Passive Steady State RF Fingerprinting: A Cognitive Technique for Scalable Deployment of Co-channel Femto Cell Underlays

Irwin O. Kennedy, Patricia Scanlon

Bell Laboratories
Alcatel-Lucent
Blanchardstown, Dublin 15
Republic of Ireland
irwinkennedy@alcatel-lucent.com

Milind M. Buddhikot

Bell Laboratories
Alcatel-Lucent
Murray Hill, New Jersey
USA

Abstract

Recently, cellular operators have begun evaluating femto cells that aggressively reuse spectrum to cover a small spatial footprint (10m radius). A large-scale femto cell underlay network will increase an operator's total number of cells by two to three orders of magnitude and presents significant scaling problems in spectrum reuse. We argue that to achieve this scale and interoperability with the existing UMTS network and handsets, the femto cell must use novel yet simple cognitive techniques: sensing, smart handover and idle mode cell camping. We report in detail on the problem of signaling storms caused by increased core network-signaling load due to idle mode cell camping. We provide a novel passive RF fingerprinting technique as a solution to tackle the problem. The technique is based on frequency domain characteristics. Our technique detects the unique characteristics imbued in a signal as it passes through a transmit chain. We are the first to propose the use of discriminatory classifiers based on steady state spectral features. In laboratory experiments, we achieve 91% accuracy at 15dB SNR based on seven different models of UMTS user equipment. In the largest known laboratory experiment of its kind, we report an accuracy of 85% using our technique on twenty UMTS user equipment. This large test set includes 10 identical devices. Our technique can be implemented using today's low cost high-volume receivers and requires no manual performance tuning.

1. Introduction

Recently, large cellular service providers have started considering the deployment of low power, indoor, "femto" base stations. The first generation femto cells deployments will rely on static allocation of spectrum wherein a portion of the total spectrum licensed to the operator will be reserved for femto cells. This form of spectrum usage that is mutually exclusive with the macro-cells ensures that carefully engineered macro-cells are not impacted. However, such a model of deployment is not advisable in the long run for several reasons. First, in several countries, especially in Europe, available 3G spectrum in which UMTS technologies are currently deployed is very small. In fact it is often limited to a single 5 MHz carrier required for UMTS. As such reserving 5 MHz carrier for UMTS femtos is either impossible or not advisable due to loss of macro-cell capacity. As air interface standards evolve to wider bands (e.g.: 20 MHz in WiMAX, LTE), static allocation becomes expensive. Therefore, the femto cells must concurrently use the same spectrum that macro-cells use. This aggressive form of spatial-temporal reuse of spectrum, often called "concurrent co-channel reuse" represents the *private commons* model of spectrum access [2, 1].

However, realizing such spectrum sharing or access poses significant challenges some of which have been addressed in the context of UMTS co-channel femto cells in literature [3, 9]. First problem is that dense deployments of femto cells wherein 1000s of femto cells are deployed per macro-cell can lead significant femto-to-macro interference and if not managed can lead to reduction in macro-cell capacity and performance. Claussen et al.[3] demonstrated via a realistic simulation study of femto deployment in an east-London suburb that femto-to-macro interference can be controlled via appropriate power management and their impact on macro-cell performance metrics such call drops can be made marginal.

The second more pressing problem femtos pose is their adverse impact on network signaling traffic which is an artifact of two design requirements: (1) end-user handsets are

978-1-4244-2017-9/08/$25.00 ©2008 IEEE

unchanged and see femto cell base stations same as macro-cell base stations. (2) Femto cells are retrofitted into the legacy macro-cellular location and paging architecture by assigning them a different Location Area Code (LAC) than that of macro-cells they embed in.

The two requirements cause naive deployments of co-channel femto cells to suffer from "signaling storms" and weaken the privacy and security mechanisms. New techniques are necessary to enable a femto base station to rapidly detect the end-user handsets without excessive interaction with and modifications to the rest of the macro-cell infrastructure and reject the UEs that are not to be served. Such a technique will be crucial to enabling scalable deployment of femto underlays.

1.1 Our Contributions

In this paper, we present a novel cognitive technique called *Passive RF Fingerprinting* designed to enable scalable deployment of femto underlays. We present our technique and its experimental evaluation in the context of UMTS femto cells. However, the technique will be equally applicable to other 3G/4G cellular technologies as it exploits signaling structure commonly used in most cellular air interface protocols. The desirable properties of our technique are as follows: (1) By operating on physical layer (signals), our technique is agnostic to higher layer protocols. (2) It requires no modifications to the standards used in the macro-cells (specifically 3GPP in our design). (3) It requires no changes to the handsets. (4) It is very easy to provision and manage and provides a significant reduction in the 5x-40x rise in signaling traffic created by a naive deployment.

1.2 Outline of the Paper

The remainder of the paper is organised as follows. In Section 2 we provide the motivation and context for this work by describing the challenges in co-channel femto cell deployments. In Section 3 we give a high-level system level description of our techniques application before reviewing the literature in the area of transmitter identification. The method of our technique is described in Section 4. We then describe the experimental apparatus in Section 5 before providing the results in Section 6 and concluding in Section 7.

2 Understanding Challenges in Co-channel Femto Deployments

In this section we elaborate on the challenges posed by co-channel femto cells.

2.1 Signaling storms

To understand what is meant by a signaling storm, lets consider the scenario depicted in Figure 1

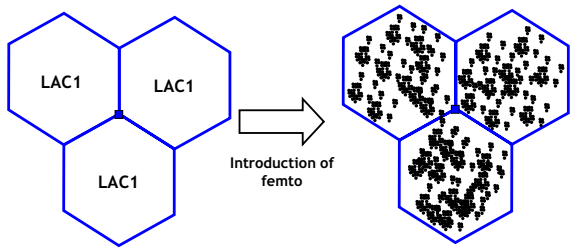

Figure 1. Dense deployment of femto cells

On the left side of the figure we depict cluster of three macro cells assigned location area code (LAC) equal to one. In absence of femto cells, when a UE moves between the three cells, LAC does not change and UE does not perform any location area updates. When the network needs to page the UE (e.g. to notify the UE of an incoming phone call), it broadcasts a paging message in all cells corresponding the last known LAC. In our example, the broadcast will be sent to the three macro cells. Now consider the scenario on the right hand side showing the same three macro cells with hundreds or thousands of femto cells (shown by dots) within it. Since the core network must know which network (femto or macro) to route the paging messages, the femto LAC must be different to the macro-cell LAC. The result is that every time a UE moves between camping on a macro to camping on a femto (and vice versa) it must perform a location area update. Prior to successful location update, the handset needs to be authenticated by the femto base station. This generates additional authentication traffic to the VLR/HLR/HSS units in the core network. In case, a femto cell is configured as a private femto cell allowed to serve select few handsets of end-users (e.g. those who reside in the home), after authentication the femto cell may deny the handset from camping on it in idle mode. However, this negative decision reached after exchanging large number of signaling messages. In very dense femto deployments and in presence of large number of UEs, the amount of the resulting signaling traffic can be very high and in the event a large number of residential femto cells are configured private, a large portion of the traffic is wasted. Simulations indicate that the introduction of femtos could result in a 5x-40x increase in core network signaling versus an existing macro cellular network[3].

One specific worst case scenario is shown in Figure 2 which shows a large macro-cell with a subset of embedded femto cells formed by femto base stations deployed in row-houses which are close to the street.

This scenario will be common is a large number of cities

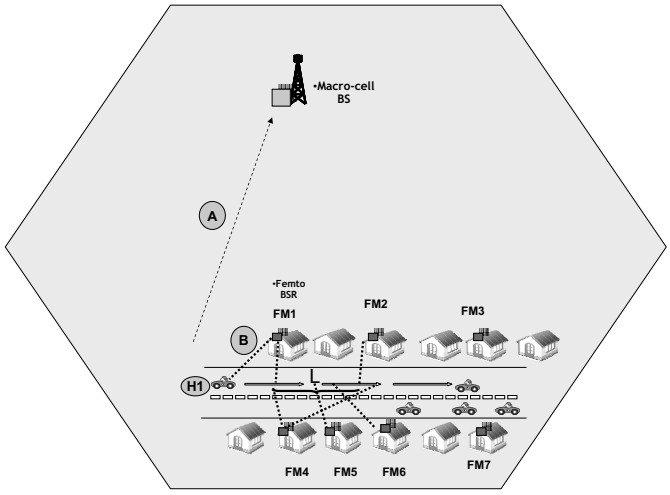

Figure 2. Worst case scenario for signaling storms

worldwide. In our example, the representative handset of end-user in the car is marked "H1" is initially associated with the macro-cell base station BS (shown by signal path "A"). As the car drives down the street with row-houses, the handset detects the pilot signal of the femto FM1. As it continues to drive down the street it sees signal from various femto BSs from FM1 to FM7. The set of visible femtos rapidly changes as an example, in the stretch of street marked L, the handset sees FM2, FM3, FM4, FM5, and FM6 BSs. In order to enable correct paging, the femto BSs are assigned a different location area code (LAC) than the macro cell. If the handset H1 is currently in the idle mode, it sends a location update to the network indicating that its location has changed. When we consider a large number of such handsets in a macro-cell, the aggregate updates can happen rapidly leading to a significant increase in the signaling load a phenomena that can be termed as a "signaling storm".

We have considered cars in this example. Pedestrians are also valid – any UE passing by the femtos may attempt to camp on them. The simulations mentioned earlier which indicate a 5x-40x increase[3] in core network location area update signaling load consider pedestrian traffic only (moving at a speed of less than 4Km/Hr.) It may be possible for the femto to detect faster moving UEs through the Doppler effect and automatically refuse their request to camp on the femto. This can be assumed because they are moving too quickly to require the femto's local service and should be kept on the macro-cellular network.

2.2 Impairment of security mechanisms

Another problem that occurs in the above scenario is that the permanent identifier for the end-user device called IMSI (International Mobile Subscriber Identification number) may be forced to be exchanged in plain text far more often than is done otherwise in absence of femtos and therefore, weaken the security/privacy in the network.

To understand this, we note that in a UMTS network there are two primary ways to identify the UE (user equipment) and each has its own drawback:

(1) **Mapping a temporary TMSI (Temporary Mobile Subscriber Identification number) identity of the handset to its IMSI**. When the UE performs a location area update, it sends its TMSI. Using option 2 above, if the femto base station knows the IMSI to which the TMSI belongs, it can identify the UE. If not, the femto base station must contact a node in the core network to resolve the mapping from TMSI to IMSI. This will result in a very large increase in signaling load on the network equipment that provides the mapping between TMSI and IMSI. The TMSI is changed by the network periodically to protect privacy so a previously stored mapping at the femto can become invalid.

(2) **Identity request message-response:** In this, the femto spoofs an identity request message by the MSC to get the UE to send its IMSI. The IMSI is unique to the UE, so the femto is then able to identify the UE. However, the IMSI is sent in plain text (unencrypted) over the air. Therefore, it allows an eavesdropper to easily identify the location of a UE. This weakens the UE privacy and security dramatically. The 3GPP UMTS standards mention that the identity-request mechanism should be invoked by the serving network whenever the user cannot be identified by means of a temporary identity. It also mentions that "This represents a breach in the provision of user identity confidentiality".

3 Passive RF Fingerprinting

3.1 Overview of Our technique

Our technique provides a third method for rapidly identifying the UE. It seamlessly co-exists with the existing 3GPP standards and is very easy to deploy and provision. The salient idea in our technique called *Passive Radio Frequency Fingerprinting* (Figure 3) relies on extracting the unique signature of the transmit chain electronics in the UE when a well known signal is transmitted by the UE. The femto cell creates a model or template of the unique signature and stores this locally. When it encounters the known signal transmitted by a UE, it computes the signature and cross-checks with its local model. If a sufficiently close match is found the UE is accepted for service, otherwise it denies service. To eliminate the possibility of rejecting

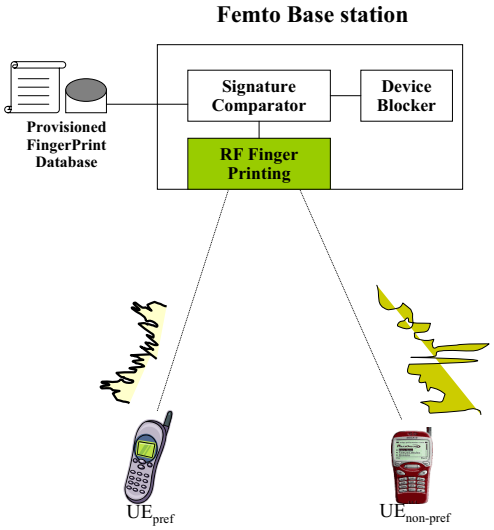

Figure 3. Using Passive RF Fingerprinting

a valid UE from gaining access, the technique errs on the conservative side and admits the device. In the case of false admission, higher level protocols still rightly block the device from gaining access.

3.2 Previous Work on Transmitter Identification

Signal detection and device identification has been one of the major areas of research in the development of wireless. Several researchers have reported the possibility of identifying a radio transmitter by analysing the received signals. This work stretches back to 1940s and radar transmitter identification [12]. The majority of techniques focus on transmissions at the physical layer.

Physical layer fingerprinting techniques may be split into two groups: transient signal techniques and steady state signal techniques. A transient signal is transmitted upon transmitter stage power up and power down. It is the short period (typically micro seconds) during which capacitive loads charge or discharge, the power amplifier ramps its power output and in some cases, where the frequency synthesizer makes the transition between steady state frequency generation and being powered off. The steady state period of signal transmission is defined here as the period between the start and end transients.

Almost every radio radiates a transient signal upon switch-on and switch-off. For this reason, transient analysis has enjoyed the most attention in the literature [8, 15, 16, 17, 18]. Transient analysis discriminates using the minor amplitude variations that occur upon transmitter switch on. Due to the short duration of transient signals, very accurate and consistent detection of the transient part of the signal is important for good identification performance. However, it also poses the most significant challenge. The receiver architecture is unusual since it must be capable of digitising at extremely high sample rates. This is necessary to provide the resolution of amplitude information required for the transient feature extraction algorithms. For example 5 GSamples/s is used by Serinken et al. [5, 16] and 500 MSample/s is used by Hall et al. [7]. The two key approaches are the threshold [15] and Bayesian step change detector [17, 18]. Both rely on reception at high SNR and an abrupt change at the start of the transient - both of which may not exist in practice. A third approach based on frequency domain analysis was recently proposed by Hall et al. [8]. Rather than relying on amplitude characteristics for start and end time estimation, the authors were able to produce reasonable estimates by analysing the variance of its spectral components under high SNR conditions. However, as noted by the authors, it is not yet known to what extent it is possible to find distinguishing characteristics in the transients in larger device sets. Others have reported that the level of difference between identical transmitters manufactured by the same company may not be distinguishable using transient analysis [5]. Where it is possible to reliably detect the transient start and end points, several researchers have reported good classification performance. In excess of 90% for high SNR environments [7, 15, 16].

We are only aware of one example in the literature of studying the steady state signal [6]. The main reason for this is the apparent lack of a steady state signal common to all devices. That is, a steady state signal is either unmodulated or contains the exact same data modulation. This property is important since the signal provides a benchmark for discovering difference between transmitters. The lack of a steady state signal common to all transmitters is no longer the case in modern transmitters. Today's digital transmitters intentionally introduce repetitive sequences such as preambles to simplify receiver design. This makes steady state signal analysis feasible today. Recently Gerdes et al. [6] proposed that analysing the steady state signal may provide the ability to distinguish between same model cards manufactured by the same company. They argue that the transient signal is so short that it cannot contain enough information to discriminate between similar devices. Their focus is on wire line transmitters where similar principles apply. A portion of the IEEE Ethernet 802.3 frame preamble common to all devices was identified and used to construct a device fingerprint. They use a matched filter implementation and simple thresholding to perform classification. Training involves characterising the matched filter's output to determine the output magnitude that corresponds to a match for a particular transmitter. The discriminatory capabilities of this approach are unclear. No overall level of accuracy is provided. It appears that the thresholding decision for device identification can result in more than one device being

identified. The result is many false-positive identifications. Their system also requires many ad hoc steps to tune the performance. For example, the discriminatory performance was manually refined through a combination of bandpass filtering, creating an ensemble of matched filters and time domain trimming.

We note that some researchers have successfully demonstrated identification using techniques at higher layers of the network stack [4, 11, 13]. However they do not offer the same level of discrimination as physical layer based techniques and have longer convergence times. For example, in [13], the researchers discovered a way to use the TCP time stamp field to reveal the identity of the transmitting node. The measurements are performed over the Internet via wired connections. The technique relies on many thousands of TCP packets being received - so it is not feasible where a technique is required to assist in authentication for network access.

In summary, physical layer fingerprinting offers promise for passive discrimination between a large set of wireless transmitters. We note that:

1. Transient analysis offers good classification performance only where the beginning and end of the transient can be reliably identified.

2. It has been reported in [5, 6], that transient analysis is not always able to distinguish between same manufacturer/same model variants.

3. The very high sample rates demanded by transient analysis requires more sophisticated receiver architectures than are otherwise required for communication.

Steady state signals offer a relatively unexplored alternative to transient analysis. We note that if discrimination is possible in the frequency domain, the use of standard low cost ADC sample rates and receiver architectures will be made possible.

4. Method

Our approach to RF fingerprinting uses frequency domain analysis combined with traditional discriminatory classifiers to perform device identification. When compared to the only known previously published steady state technique[6], our technique offers a significant performance improvement through more flexible feature extraction approach and the use of a k-NN discriminatory classifier. Our work is distinguished from the large body of previous work on transient based analysis by its focus on the steady state portion of the signal. The main advantage over transient analysis is that it can be implemented using today's low-cost radio receiver front ends e.g. Ethernet access points or femto cells. These radio receivers capture the signal at sufficiently high sample rates for our proposed approach. By contrast, transient based approaches require very high sample rates to capture the amplitude fluctuations of the transient part of the signal.

We propose that our approach takes advantage of the bulk effect of small differences that exist at all stages of transmitter manufacturer. Differences in component design (filters, power amplifiers, inductors, capacitors), same component manufacturing tolerance spread, PCB materials and PCB soldering etc. These differences are imbued upon the transmitted signal and their bulk effect can be detected at the receiver. Since the receiver's frequency response can be assumed to stay constant, the only differences in the baseband received signal are due to different transmitters.

Figure 4 illustrates the processing steps involved in device identification. The input to the preprocessing stage is the received RF signal from the transmitter. For convenience we constrain the higher layers of the communications system to transmit exactly the same signal every time. For example, for the random access channel (RACH) preamble in UMTS, the signature and up-link scrambling code pair are constrained to a single combination, rather than the usual 16-48 different combinations. A standard radio receiver architecture is employed, down converting the transmit band to baseband, before being bandpass sampled by the ADC at the Nyquist rate.

The next step in Figure 4 is carrier frequency offset correction. Captured preambles are separated in time by a period of no transmission. The preamble sequences were extracted from the signal prior to down sampling using a sum of the absolute values window function. The window has length equal to the number of samples in a preamble. It is shifted across the file in 10 sample increments, with the total energy recorded for each window. For every set of samples between two periods of no transmission, the window with the maximum energy is extracted as the preamble. No attempt is made to distinguish between transient and steady state portions of the signal. We estimate the complete preamble is extracted with better than 99.99% accuracy relative to its total energy content.

We do not apply carrier frequency correction since the UMTS handset disciplines its local quartz crystal based timing source to the down link broadcast carriers. The timing in the base station is based upon a high quality rubidium source. To remove time domain amplitude characteristics, we normalise the data so that the signal energy is equal to one.

The frequency domain representation of the complex baseband preamble signal is obtained using Fast Fourier Transform (FFT). It is in the frequency domain that the feature vector is constructed for input to the classifier. The output of the FFT is of extremely high dimensionality. Based

on the curse of dimensionality, high dimensionality can lead to poor classifier performance if insufficient training data is provided. Given the limited size of the training data set in these experiments, the feature vector dimensionality is reduced. To reduce the dimensionality and to construct the feature set presented to the classifier, we reduce the number of bins by taking the mean value of multiple FFT bins to form a single new bin. A set of log-spectral-energy features is finally output to the classifier. Prior to the spectral analysis and to remove amplitude variations that may occur each time the signal is transmitted, the time domain samples are energy normalised.

The output of the spectral analysis stage feeds into the final device identification stage in Figure 4. We feed the output features of the feature extraction stage into the classifier. The goal of the classifier is to identify a particular UE based on its frequency domain features. The classifier is trained using a set of labeled training examples i.e. each preamble acquired is labeled according to which board the preamble was extracted from. The classifier then attempts to identify previously unseen preambles and ascertain which board (or class) the preamble was transmitted from. The performance of the classifier is evaluated based on the number of correctly identified UEs.

The data collected from each board is divided into two sets. The first set is used in the classification training step and the second test set is omitted from training and used to test performance of the system. In the k-NN algorithm the training preambles are mapped into multidimensional feature space which is partitioned into regions based on the class labels. The frequency domain features extracted from a preamble are said to belong to a particular board or class if it is the most frequent class label among the k nearest training samples, where distance is determined using the Euclidean distance metric. In these experiments $k = 5$ is used. In this RF Fingerprinting system each class represents one of the 20 possible UMTS UEs. The system is presented with a previously unseen preamble and attempts to discriminate between each of the twenty candidate classes to determine from which UE this preamble is obtained from. Classification accuracy is given as the percentage of correctly identified boards.

The complete system can be described in such a way to enable autonomous operation - requiring no user assistance. For example, in an application for monitoring access attempts to a UMTS femto cell, training data can be captured when a new client is installed. The SNR at which the data is captured can be calculated by examining the sampled data. Several classifier models for different SNRs can be constructed by adding additive white Gaussian noise to the captured data. When an access request from an UMTS client is received, the access point measures the SNR and the appropriate model is selected. To increase classification accuracy, several transmissions may be solicited by the access point. The way in which this is achieved is system specific, for example in UMTS, if the base station does not respond, an UMTS client will retransmit the message.

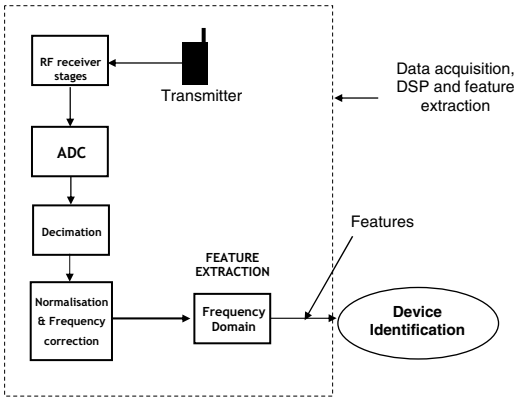

Figure 4. Processing Chain

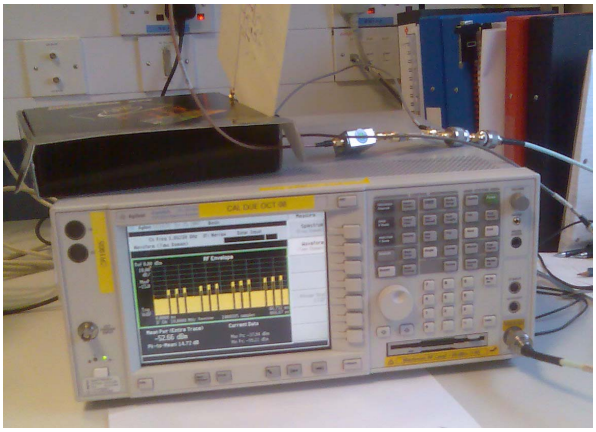

Figure 5. Agilent E4440A PSA Spectrum Analyser.

5. Experimental Apparatus

The test equipment used for capturing the digital I/Q samples is an Agilent E4440A PSA spectrum analyser with option B7J. The Agilent analyser can be seen in the photograph in Figure 5. In the photograph, the analyser's display shows a capture of the RF power envelope against time for a UMTS RACH preamble ramp sequence. All captures are performed at 15MSamples/s. The UMTS User Equipment (UE) used was a mixture of UMTS class 4 PCMCIA data cards and UMTS mobile phone handsets. From an RF perspective, both PCMCIA cards and handsets comply with the same set of standard UE specifications. For the purposes of these experiments, whether the transmitter is in a

PCMCIA card or mobile phone handset should not have any significance. A laptop with linux loaded provided the power supply for the PCMCIA cards to power up, boot their internal flash based firmware and transmit the RACH preamble power ramp sequences. No driver software was required.

We list the UEs used in the experiments in Table 1. As can be seen from the table, the group is composed of 4 mobile handsets and 16 PCMCIA cards. The group was simply composed from what was available in the laboratory and previously used for other purposes. The UE ID sequence has some entries missing due to the UE either not working or not being available at the time of measurement. There are 5 different manufacturers - Nokia, Samsung, Sierra Wireless, Novatel Wireless and Option. There are 10 identical Novatel U530 models, making for a challenging discrimination problem. Several differentiating markings can be observed on the case of the UEs. Some of these appear to be HW model revisions (including what is possibly the date of manufacturer) and serial numbers. We list these in the table under the columns HW Version and Other Details.

A number of SIM cards programmed for the UMTS base station are available. When a UE is having measurements performed, a SIM card is inserted and the UE is placed on a chair about 0.4m above the floor. The base station is configured to transmit with very low power (less than 100mW). The base station and the UE are separated by lab partitions and desks at a distance of about 5 metres, without any line of sight. The laboratory is on the first floor of a four storey building. The labs dimensions are approximately 6m by 4m. It has one external wall with a window looking over a large car car park and the walls are light partition walls, separating the laboratory from a large open plan office space.

We have full access to the software load on an Alcatel-Lucent 2100 MHz UMTS base station. In the UMTS standard, system information blocks (SIBs) are broadcast in the cell down link channels. The UE reads the SIBs and uses them to configure its operation. Of particular interest to us in this work is the ability to restrict the UE to use only a single combination of RACH preamble signature and scrambling code. This means that every RACH preamble transmission from each of the 20 UEs contains the same digital I/Q content at the transmitter. Typically, the UE would randomly select from a total of 16 RACH preamble signatures and 16 scrambling codes to produce a total of 256 different possible transmissions. The result of constraining all UEs to transmit the same RACH preamble is that the only difference between the signals radiated by the different UEs is due to the analogue transmit stages. Differences in the frequency response at all stages of the transmitter circuitry will combine to influence the power spectral density of the radiated signal and hence the signal received. For example, the clock generation circuitry, power amplifier, mixers, intermediate frequency filters, transmit frequency filter and antenna.

We also edited the basestation software load so the basestation would never respond to a RACH preamble. This meant the UE would never receive a response to its RACH preamble, therefore it would continue to retransmit its RACH preamble ramp sequence. This was done to make it straightforward to capture the examples required for classifier training and testing purposes.

As observed above, the important property of the preamble is that it is always identical and is repeated often. The UMTS preamble occupies a bandwidth of 5MHz and consists of 4096 chips at a rate of 3.84Mcps. The result is a 4096 chip pseudo random quadrature phase shift keying (QPSK) signal. It is root raised cosine filtered with an excess bandwidth $BT = 0.22$.

The Agilent PSA is driven using an ANSI C program running on a laptop. The laptop is connected to the analyser via 100MBit Ethernet. The C program uses SCPI commands to command the PSA over the free to download VISA interface[10]. This enabled the collection of RACH preambles to be automated. The C program can set the capture centre frequency, the sample rate, the RF energy trigger level etc. The analyser is then instructed to single sweep on RF trigger. The program polls the analyser for the trigger event. When a trigger occurs, the program transfers the contents of the anslysers I/Q memory to the laptop hard disc. This is repeated until enough RACH preambles have been captured. Figure 6 shows part of a capture file. The figure was generated in Matlab by plotting the normalised log of the absolute value of the I/Q samples. We can observe the power step used by the UE is 4dBm and the length of each RACH preamble burst is just over 1ms. This power step value can be set in the basestation SIB broadcasts.

We captured around 200 preambles for each UE. The preamble power ramp means that each preamble in a ramp sequence will have higher power than the previous. This means the preambles must be sorted by power for SNR based performance analysis of our identification technique. The 200 preambles per UE produced a total of around 4000 preambles. Around half are used for training purposes with the remainder used to test the technique.

With the raw data file written to disc by the C program, the remaining processing was performed in Matlab. The freely available Netlab machine learning library[14] was used to provide the kNN functionality in Matlab. All other functionality as described in Section 4 was implemented by the authors.

6. Results

We conducted two sets of experiments to test and explore the performance of our technique. The first set of experiments took all 20 UEs and attempted to discriminate

ID	Type	Manufacturer	Model	HW Version	Other Details
1	PCMCIA	Sierra Wireless	Aircard 875		
2	PCMCIA	Sierra Wireless	Aircard 850		
3	PCMCIA	Sierra Wireless	Aircard 850		
4	PCMCIA	Option	Globetrotter		
11	PCMCIA	Novatel Wireless	Merlin U740	HW 04.04 20060927	358661-00-121145-4
17	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 040410	35301800-46147-0
19	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050406	35301800-265044-3
21	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050405	35301800-264904-9
22	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050406	35301800-265504-6
24	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 040713	35301800-099670-7
25	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050406	35301800-265671-3
26	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050407	35301800-265315-7
27	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050406	35301800-265035-1
28	PCMCIA	Novatel Wireless	Merlin U530	HW 01.08 050406	35301800-265346-2
30	Handset	Samsung	SGH-Z107		SN R4WY280389H
31	Handset	Samsung	SGH-Z107		SN R3XXC96948R
32	PCMCIA	Novatel Wireless	Merlin U740	HW 4.4.20060927	358661-00-12134-6
33	Handset	Nokia	6650		350989/10/058035/6
34	Handset	Nokia	6650		350989/10/068406/7
35	PCMCIA	Novatel Wireless	U530(?)	HW Beta 030324	004400-00-303443-4

Table 1. Details of UMTS User Equipment used in experiments.

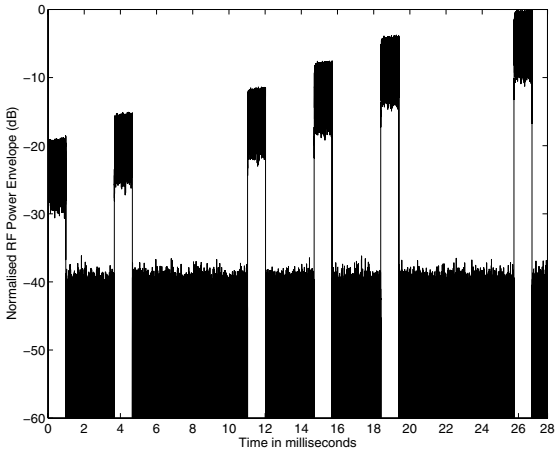

Figure 6. Sequence of UMTS RACH Preambles.

between them. The second set of experiments attempted to discriminate between only different handset models - seven UEs in total. All results presented are based on measurements performed in an indoor multi path environment.

Before we present the results, it is useful to consider Figure 7. It plots the energy normalised power spectral density against frequency for UE 4 and UE 11. The two UEs are from different manufacturers, so intuitively we would expect them to have a different design and to use different components. From the graph, we can visually observe the difference. The gap between the two curves represents the difference imbued by the transmitters on the radiated signal. Our techniques aim is to reliably discern this difference.

Figure 9(a) shows the system's classification performance when discriminating between all twenty different UEs. The Figure plots the percentage classified correctly against the number of bins. A separate line is plotted for six different SNR environments. The binning functions purpose is to reduce the dimensionality of the spectral features fed into the classifier. If the number of bins is set to one, a single feature, the mean energy of the complete spectrogram forms the classifier input. The performance is recorded every 5 bins below 200 bins and every 50 thereafter. As we increase the number of bins, the frequency resolution is increased. As can be seen in the graph, the performance increases rapidly as the number of bins is increased from one to thirty. From thirty onwards the performance levels off.

We notice that the lower SNR conditions demand more bins to achieve the maximum classification performance.

We believe that at high SNR, even the smallest spectral energy differences can contribute to discrimination. As the SNR decreases the small spectral energy differences start to be destroyed. By increasing the number of bins, we increase the number of spectral energy features. In doing so, we reduce the chance that all features have had their complete discriminatory value destroyed by noise. We note that noise will reduce the discriminatory value offered by the smaller bins, hence even after dividing the energy across many bins, lower overall classification performance is still to be expected at low SNR.

The order of the high SNR curves is not as expected and requires careful consideration. We would expect the curves to be in ascending order of ascending SNR environments. However, we observe that 15dB is the highest performing, followed by 20dB in descending order, 30dB, 25dB, 10dB, 5dB and 0dB. It is perhaps easier to observe in Figure 8(a). This shows the peak classification performance at each SNR. The peak of 85% is obtained at an SNR of 15dB. At higher SNR the performance declines. This is counter intuitive - a higher SNR should result in less variance in the value of a bin. We strongly suspect the explanation for this result is a shortcoming of our experiment. It is well known that machine learning techniques can be adversely affected by a lack of training data. Also it is important to use the same number of examples from each UE when training the kNN. Otherwise there is a danger that the resulting classifier will have a bias towards those UEs with more training examples. We always ensured that the same number of training examples were used per UE. However, we found that for some UEs we had fewer captures of high power RACH preambles. This meant we had fewer training examples above 15dB. We suspect this explains the unexpected results contained in Figure 9(a) and Figure 8(a).

The confusion matrix is a useful way of exploring a classifiers ability to discriminate. Table 2 shows the confusion matrix for all twenty UEs at the top performing 15dB SNR. The rows and columns of the table are the UE IDs. Table 1 maps the UE ID to the UE description. The confusion matrix shows both the correct classifications and incorrect classifications for all the test data. The diagonal running from top left to bottom right of the table contains the number of correct classifications of each UE. To the left and right of the row the diagonal score is contained in, we can see the misclassifications. The number is a score out of the total number of test cases of 99. For classifications, high is good, for misclassifications, low is good. We note that the identical Merlin U530 models, located in the centre of the table at ID 17 to 28, have significant confusion between one another. This is to be expected since they will have very similar designs and set of components. As should be expected,

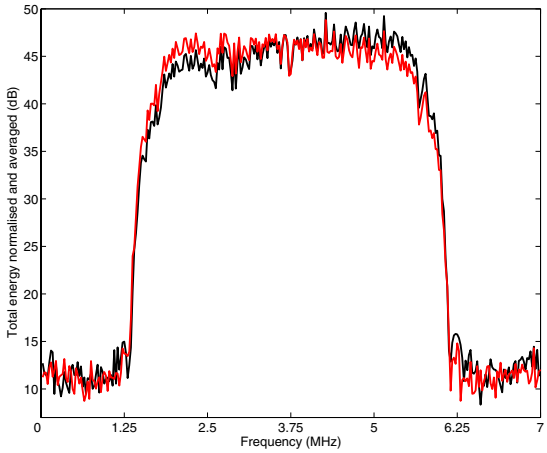

Figure 7. Power Spectral Density of User Equipment ID 4 and ID 11

the two identical model Samsung UEs, ID 30 and 31, show reasonable confusion between one another, but still keep a classification performance above 80%. We notice that UE 4 causes significant confusion for UEs 11, 17, 19, 21 and 25. This is not straightforward to explain, since UE 4 is branded Option and the confusion UEs are all branded Novatel Wireless. Perhaps the products are OEMed to the same external manufacturer? Perhaps the same transmit filter component is used in both designs? An alternative explanation is that significant difference does exist between these UEs and our experiment does not detect it reliably. Difference is discernible in the majority of test cases. Perhaps improvements could be made to the algorithm and experimental setup to increase the discriminatory performance.

The second set of experiments examined the ability to discriminate between the seven different models. One UE was selected to represent each model: 1,2,4,11,17,30,33. Figure 9(b) shows the results of this experiment. It plots the percentage classified correctly against the number of bins from 1 to 100. As noted above during the discussion of the 20 UE results, the order of the curves is again not as expected.

The peak performance can easily be seen in Figure 8(b). At 15dB a peak of 91% classification performance is achieved. This is a 6% improvement over the performance on the complete set of twenty phones. An improvement is expected, since the classifier is no longer expected to discriminate between identical models.

UE	1	2	3	4	11	17	19	21	22	24	25	26	27	28	30	31	32	33	34	35	Total
1	97	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	99
2	1	94	0	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	99
3	0	0	92	3	0	0	0	0	0	1	2	0	0	1	0	0	0	0	0	0	99
4	1	1	1	88	0	1	0	0	0	0	4	0	0	0	0	0	2	1	0	0	99
11	0	2	1	14	79	1	0	0	0	0	0	0	0	1	0	0	0	0	1	0	99
17	0	0	0	12	0	81	0	1	0	0	0	0	0	1	0	2	0	0	1	1	99
19	0	0	0	15	0	13	66	2	0	0	0	1	0	1	0	0	0	0	0	1	99
21	0	1	0	9	0	2	0	75	0	2	5	3	0	0	0	0	0	0	0	2	99
22	0	0	0	0	0	0	0	7	80	0	0	5	1	2	0	0	1	0	0	3	99
24	0	0	1	1	0	0	0	4	0	90	0	0	0	2	0	0	0	0	1	0	99
25	0	0	0	5	0	0	0	5	0	7	78	2	0	0	0	0	1	0	0	1	99
26	0	0	0	2	0	1	3	4	1	7	1	75	1	0	0	0	3	0	0	1	99
27	0	0	0	0	0	0	0	0	3	6	0	0	90	0	0	0	0	0	0	0	99
28	0	0	0	0	0	3	1	0	3	8	0	0	9	72	0	0	0	0	0	3	99
30	0	0	0	0	0	0	0	0	0	7	0	0	0	0	85	7	0	0	0	0	99
31	0	0	0	0	0	2	0	1	0	8	0	1	0	0	4	82	0	0	1	0	99
32	2	1	0	1	2	0	0	1	0	7	1	4	1	0	0	0	78	1	0	0	99
33	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	96	1	0	99
34	0	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	1	96	0	99
35	0	0	0	0	0	0	0	2	3	0	1	2	1	2	0	1	0	0	0	87	99
Total	101	99	97	151	86	104	70	103	90	143	93	93	103	82	89	92	85	99	101	99	1980

Table 2. Confusion matrix for all 20 UEs at 15dB SNR

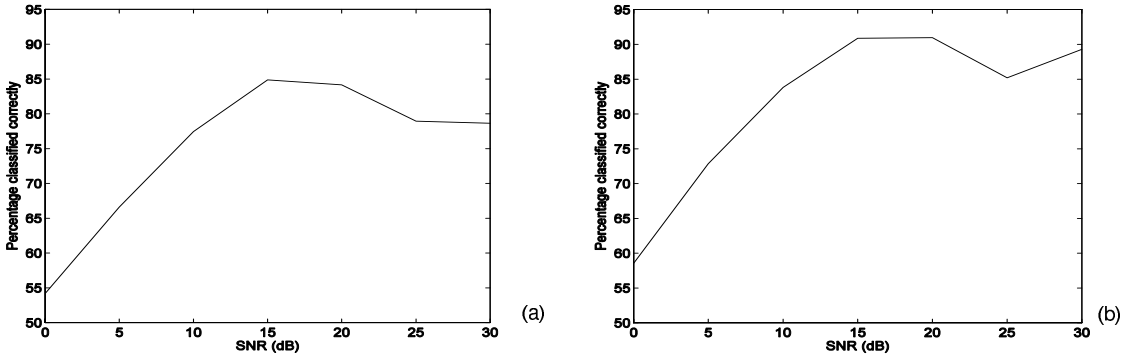

Figure 8. Best classification performance against SNR.

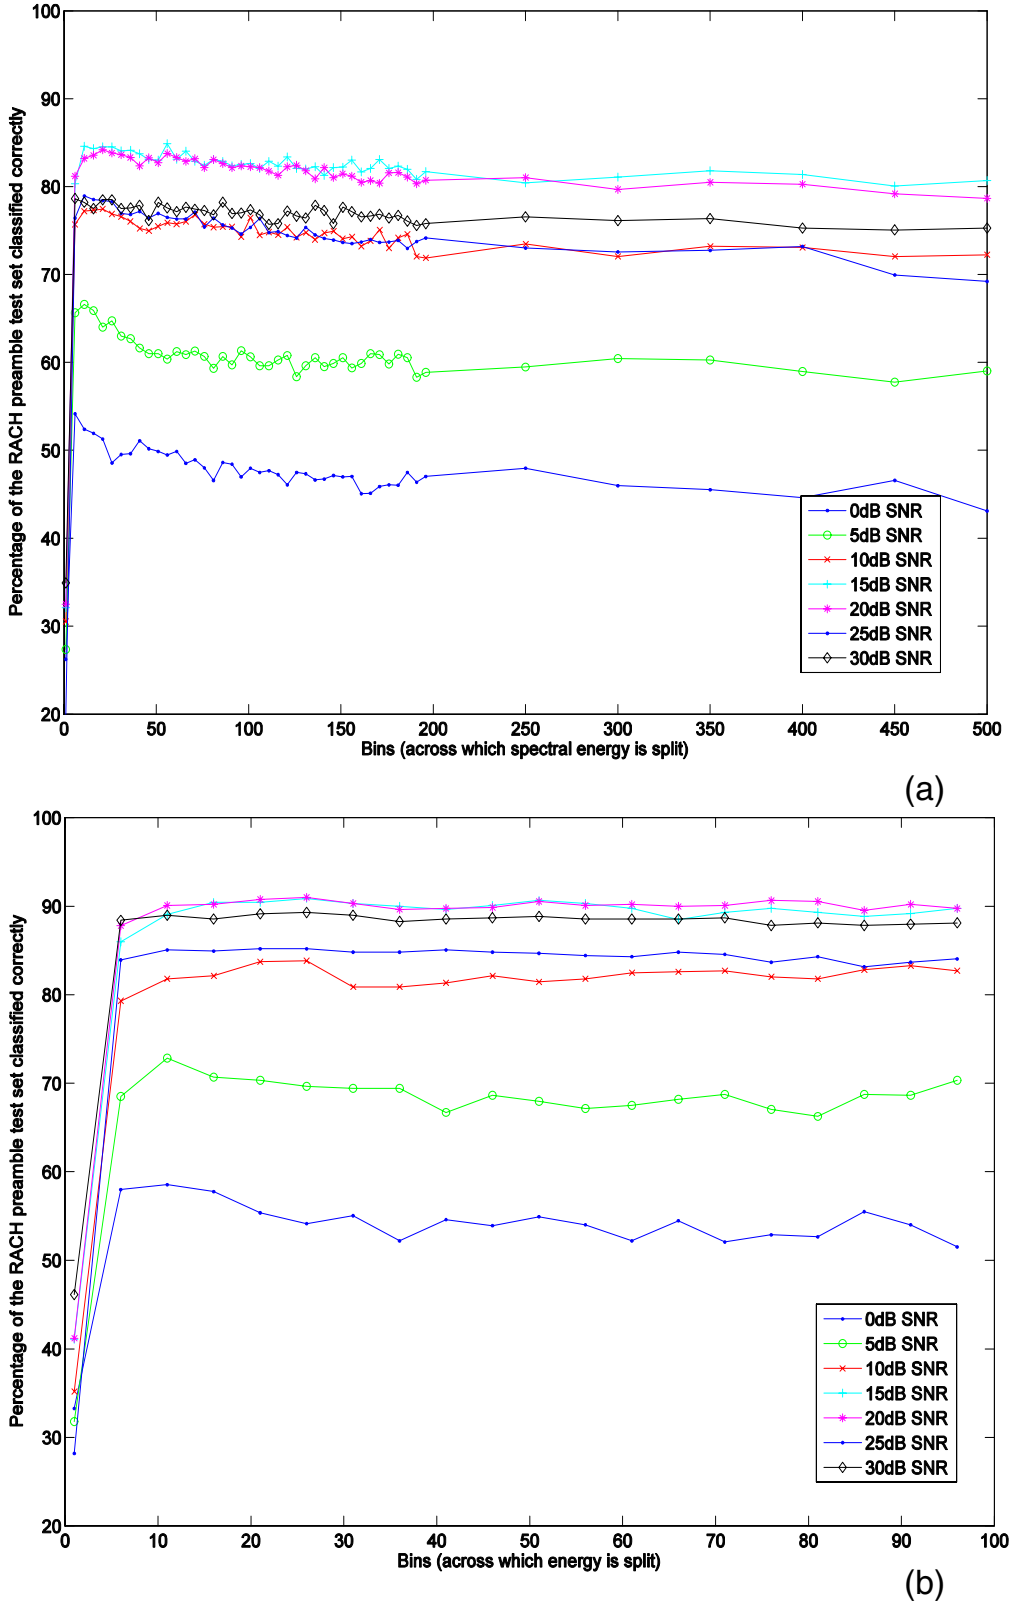

Figure 9. Best classification performance against SNR.

7. Conclusions

We have presented a novel, low-cost approach to transmitter identification using RF Fingerprinting. Our approach performs very well - being able to distinguish between seven different model transmitters with 91% accuracy at 15dB SNR. We report on the largest known test case in the literature, where we show good identification performance within a group of 20 UMTS UEs. We achieve 85% identification performance at 15dB, even with 50% of UEs being of the exact same model. Our use of the k-NN discriminatory classifier automates creation of the classification engine and the use of the FFT introduces great flexibility into spectral feature selection. Our system is capable of working with common low-cost receiver architectures with no hardware modifications. It therefore offers a lower cost solution to previously proposed transient based approaches which require very high speed ADCs.

From a UMTS systems perspective these results are very encouraging. To reduce the amount of core network signaling, the challenge is to determine whether a phone belongs to a particular group of phones or not. This is a verification problem rather than an identification problem. Obviously we will use very similar DSP techniques to explore the verification problem, but our measure of system performance will be quite different.

In future work we plan to examine the UE verification problem in more depth, investigate techniques to improve performance in low SNR environments and other non-ideal environments.

8 Acknowledgments

This work has received support from IDA Ireland.

References

[1] –. Promoting Efficient Use of Spectrum Through Elimination of Barriers to the Development of Secondary Markets. *Federal Communications Commission FCC 04-167*, Sept 2004.

[2] M. Buddhikot. Understanding Dynamic Spectrum Access: Models, Taxonomy and Challenges. *Proceedings of IEEE DySPAN07*, April 2007.

[3] H. Claussen. Performance of Macro and Co-channel Femtocells in a Hierarchical Cell Structure. *Proceedings of IEEE Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2007)*, 2007.

[4] C. Corbett, R. Beyah, and J. Copeland. A Passive Approach to Wireless NIC Identification. *Proceedings IEEE International Conference on Communications*, 2006.

[5] K. Ellis and N. Serinken. Characteristics of radio transmitter fingerprints. *Journal of Radio Science*, pages 585–597, 2001.

[6] R. Gerdes, T. Daniels, M. Mina, and S. Russell. Device Identification via Analog Signal Fingerprinting: A Matched Filter Approach. *ISOC Network and Distributed System Security Symposium*, 2006.

[7] J. Hall, J. Barbeau, and E. Kranakis. Detection of Transient in Radio Frequency Fingerprinting using Signal Phase. *Proceedings Wireless and Optical Communications*, 2003.

[8] J. Hall, M. Barbeau, and E. Kranakis. Detecting rogue devices in Bluetooth networks using Radio Frequency Fingerprinting. *Proceedings of the International Conference on Communications and Computer Networks*, 2006.

[9] L. Ho and H. Claussen. Effects of User-Deployed, Co-Channel Femtocells on the Call Drop Probability in a Residential Scenario. *IEEE International Symposium on Personal, Indoor and Mobile Communications*, 2007.

[10] N. Instruments. National Instruments VISA. URL, 2008. http://www.ni.com/visa/.

[11] J.Franklin, D. McCoy, P. Tabriz, V. Neagoe, J. V. Randwyk, and D. Sicker. Passive Data Link Layer 802.11Wireless Device Driver Fingerprinting. *Proceedings of the 15th USENIX Security Symposium*, pages 167–178, 2006.

[12] R. Jones. *Most Secret War*. Hamilton, 1978.

[13] T. Kohno, A. Briodo, and K. Claffy. Remote Physical Device Fingerprinting. *IEEE Transactions on Dependable and Secure Computing*, pages 93–108, 2005.

[14] I. Nabney and C. Bishop. Netlab Toolbox. URL, 2004. http://www.ncrg.aston.ac.uk/netlab/index.php.

[15] D. Shaw and W. Kinsner. Multifractal Modelling of Radio Transmitter Transients for Classification. *Proceedings Conference on Communications, Power and Computing*, pages 306–312, 1997.

[16] O. Tekbas, N. Serinken, and O. Ureten. An experimental performance evaluation of a novel radio-transmitter identification system under diverse environmental conditions. *Canadian Journal Computer Engineering*, 2004.

[17] O. Ureten and N. Serinken. Bayesian detection of transmitter turn-on transients. *Proceedings NSIP99*, pages 830–834, 1999.

[18] O. Ureten and N. Serinken. Detections of radio transmitter turn-on transients. *Electronic Letters*, pages 1996–1997, 1999.