

DETECTING ROGUE DEVICES IN BLUETOOTH NETWORKS USING RADIO FREQUENCY FINGERPRINTING

Jeyanthi Hall Michel Barbeau and Evangelos Kranakis

School of Computer Science

Carleton University

1125 Colonel By Drive

Ottawa, Ontario, Canada

email: jeyanthihall@rogers.com, barbeau.kranakis@scs.carleton.ca

ABSTRACT

Unauthorized Bluetooth devices or rogue devices can impersonate legitimate devices through address and link key spoofing. Moreover, they can infiltrate a Bluetooth network and initiate other forms of attacks. This paper investigates a novel intrusion detection approach, which makes use of radio frequency fingerprinting (RFF) for profiling, Hotelling's T^2 statistics for classification and a decision filter, for detecting these devices. RFF is a technique that is used to uniquely identify a transceiver based on the transient portion of the signal it generates. Moreover, the use of a statistical classifier proves advantageous in minimizing requirements for memory. Finally, the detection rate is also improved by incorporating a decision filter, which takes the classification results of a set of events into consideration, prior to rendering the final decision. The average False Alarm Rate of five percent and Detection Rate of ninety-three percent support the feasibility of employing these components to address the aforementioned problem.

KEY WORDS

Bluetooth rogue devices, Intrusion Detection, Radio Frequency Fingerprinting, Network Security, Wireless Networks, Hotelling's T^2 statistics.

1 Introduction

Of the many forms of attacks, which continue to be problematic in WiFi/802.11 or cellular networks, the rogue access points (RAP) is considered to carry a high risk, as indicated by Barbeau, Hall and Kranakis [1]. A RAP is first programmed with the identifier of an authorized AP. It is then deployed in order to obtain confidential information, such as passwords and credit card numbers, from unsuspecting victims. This information is used to impersonate users and to initiate various attacks on the network.

Bluetooth (BT) networks [2] are by no means impervious to this form of attack. Within this domain, an unauthorized rogue device R can be programmed to assume the identity of a legitimate device. For example, let us consider a scenario, whereby a BT device M (master) communicates periodically with one or more authorized devices (slaves). Moreover, as they are memory-constrained, their unit keys

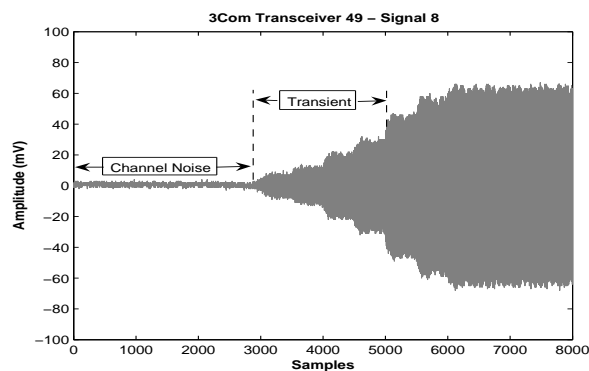


Figure 1. Signal from a BT transceiver

serve as link keys for the purpose of authentication and encryption. Finally, in order to expedite the authentication process in subsequent sessions, device M stores the link keys of all the slaves.

As aforementioned, a device impersonation attack proceeds as follows. In order to initiate the attack, device R requires the address and unit key of one of the slaves, e.g. device A. The first requirement is fulfilled by capturing and programming the 48-bit address of device A into device R, as demonstrated by Hager and Midkiff [3]. As far as the unit key is concerned, it is obtained, by device R, by initiating and establishing a communication link with device A. At the end of the session, device R would have obtained the unit key of device A, which is used henceforth as the link key [4]. Equipped with both pieces of information, device R is now capable of authenticating itself, as device A, to M, and thus gaining access to the network.

Given that the likelihood of this attack is possible, i.e. no major technical challenges to overcome, and the impact can be high, e.g. initiation of denial of service attacks, the resulting risk can be considered major. According to the specification from ETSI [5], this level of risk warrants the development and implementation of countermeasures.

What would prove useful is a mechanism for detecting rogue devices, in BT networks. Pioneered by the military to track the movement of enemy troops and subsequently implemented by some cellular carriers (e.g. Bell

Nynex) to combat cloning fraud [6], radio frequency fingerprinting (RFF) has been used to uniquely identify a given transceiver, based on its transceiverprint. A transceiverprint consists of features, which have been extracted from the turn-on transient portion of a signal [7]. Figure 1 illustrates the location of the transient, using a signal from a 3Com BT transceiver. The x-axis represents a sixteen μsec window of the signal, which has been sampled at a rate of 500 million samples per second. What can also be observed are the three segments of the signal, namely channel noise, transient and data transmission. The key benefit of this technique is that a transient reflects the unique hardware characteristics of a transceiver and thus cannot be easily forged, unless the entire circuitry of a transceiver can be accurately replicated, e.g. by theft of an authorized device.

In this paper, a novel approach, which makes use of RFF, Hotelling's T^2 statistics (a statistical classifier) and a decision filter, is proposed for anomaly-based intrusion detection (ABID) in BT networks. First, RFF is used to create a profile of each authorized BT device/transceiver and to associate each profile with the address of the corresponding device. Next, the statistical classifier is invoked in order to determine if an observed transceiverprint is normal, i.e. it matches the profile of an authorized device with a given address. Once a set of transceiverprints have been classified, the decision filter is applied. If a predefined percentage of these transceiverprints have been classified as normal, then there is a high probability that the signals did originate from an authorized device. Otherwise, an intrusion by a rogue device, with a spoofed address, is suspected.

An application of this technique to the scenario, presented previously, is as follows. Once the profile of device A (slave) has been created and stored in the device database [8] within device M (master), it is cross referenced with the address of device A. This configuration permits device M's security manager, which is responsible for ensuring a predefined level of security, to perform both device authentication and intrusion detection. Whereas device authentication is carried out using a challenge-response mechanism, intrusion detection is initiated by invoking the classifier.

The remaining sections of the paper are organized as follows. The details of the proposed ABID components are presented in Section 2, followed by evaluation results in Section 3. Section 4 briefly summarizes other related work in the area of RFF. Finally, the conclusions drawn are reported in Section 5.

2 Novel Approach: ABID using RFF

This section describes the framework and key activities that are undertaken to fulfill two primary objectives: the creation of a profile for each BT transceiver and specification of the classification system.




Figure 2. Test case for TDPC

2.1 Intrusion Detection Framework

The intrusion detection framework is designed to classify an observed transceiverprint as normal or anomalous. A classification result of *normal* indicates a strong probability that the underlying signals did originate from the transceiver of an authorized device. Likewise, an *anomalous* result serves as an indicator of a potential intrusion. When an alarm is raised, an appropriate response, based on the security requirements of applications, is initiated.

The flow of information begins with the conversion of an analog signal to a digital signal (will not be covered in detail). Once in a digital form, the transient portion of the signal is extracted by the *transient extractor*. Upon isolating the transient, the amplitude, frequency and phase components of the transient are extracted by the *feature extractor*. In turn, these components are used for the extraction of specific features that define a transceiverprint. The *statistical classifier* is used to determine if a given transceiverprint is normal or anomalous. Finally, the *decision filter* is applied to the classification results of a set of transceiverprints, in order to render a final decision regarding the status, e.g. authorized or intruder, of a BT device.

A transceiver profile is created by extracting the transceiverprints from a subset of the digital signals and storing the corresponding centroid and covariance matrix [9]. This exercise is undertaken prior to the classification or detection process. In the case of BT networks, the profiles could be created during the pairing process.

2.2 Transient Extractor

As the unique characteristics of transceivers are manifested in the transient portion of a signal, a key objective is to isolate and to extract the transient. The challenge, however, is identifying the starting point of the transient. We provide a brief overview of the technique, referred to as Transient Detection using Phase Characteristics (TDPC) by Hall, Barbeau and Kranakis [7]. It exploits the phase characteristics of a signal. As far as the end point of a transient is concerned, it is identified in an experimental manner and in consultation with the BT specification [10].




Figure 3. Histogram of detection errors for TDPC

Figure 2 depicts the application of TDPC using signal 8 from transceiver 49. More specifically, the second plot supports the fact that the difference in phase variance becomes constant at the starting point of a transient, a key requirement for successful detection. As indicated by the vertical line, in the first plot, the performance of TDPC is good. The expected starting point is within 150 samples prior to the actual value.

Figure 3 presents the distribution of detection errors. The x-axis represents the spectrum of detection errors, which signify the difference between the actual starting point of transients and the results produced by this algorithm. The mean of 21 and standard distribution of 51 for TDPC support the use of phase variance for transient detection. As a side note, the actual starting points are established through visual inspection, due to the absence of highly robust algorithms and the need to calculate detection errors in a precise manner.

2.3 Feature Extractor

Once the transient has been isolated, the next requirement is to extract the three primary components. We have opted to make use of all three components in order to enhance the characterization of transceivers. Figure 4 presents those associated with signal 8 from transceiver 49.

The amplitude envelope and instantaneous phase are obtained using standard algorithms [11] and are illustrated in the first and second plot respectively. The preferred approach for obtaining the frequency characteristics of a non-stationary signal, e.g. transient, is the application of the Discrete Wavelet Transform (DWT) [12]. Due to its low computational complexity, as defined by Choe et al. [13], the Daubechies filter is used to obtain the DWT coefficients, depicted in the third plot. This information proves useful in detecting variations in the frequency spectrum of devices.

Once these components have been extracted, a feature vector, also referred to as a transceiverprint, is created using a set of features F1-F15. Details re-

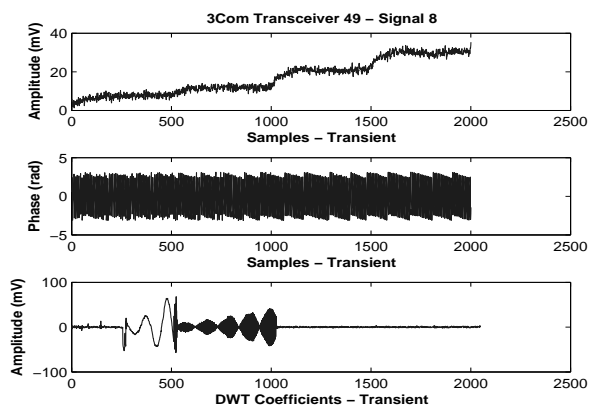


Figure 4. Components of a transient

garding the selection of these features are available in the paper by Hall, Barbeau and Kranakis [14]. In order to refrain from a detailed treatment of each feature, an intuitive interpretation of three features is provided next:

Normalized DWT coefficients (F1) This feature represents the standard deviation of the normalized amplitude of DWT coefficients. Normalization is achieved using the maximum amplitude of these coefficients. **Power per section (F8)** For this feature, power is calculated for each consecutive segment of a signal. This provides a trajectory of the power level during the turn-on ramp. **Normalized DWT coefficients by levels** Unlike F1, the standard deviation of the normalized amplitude of coefficients is obtained for levels one to six. Whereas level six is associated with the highest range of frequencies, specificities of the lower frequency bands can be identified using the remaining levels. What is interesting is that the level of details increases as the levels become lower.

2.4 Profile Definition

In order to classify a transceiverprint, a profile of the source is required. A subset of the transceiverprints, obtained from the captured signals, is selected using k-means clustering. This technique creates k subsets (or clusters) of data elements using the initial set of data. A number of iterations are executed whereby elements from one cluster are moved to another. This process continues until the similarity measure, i.e. the difference between the average data element and others, within each cluster is optimal. We employ this technique to select a representative set (one from each cluster) of transceiverprints. This set is used to create the key elements of a profile: the centroid and covariance matrix.

These two elements collectively represent the *intra-transceiver* (within a transceiver) and *inter-transceiver* (between transceivers) variability, which dictate the False Alarm Rate (FAR) and Detection Rate (DR). FAR is defined as the number of transceiverprints from X classi-




Figure 5. Inter-transceiver Variability

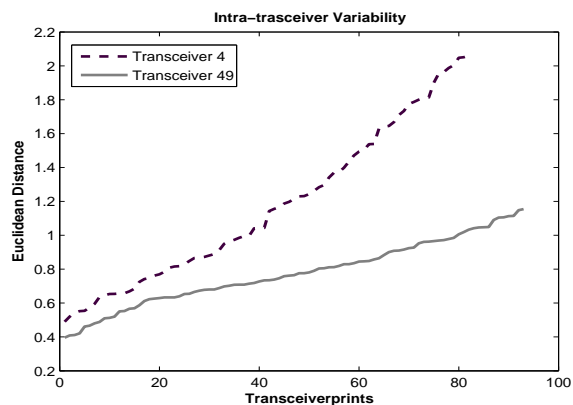


Figure 6. Intra-transceiver Variability

fied as belonging to Y divided by the total number of transceiverprints from X . On the other hand, DR is defined as the number of transceiverprints from Y classified as an intruder (from X 's perspective) divided by the total number of transceiverprints from Y .

As far as the inter-transceiver variability is concerned, the greater the variability, the higher the DR. This type of variability can be assessed by comparing the mean (centroid) and/or standard deviation of each of the features of the profiled transceiverprints.

Figure 5 illustrates the inter-transceiver variability between transceiver 49 from 3Com and transceiver 4 from Ericsson. What can be observed is that features 4, 6, and 8-13 exhibit different characteristics, and hence contribute, to a higher degree, towards inter-transceiver variability. This level of variability permits the classifier to distinguish between these two transceivers.

While high inter-transceiver variability is desirable for obtaining a high DR, it is also essential to have low intra-transceiver variability, a prerequisite for low FAR.

Figure 6 depicts the range of intra-transceiver variability of the transceiverprints, associated with transceivers

49 and 4. In particular, the variability between transceiverprints is established by obtaining the Euclidian Distance (ED) between the centroid and each of the transceiverprints. In the case of transceiver 49, the short range (0.4 to 1.0) and gradual slope are indicative of the consistency of its signals. These characteristics will prove useful in minimizing the FAR. On the other hand, transceiver 4 is characterized by a wide range (0.5 to 2) and a steeper slope.

2.5 Statistical Classifier

After having created the transceiver profiles and having obtained a transceiverprint, the statistical classifier is invoked. The Hotelling's T^2 statistics [9], is used for determining the degree of similarity between an observed transceiverprint and the profile of a BT transceiver, with a given address.

The memory per profile (MPP) of this classifier is very modest and is defined by Eq.1

$$MPP(m, n) = mn + m(n^2) \quad (1)$$

where n is the number of features, m is the size in bytes, and mn and mn^2 represent the memory requirement for the centroid and covariance matrix respectively.

In order to determine whether or not a given degree of similarity is normal, the T^2 value is transformed to follow an F distribution. The transformation is carried out by multiplying it by $n(n-p)/(p(n+1)(n-1))$, where n represents the sample size and p is the number of features in a transceiverprint. If the transformed value is greater than the F value of 2.20 (for a ninety-five percent confidence interval), the transceiverprint is classified as anomalous.

2.6 Decision Filter

In a wireless environment, characterized by noise and interference, there is a potential for increased variability between signals that are transmitted by the same transceiver. Hence, a decision, based on the classification results of a *single* transceiverprint, is likely to produce sub optimal results. Therefore, a decision filter is used to compensate for this type of error. First, a *set* of transceiverprints is independently classified. Then, the filter is applied to the classification results of the entire set. If eighty percent or more (used in this iteration) of the transceiverprints, has been classified as normal, then a final decision of normality is rendered. Typically, this threshold would be established based on the specific requirements of the application.

3 Evaluation

The purpose of the evaluation exercise is to assess the composition of a transceiverprint based on the classification success rate. The following steps were carried out using a set of signals captured from BT transceivers:

For each transceiver being profiled, the aforementioned features were extracted from the transients. Once

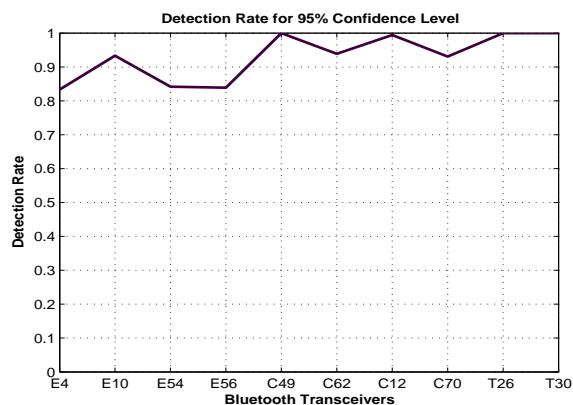


Figure 7. Detection Rate

the outliers (approximately 5-10) were removed, a subset (approximately 30-40) of the transceiverprints was selected using k-means clustering and subsequently used to calculate a centroid and covariance matrix. The remaining transceiverprints (50) were used for testing purposes.

3.1 Details of Evaluation

In order to evaluate the profiling and classification aspects of the proposed technique, signals from each of the 10 BT transceivers (3COM-4, Ericsson-4, Test Radios-2) were captured. All subsequent processing and evaluations were carried out using the Matlab software and associated tools.

3.2 Evaluation Results - RFF and Statistical Classifier

The FAR and DR served as our primary metrics. Additionally, 40 iterations were used for the purpose of assessing these metrics.

False Alarm Rate

All of the 10 transceivers, with the exception of transceiver 4 from Ericsson (E4), have a FAR of zero percent. Unlike the high intra-transceiver variability of E4, see Figure 6, the others have a low to moderate level of variability. This characteristic permits the k-means clustering algorithm to select a set of transceiverprints, which accurately characterizes a transceiver via the centroid and covariance matrix. Finally, the overall mean and standard deviation are five percent and 0.03 respectively.

Detection Rate

Figure 7 illustrates the DR for the profiled transceivers. The x-axis represents the transceivers, which are identified using the term $M\#$. Whereas M represents the manufacturer, e.g. E=Ericsson, C=3Com and T=Test Radio, $\#$ is the identifier of the transceiver.

There are a few observations that are of interest. First, the overall mean of ninety-three percent and standard de-

viation of seven suggest the presence of inter-transceiver variability between the 10 BT transceivers. Second, the average DR for the test radios is the highest at one hundred percent, followed by those from 3Com at ninety percent and Ericsson at eighty-seven percent. These results support the different levels of both inter-transceiver and intra-transceiver variability. In particular, it is interesting to note that E4 has one of the lowest DR. This should not come as a surprise, since, as previously stated, it also has a large intra-transceiver variability. Finally, although not depicted in the figure, there is a degree of similarity between the transceivers, i.e. signals, from the same manufacturer.

We compare the performance of this technique with that proposed by Choe [13]. Although the underlying frameworks are different, the overall concept is similar to some degree. However, the number of profiled transceivers was limited to three (2-Motorola HT-220, 1-Motorola MX-330) in comparison to the 10 BT transceivers used in this research project. Despite the increased complexity, the average success rate, represented by the DR of ninety-three percent, is consistent with their rate of ninety-four percent.

4 Related Work

This section provides a brief overview of the various research initiatives that have been undertaken to address the requirements of the RFF process.

In the paper by Ellis and Serinken [15], the authors examine the amplitude and phase components of signals and arrive at the conclusion that all transceivers do possess some consistent features. The detection of transients, based on the variance in amplitude, is proposed by Shaw and Kinsner [16]. In terms of classification, different approaches have been proposed. In the paper by Somervuo and Kohonen [17], the authors make use of the Self-Organizing Map and a Learning Vector Quantization (LVQ) algorithm to support variable-length feature sequences used for classification. While the use of DWT coefficients is explored by Hippenstiel and Payal in [18], Toonstra and Kinsner [19] exploit the properties of genetic algorithms for classification purposes.

5 Conclusion

Based on preliminary evaluation results, i.e. average FAR of five percent and DR of ninety-three percent, the use of RFF, Hotelling's T^2 statistics, and a decision filter, for anomaly-based intrusion detection in BT networks, is technically feasible.

More specifically, the characterization of transceivers, using multiple features, has resulted in a high DR. In addition, the use of a statistical classifier, that is memory conscious, could achieve sufficient performance for supporting various applications or services in BT networks. Finally, delaying the final decision until a sufficient number

of transceiverprints have been classified, increases both the confidence level and classification success rate.

Nevertheless, there are some issues, which warrant further attention. First and foremost, the success rates can be further increased by optimizing the composition of the transceiverprints and validating them using a larger set of transceivers from the same manufacturer. Second, it would prove useful to repeat the profiling exercise periodically in order to determine the impact of various factors, e.g. transceiver aging, on the classification success rate. Third, as far as scalability is concerned, further research is required to determine the maximum number of transceiver profiles, which can be supported by a node in a BT network. Finally, field tests should be carried out in order to assess the true performance of the proposed intrusion detection framework, in particular, the digital signal processing component.

6 Acknowledgments

The authors graciously acknowledge the financial support received from the following organizations: Natural Sciences and Engineering Research Council of Canada (NSERC) and Mathematics of Information Technology and Complex Systems (MITACS). In addition, the authors also acknowledge the support received from Ericsson and 3COM.

References

- [1] M. Barbeau, J. Hall, and E. Kranakis. Detecting Impersonation Attacks in Future Wireless and Mobile Networks. In *Proceedings of MADNES 2005 - Workshop on Secure Mobile Ad-hoc Networks and Sensors - Held in conjunction with ISC'05*, Singapore, September 20-22 2005. SVLNCS.
- [2] Bluetooth Special Interest Group. Specification of the Bluetooth System Version 1.0B. <http://www.bluetooth.org>, December 1999. Accessed August 25 2004.
- [3] C.T. Hager and S.F. Midkiff. Demonstrating vulnerabilities in bluetooth security. In *Proceedings of the Conference on GLOBECOM*, pages 1420–1424. IEEE, 2003.
- [4] C.T. Hager and S.F. Midkiff. An analysis of Bluetooth security vulnerabilities. In *Proceedings of the Conference on Wireless Communications and Networking*, volume 3, pages 1825–1831. IEEE, March 2003.
- [5] ETSI. Telecommunications and internet protocol harmonization over networks TIPHON release 4; protocol framework definition; methods and protocols for security; part 1: Threat analysis. Technical Specification ETSI TS 102 165-1 V4.1.1, 2003.
- [6] M. J. Riezenman. Cellular security: better, but foes still lurk. *IEEE Spectrum*, pages 39–42, June 2000.
- [7] J. Hall, M. Barbeau, and E. Kranakis. Detection of transient in radio frequency fingerprinting using signal phase. In *Wireless and Optical Communications*, pages 13–18. ACTA Press, July 2003.
- [8] Bluetooth Special Interest Group. Bluetooth Security Architecture volume 1. <http://www.bluetooth.org>, 1999. Accessed in January 2006.
- [9] M. Natrella. NIST/SEMATECH e-Handbook of Statistical Methods. <http://www.itl.nist.gov/div898/handbook>, 2001. Accessed in June 2006.
- [10] Bluetooth Special Interest Group. Bluetooth Radio Specification part a. <http://www.bluetooth.org>, February 2001. Accessed in January 2006.
- [11] J. G. Proakis and D. G. Manolakis. *Digital Signal Processing*. Prentice Hall PTR, 1996.
- [12] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [13] H. Choe, C.E. Poole, A.M. Yu, and H.H. Szu. Novel identification of intercepted signals from unknown radio transmitters. *SPIE*, 2491:504–516, 1995.
- [14] J. Hall, M. Barbeau, and E. Kranakis. Enhancing intrusion detection in wireless networks using radio frequency fingerprinting. In *Proceedings of the 3rd IASTED International Conference on Communications, Internet and Information Technology (CIIT)*, pages 201–206, St. Thomas, U.S. Virgin Islands, November 2004.
- [15] K.J. Ellis and N. Serinken. Characteristics of radio transmitter fingerprints. *Radio Science*, 36:585–597, 2001.
- [16] D. Shaw and W. Kinsner. Multifractal modelling of radio transmitter transients for classification. In *Communications Power and Computing*, pages 306–312, Winnipeg Manitoba, May 1997. IEEE.
- [17] P. Somervuo and T. Kohonen. Self-Organizing Maps and Learning Vector Quantization for Feature Sequences. *Neural Processing Letters*, 10:151–159, 1999.
- [18] R.D. Hippenstiel and Y. Payal. Wavelet based transmitter identification. In *International Symposium on Signal Processing and its Applications*, Gold Coast Australia, August 1996.
- [19] J. Toonstra and W. Kinsner. Transient analysis and genetic algorithms for classification. In *WESCAN*. IEEE, 1995.