-vehicles-from-2040/>. 2018.
- [141] BBC. *Petrol and diesel ban: How will it work?* <https://www.bbc.co.uk/news/uk-40726868>. 2017.
- [142] Zeinab Rezvani, Johan Jansson, and Jan Bodin. “Advances in consumer electric vehicle adoption research: A review and research agenda”. In: *Transportation research part D: transport and environment* 34 (2015), pp. 122–136.
- [143] Mark Chediak. *Electrify America Plans \$200 Million for California Clean Cars*. <https://www.bloomberg.com/news/articles/2018-10-03/electrify-america-plans-200-million-for-california-clean-cars>. 2019.
- [144] International Energy Agency. *Global EV Outlook 2018*. 2018. URL: <https://www.iea.org/gevo2018/>.
- [145] Cristina Alcaraz, Javier Lopez, and Stephen Wolthusen. “OCPP protocol: Security threats and challenges”. In: *IEEE Transactions on Smart Grid* 8.5 (2017), pp. 2452–2459.
- [146] Kaibin Bao et al. “A threat analysis of the vehicle-to-grid charging protocol ISO 15118”. In: *Computer Science-Research and Development* 33.1-2 (2018), pp. 3–12.
- [147] Daniel Zelle et al. “Anonymous Charging and Billing of Electric Vehicles”. In: *Proceedings of the 13th International Conference on Availability, Reliability and Security*. ACM. 2018, p. 22.
- [148] Achim Friedland. *Security and Privacy in the Current E-Mobility Charging Infrastructure*. <https://blog.deepsec.net/deepsec2016-talk-security-privacy-current-e-mobility-charging-infrastructure-achim-friedland/>. DeepSec, 2016.
- [149] Matthias Dalheimer. *Ladeinfrastruktur für Elektroautos: Ausbau statt Sicherheit*. https://media.ccc.de/v/34c3-9092-ladeinfrastruktur_fur_elektroautos_ausbau_statt_sicherheit. 2017.

- [150] European Alternative Fuels Observatory. *Electric vehicle charging infrastructure*. 2018.
- [151] U.S. Department of Energy Alternative Fuels Data Centre. *Alternative Fueling Stations*. <https://www.afdc.energy.gov/stations/>. 2018.
- [152] CharIN e.V. *What is the Combined Charging System?* <https://www.charinev.org/ccs-at-a-glance/what-is-the-ccs/>. 2018.
- [153] Jonathan M. Gitlin. *Electrify America will deploy 2,000 350kW fast chargers by the end of 2019*. <https://arstechnica.com/cars/2018/04/electrify-america-will-deploy-2000-350kw-fast-chargers-by-the-end-of-2019/>. 2018.
- [154] Share&Charge. *Share&Charge*. <https://shareandcharge.com/>. 2019. (Visited on 01/17/2019).
- [155] ZF Car eWallet GmbH. *Car eWallet*. <https://car-ewallet.de/index.php/what-we-do/>. 2019. (Visited on 01/17/2019).
- [156] AMO Labs. *AMO Labs Preparing to Enter the European Market with Gridwiz!* 2018. URL: <https://medium.com/amo-blockchain/amo-labs-preparing-to-enter-the-european-market-with-gridwiz-7fbf6f19c652> (visited on 01/06/2019).
- [157] ElaadNL. *IOTA Charging Station*. 2018. URL: <https://www.elaad.nl/projects/iota-charging-station/> (visited on 01/06/2019).
- [158] *Road vehicles — Vehicle to grid communication interface — Part 1: General information and use-case definition*. Standard. Geneva, CH: International Organization for Standardization, 2013.
- [159] EcoG. *Providing a customized electric vehicle (EV) fast charging experience through a PaaS for value added services & shared revenue streams*. 2019. URL: <https://www.ecog.io/> (visited on 01/06/2019).
- [160] Organisation Internationale des Constructeurs d’Automobiles. *World Motor Vehicle Production: World Ranking of Manufacturers*. 2016. (Visited on 2018).
- [161] InsideEVs. *Tesla Model 3 With CCS Combo Inlet, S & X With CCS Adaptor In Europe*. <https://insideevs.com/tesla-model-3-ccs-combo-s-x-adaptor/>. 2019.
- [162] HomePlug Powerline Alliance. “HomePlug Green PHY Specification”. In: *HomePlug, June* (2010).
- [163] Johan Peeters. “Fast charging just got faster”. Presentation at eMove360 Conference 2017.
- [164] Fastned. *Autocharge*. <https://support.fastned.nl/hc/en-gb/articles/115012747127-Autocharge->. 2019. (Visited on 01/06/2019).
- [165] Open Fast Charging Alliance. *Automatic charging start and authorization of electric vehicles*. 2017. URL: <https://github.com/openfastchargingalliance/openfastchargingalliance/blob/master/autocharge-final.pdf> (visited on 01/06/2019).

- [166] Seokcheol Lee et al. “Study on analysis of security vulnerabilities and countermeasures in ISO/IEC 15118 based electric vehicle charging technology”. In: *IT Convergence and Security (ICITCS), 2014 International Conference on*. IEEE. 2014, pp. 1–4.
- [167] Marc Mültin. “Das Elektrofahrzeug als flexibler Verbraucher und Energiespeicher im Smart Home”. PhD thesis. KIT-Bibliothek, 2014.
- [168] *Road vehicles — Vehicle to grid communication interface — Part 2: Network and application protocol requirements*. Standard. Geneva, CH: International Organization for Standardization, 2014.
- [169] *Electromobility - Digital communication between a d.c. EV charging station and an electric vehicle for control of d.c. charging in the Combined Charging System*. Standard. Deutsche Institut für Normung, 2014.
- [170] Rainer Falk and Steffen Fries. “Electric vehicle charging infrastructure security considerations and approaches”. In: *Proc. of INTERNET (2012)*, pp. 58–64.
- [171] Christina Höfer et al. “POPCORN: Privacy-preserving charging for eMobility”. In: *Proceedings of the 2013 ACM workshop on Security, privacy & dependability for cyber vehicles*. ACM. 2013, pp. 37–48.
- [172] Cesar Bernardini, Muhammad Rizwan Asghar, and Bruno Crispo. “Security and privacy in vehicular communications: Challenges and opportunities”. In: *Vehicular Communications (2017)*.
- [173] Flavio D Garcia, GT de Koning Gans, and Roel Verdult. “Exposing iClass key diversification”. In: *Proceedings of the 5th USENIX Workshop on Offensive Technologies (WOOT’11)*. 2011.
- [174] Ross Anderson. *Security engineering*. John Wiley & Sons, 2008.
- [175] Monjur Alam et al. “One&Done: A Single-Decryption EM-Based Attack on OpenSSL’s Constant-Time Blinded RSA”. In: *27th USENIX Security Symposium*. 2018, pp. 585–602.
- [176] Giovanni Camurati et al. “Screaming channels: When electromagnetic side channels meet radio transceivers”. In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2018, pp. 163–177.
- [177] Peter Smulders. “The threat of information theft by reception of electromagnetic radiation from RS-232 cables”. In: *Computers & Security 9.1 (1990)*, pp. 53–58.
- [178] Matthias Schulz et al. “Trust the wire, they always told me!: On practical non-destructive wire-tap attacks against Ethernet”. In: *Proceedings of the 9th ACM Conference on Security & Privacy in Wireless and Mobile Networks*. ACM. 2016, pp. 43–48.
- [179] Mordechai Guri, Matan Monitz, and Yuval Elovici. “Usbee: Air-gap covert-channel via electromagnetic emission from usb”. In: *Privacy, Security and Trust (PST), 2016 14th Annual Conference on*. IEEE. 2016, pp. 264–268.
- [180] Jean-Pierre Hubaux, Srdjan Capkun, and Jun Luo. “The security and privacy of smart vehicles”. In: *IEEE Security & Privacy 3 (2004)*, pp. 49–55.

- [181] Mohammad Khodaei, Hongyu Jin, and Panagiotis Papadimitratos. “SECMACE: Scalable and robust identity and credential management infrastructure in vehicular communication systems”. In: *IEEE Transactions on Intelligent Transportation Systems* 19.5 (2018), pp. 1430–1444.
- [182] Kexiong (Curtis) Zeng et al. “All Your GPS Are Belong To Us: Towards Stealthy Manipulation of Road Navigation Systems”. In: *27th USENIX Security Symposium*. Baltimore, MD: USENIX Association, 2018, pp. 1527–1544. URL: <https://www.usenix.org/conference/usenixsecurity18/presentation/zeng>.
- [183] Andrés Molina-Markham et al. “Private memoirs of a smart meter”. In: *Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building*. ACM, 2010, pp. 61–66.
- [184] David Varodayan and Ashish Khisti. “Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage”. In: *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 1932–1935.
- [185] DBT. *Quick Charger Dual DC Product Datasheet*. 2014. URL: <http://www.dbtcev.fr/wp-content/uploads/2014/01/TEQC0c0146-EN-Quick-Charger-Dual-DC.pdf> (visited on 03/28/2018).
- [186] Chargemaster Ltd. *Polar Network*. <https://chargemasterplc.com/polar/>. 2018.
- [187] ABB. *Terra 53 Product Leaflet*. 2017. URL: https://www.franklinenergy.co.uk/wp-content/uploads/2017/07/4EVC204308-LFEN_Terra53C-CT-CJ-CJG.pdf (visited on 03/28/2018).
- [188] POD Point. *Open Charge Electric Car Charging Stations*. <https://pod-point.com/open-charge>. 2018.
- [189] Ecotricity. *Electric Highway*. <https://www.ecotricity.co.uk/for-the-road>. 2018.
- [190] InstaVolt. *Our Technology*. <https://instavolt.co.uk/about-us/our-technology/>. 2018.
- [191] InstaVolt Ltd. *About InstaVolt*. <https://instavolt.co.uk/>. 2018.
- [192] BP Chargemaster. *Chargemaster Ultracharge 500S Datasheet*. 2019. URL: <http://bpchargemaster.com/wp-content/uploads/2019/01/Ultracharge-brochure.pdf> (visited on 02/02/2019).
- [193] EVTRONIC. *Quickcharger Product Datasheet*. 2016. URL: http://evtronic-industries.com/wp-content/uploads/2015/10/EVTRONIC_QUICKCHARGER_Product-Sheet_2016-11-17-1.pdf (visited on 03/28/2018).
- [194] Chargepoint Services. *GeniePoint*. <https://www.chargepointservices.co.uk>. 2019. (Visited on 02/02/2019).
- [195] Efacec. *Efacec QC45 Datasheet*. 2016. URL: http://electricmobility.efacec.com/wp-content/uploads/2016/10/CS119I1307D1_QC45.pdf (visited on 02/02/2019).

- [196] Shell Plc. *Welcome to Shell Recharge*.
<https://www.shell.co.uk/motorist/welcome-to-shell-recharge.html>.
2019. (Visited on 02/02/2019).
- [197] Efacec. *QC45S Product Page*.
<https://electricmobility.efacec.com/ev-qc24s-quick-charger/>. 2019.
(Visited on 02/15/2019).
- [198] Harald T Friis. “A note on a simple transmission formula”. In: *proc. IRE* 34.5 (1946), pp. 254–256.
- [199] CharIn. *Target Grid Integration Levels*.
<https://insideevs.com/ccs-combo-standard-v2g-2025/>. 2019.
- [200] INSYS MICROELECTRONICS GmbH. *INSYS Powerline GP Manual*.
https://256.insys-icom.com/bausteine.net/f/10637/HB_en_INSYS_Powerline_GP_1711.pdf?fd=0. 2017. (Visited on 05/14/2019).
- [201] devolo AG. *dLAN Embedded PLC Module Datasheet*.
https://www.codico.com/shop/media/datasheets/Devolo_dLAN_Green_PHY_Module_20130713_en_data_sheet_019.pdf. 2012. (Visited on 05/14/2019).
- [202] Recargo Inc. *PlugShare*. <https://www.plugshare.com/>. 2018.
- [203] ElaadNL. *Update Global EV Charging Test: PKI Workshop*. 2018. URL: <https://www.elaad.nl/news/update-global-ev-charging-test-pki-workshop/> (visited on 11/02/2018).
- [204] Paul Klapwijk and Lonneke Driessen-Mutters. *Exploring the public key infrastructure for ISO 15118 in the EV charging ecosystem*. ElaadNL, 2018. URL: <https://www.elaad.nl/news/publication-exploring-the-public-key-infrastructure-for-iso-15118-in-the-ev-charging-ecosystem/>.
- [205] Sébastien Dudek, Jean-Christophe Delaunay, and Vincent Fargues. “V2G Injector: Whispering to cars and charging units through the Power-Line”. In: *Symposium sur la sécurité des technologies de l’information et des communications (SSTIC 2019)*. 2019.
- [206] Bastian Bloessl et al. “An IEEE 802.11 a/g/p OFDM Receiver for GNU Radio”. In: *Proceedings of the second workshop on Software radio implementation forum*. ACM. 2013, pp. 9–16.