



(51) International Patent Classification:

H04B 17/20 (2015.01) H04W 12/12 (2009.01)

G01R 23/16 (2006.01) H04W 24/08 (2009.01)

(21) International Application Number:

PCT/US2018/022639

(22) International Filing Date:

15 March 2018 (15.03.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/472,387 16 March 2017 (16.03.2017) US

(71) Applicants: BATTELLE ENERGY ALLIANCE, LLC

[US/US]; Idaho National Laboratory P.O. Box 1625, Idaho Falls, Idaho 83415 (US). UNIVERSITY OF UTAH RESEARCH FOUNDATION [US/US]; 615 Arapen Dr., Suite 310, Salt Lake City, Utah 84108 (US).

(72) Inventors: DERR, Kurt W.; 151 Fair Hills Circle, Idaho

Falls, Idaho 83401 (US). RAMIREZ, Samuel; 841 N. 1100

E., Shelley, Idaho 83274 (US). KASERA, Sneha K.; 2963 Brookburn Road, Salt Lake City, Utah 84109 (US). BECKER, Christopher D.; 404 4th St. NW, Aitkin, Minnesota 56431 (US). BASET, Aniqua Z.; 224 University Village, Salt Lake City, Utah 84108 (US).

(74) Agent: BACA, Andrew J. et al.; 230 South 500 East, Suite 300, Salt Lake City, Utah 84102 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

(54) Title: SYSTEM, METHOD, AND APPARATUS FOR WIRELESS FREQUENCY SIGNAL IDENTIFICATION AND PROTOCOL REVERSE ENGINEERING




FIGURE 2A

(57) Abstract: Wireless signal classifiers and systems that incorporate the same may include an energy-based detector configured to analyze an entire set of measurements and generate a first single classification result, a cyclostationary-based detector configured to analyze less than the entire set of measurements and generate a second signal classification result; and a classification merger configured to merge the first signal classification result and the second signal classification result.

GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

- 1 -

**SYSTEM, METHOD, AND APPARATUS FOR WIRELESS
FREQUENCY SIGNAL IDENTIFICATION AND
PROTOCOL REVERSE ENGINEERING**

5

PRIORITY CLAIM

This application claims the benefit of the filing date of United States Provisional Patent Application Serial No. 62/472,387, filed March 16, 2017, for “System, Method, and Apparatus for Wireless Frequency Signal Identification and Protocol Reverse Engineering.”

10

**STATEMENT REGARDING FEDERALLY
SPONSORED RESEARCH OR DEVELOPMENT**

The invention was made with government support under Contract No. DE-AC07-05-ID14517, awarded by the United States Department of Energy. The government has certain rights in this invention.

15

TECHNICAL FIELD

Embodiments of the present disclosure relate, generally, to systems and methods for identifying wireless signals and protocols, and more particularly, systems and methods for protocol reverse engineering of wireless signals.

20

BACKGROUND

Wireless communications technology is becoming ubiquitous throughout society. Although Wi-Fi has grown to be the ubiquitous Internet access technology, many other wireless protocols are used, for example, wireless communication systems such as Bluetooth, Wi-Fi, cellular, Apple iBeacon, Z-Wave, and ZigBee. Wireless communications devices are widely used in residential homes, in public safety, emergency response, and critical infrastructure applications.

25

Widespread use of wireless technology raises security concerns. Unauthorized third parties may attempt to access or intrude into wireless devices and wireless networks illegally. If third-parties do access wireless devices or wireless networks, the security of computers and data is at risk.

30

- 2 -

DISCLOSURE

Some embodiments of the present disclosure relate to a computer-implemented wireless signal classification method. The method may include: receiving a first wireless signal classification, the first wireless signal classification based on blocks of radio frequency (RF) measurements of a wireless spectrum over a period of time; receiving a second wireless signal classification, the second wireless signal classification based on part of the blocks of RF measurements; weighting the first wireless signal classification and weighting the second wireless signal classification; and merging the weighted first wireless signal classification and the weighted second wireless signal classification to arrive at a classification result.

Some embodiments of the present disclosure relate to a system. The system may include an energy-based detector configured to analyze an entire set of measurements and generate a first single classification result; a cyclostationary-based detector configured to analyze less than the entire set of measurements and generate a second signal classification result; and a classification merger module configured to merge the first signal classification result and the second signal classification result.

BRIEF DESCRIPTION OF THE DRAWINGS AND EXHIBITS

Purposes and advantages of the embodiments of the present disclosure will be apparent to one of ordinary skill in the art from the specification in conjunction with the appended

Drawings and Exhibits:

FIGURE 1 is a block-diagram of a classification node according to an embodiment of the disclosure.

FIGURE 2A is a block-diagram of a classification and detection system according to an embodiment of the disclosure.

FIGURE 2B is a block-diagram for an energy-based detection path according to an embodiment of the disclosure.

FIGURE 2C is a block-diagram of a cyclostationary-based detection path according to an embodiment of the disclosure.

FIGURE 3 is a block-diagram of a classification and capture system according to an embodiment of the disclosure.

- 3 -

FIGURE 4 is a block-diagram of a protocol reverse engineering system according to an embodiment of the disclosure.

FIGURE 5A is a block-diagram of a classification and detection system according to an embodiment of the disclosure.

5 FIGURE 5B is a block-diagram of an energy-based detection path according to an embodiment of the disclosure.

FIGURE 5C is a block-diagram of a machine-learning-based detection path according to an embodiment of the disclosure.

10 MODE(S) FOR CARRYING OUT THE INVENTION

Several situations have been identified that give rise to security concerns for wireless networks and wireless communications:

First, wireless systems are deployed in critical infrastructures and the vulnerabilities of these wireless systems increases the vulnerability to these sectors and of the economy. For
15 example, wireless communication is used in critical infrastructure (CI) applications for monitoring and providing data on the status of CI components and for intelligent transportation systems. Malicious actors may compromise existing wireless devices or implant rogue wireless devices (RWDs) to feed false data to the operators of an operator station creating the potential for a catastrophe.

20 Second, business entities, government facilities, critical infrastructure, and homes are not “aware” of the wireless activities that surround them, or the vulnerabilities that such activities enable. They may not be aware that they are being hacked or compromised, or that they experience wireless interference (intentionally or unintentionally).

Third, traditionally wireless messages have moved through the free-space environment
25 on certain spectrum allocations, which have been scarce, heavily regulated, and often unattainable resources. However secondary users are increasingly permitted to use licensed bands by way of dynamic spectrum allocation (DSA) system that enable the secondary use when the primary users are not using those bands. Secondary users (and primary users) may violate spectrum allocation policies (*e.g.*, regulations).

30 Fourth, the government or other organizations/entities may acquire so-called “black

- 4 -

box” devices having wireless behavior (blind protocol) that is unknown. While spectrum analyzers may be used to analyze specific frequency ranges, first a user must know the spectrum ranges to analyze. It is difficult to capture a signal over time by a spectrum analyzer, the volume of data is large, and post processing (analyze the signal afterwards) resource
5 intensive, accordingly, it is not possible to do it in real time analysis.

In high-security and control-system environments such as nuclear plants, power plants, military facilities, and other CI, the wireless signal types may be restricted to just a limited number of authorized types for security reasons. In such environments the presence of unauthorized wireless signals or the absence of authorized signals may indicate malicious
10 activities or a problem that must be addressed quickly to avoid a breach or system failure.

Similarly, the presence of unknown signals in an enterprise building might indicate malicious activities like the presence of wireless spying devices that may compromise an organization's confidential data and/or critical assets.

Accordingly, some embodiments of the present disclosure are related, generally, to an
15 efficient wireless signal classification system capable of detecting known signal types as well as unknown signals in real-time. In some embodiments, the classification system may operate in conjunction with or be incorporated in detection systems, which detect problems in timely manner, raise alerts, and/or take appropriate actions. Embodiments of a classification system may also be used to analyze black box devices to show that the devices act as intended (or as
20 indicated), with no additional signals or interference being generated. Other embodiments relate to a real-time wireless signal classification system used in, or operating in conjunction with, shared spectrum applications for detecting the presence of incumbent transmitters and/or spectrum usage violations by secondary users. Embodiments also have a number of additional functionalities, including but not limited to signal recording, blind signal analysis, signal
25 demodulation, signal localization, and protocol reverse engineering. As noted throughout the present disclosure, embodiments of the classification system may be used either as a standalone system or as a system integrated into other systems.

Having come to understand these and similar security concerns, embodiments of the present disclosure facilitate real-time monitoring and analysis of: (1) CI applications that rely
30 on wireless communication, (2) devices to detect possible spectrum violations, and (3)

- 5 -

wireless signals in general to detect and interact with RWDs. Other benefits and advantages also exist.

Various embodiments described herein may include elements described as implemented in a “computer” or a “computer system.” Here, the terms “computer” and
5 “computer system” are to be understood to include at least one non-transitory computer readable memory and at least one processing unit. In general, the memory will store, at one time or another, at least portions of an executable program code, and the processor(s) will execute one or more of the instructions included in that executable program code. It will be appreciated that the term “executable program code” and the term “software” mean
10 substantially the same thing for the purposes of this description. It is not necessary to the practice of the various embodiments described herein that the memory and the processor be physically located in the same place. That is to say, it is foreseen that the processor and the memory might be distributed among physical pieces of equipment or even in geographically distinct locations.

15 The processing unit may be a general purpose “central processing unit,” but may use a wide variety of other technologies. Specific-purpose hardware may also be used to implement the embodiments described herein, including a microcomputer, mini-computer, mainframe computer, programmed micro-processor, micro-controller, peripheral integrated circuit element, a CSIC (Customer Specific Integrated Circuit), ASIC (Application Specific
20 Integrated Circuit), a logic circuit, a digital signal processor, a programmable logic device such as an FPGA (Field Programmable Gate Array), PLD (Programmable Logic Device), PLA (Programmable Logic Array), RFID integrated circuits, smart chip, or any other device or arrangement of devices that is capable of implementing the steps of the processes of the invention. The processing unit may consist of a single core, or may be a multi-core processor
25 that has two or more processing units that can operate executing instructions independently in parallel.

Various embodiments of the disclosure may be described in terms of “algorithms” “module(s)” and “engine(s),” which refer to the logic, embodied in hardware and/or software, to accomplish the features, functions, tasks or steps described herein. In the case of a general
30 purpose computer, the “modules” and “engines” may be embodied in software classes and

- 6 -

applications executed by processor cores, and while the modules or engines are executing as instruction on a non-transitory storage medium a general purpose computer may be thought of as a special purpose computer or a specific purpose computer. The “modules” and “engines” may also relate to specific purpose hardware, including the firmware and machine code, controlling its operation.

Further, when embodied in software, the “modules” and “engines” that enable a computer system to act in accordance with the invention may be stored on non-transitory storage mediums in any number of language forms including, but not limited to, original source code, assembly code, object code, machine language, compressed or encrypted versions of the foregoing, and any and all equivalents. Some examples of languages that may be used to write the software include, but are not limited to, C, C++, JAVA, MATLAB, MINITAB, EXPRESS, DRAKON, DYNA, PYTHON, MOOSE, and RUBY. The software programs may be further translated into machine language or virtual machine instructions and stored in a program file in that form. The program file may then be stored on or in one or more of the articles of manufacture.

Users may interact with the computer systems described herein by way of graphical user interfaces (GUI) on a display and input devices such as touchscreens, keyboards, a computer mouse, touchpads, and the like.

Embodiments of the monitoring techniques described herein, generally, comprise signal detection and signal classification. Energy-based detection (EBD) is known to the inventors of the present disclosure to provide an efficient technique to detect signals. EBD may detect a signal based on the energy observed in a received signal. The detection process can be done in time-domain as well as frequency domain.

However, it is now understood by the inventors of the present disclosure that EBD is not always accurate due to difficulties in determining the noise floor and exhibits poor performance when finding signals close to or below the noise floor. An alternative approach to EBD is cyclostationary-based detection (CBD). Generally, CBD involves extracting cyclostationary features from a signal. Cyclostationary features are periodic characteristics of a signal that result from modulation, sampling, multiplexing, and/or coding operations. These characteristics are unique for signal types and may be used to distinguish among signal types.

- 7 -

Generally, cyclostationary features may be extract from a signal by correlating a signal with a delayed version of itself, wherein a high correlation will be seen when the delay is equal to a period of a Cyclostationary feature. Cyclostationary features of a signal may be represented by Spectral Correlation Functions (SCFs), which may be computed using a time smoothing method and FFTs, an FFT accumulation method, a strip spectral correlation analyzer, etc. Additional computations involving the SCF may be used to determine which shift(s), α , provides higher correlation. A resulting N-sized array may be referred to as the α -profile that includes maximum values for all possible shifts, α . Conventional CBD is very accurate, but computationally expensive, and so not well suited to real-time signal detection.

Various embodiments of the disclosure apply both EBD and a modified CBD (M-CBD) to signals and merge the results of each process to detect and classify the signals. Such embodiments are both efficient and accurate, and maintain real-time detection capabilities. With reference to FIG. 2A, in these embodiments, the EBD path 131 is used continuously, while the M-CBD path 132 is “on” periodically but (relative to EBD) infrequently. The results from the M-CBD path 132 are provided as feedback to the EBD path 131 to adjust parameters and improve detection accuracy. In addition, the M-CBD path 132 may detect signals not detected by the EBD path 131 (*e.g.*, because the signals are close to or below the noise floor).

FIG. 1 illustrates an embodiment of a classification node 100 according to an embodiment of the present disclosure. The classification node 100 comprises a software defined radio (SDR) 110 and a computer 120. In various embodiments, the architecture of the SDR 110 is of a type known to those of ordinary skill in the art, and each module in the SDR 110 may be implemented in software, hardware, an embedded system, and combinations thereof. Software components of the SDR 110 may be executed on a general purpose computer. By way of non-limiting example, the SDR 110 software may be based on GNU Radio, GNU Radio Companion, and GNU Radio Blocks. By way of non-limiting example, hardware implementations of the SDR 110 may be based on USRP B210, X310, HackRF One, and the like. Various embodiments of the SDR 110 may utilize application programming interfaces (API) from C++ and/or the software components associated with GNU Radio (or one of the other packages noted above) to interact with the hardware of X310 (or one of the other architectures noted above).

- 8 -

The SDR 110 outputs sampled RF signals to the computer 120. In one embodiment, the computer 120 includes classification logic, including, in one embodiment, the classification and detection system 130. Embodiments of the classification and detection system 130 are illustrated in FIGS. 2A, 2B and 2C. The classification and detection system 130 illustrated in FIG. 2A includes an EBD path 131 and an M-CBD path 132. The EBD path 131 is illustrated with more detail in FIG. 2B, and the M-CBD path 132 is illustrated with more detail in FIG. 2C. The computer 120 is not limited to sampled RF signals received from an SDR, and may receive measurement information from other spectrum measurement sources, including, by way of non-limiting example, other spectrum capture devices, files (*e.g.*, stored measurement information), network equipment (*e.g.*, received over a network), and the like.

In one embodiment, the EBD path 131 may include power spectral density (PSD) calculator 152, noise floor calculator 154, dynamic cutoff calculator 156, peak detector 158, bandwidth analyzer 160, timing analyzing 162, and pattern matcher 164.

In one embodiment, the PSD calculator 152 is configured to calculate the PSD from an FFT and pass the PSD to the noise floor calculator 154 as well as the peak detector 158.

In one embodiment, the noise floor calculator 154 may be configured to receive the PSD values (bins), find the lowest n values of that PSD and compute a mean and standard deviation of those n values, thereby enabling dynamic calculation of the noise floor continuously at runtime. The mean and standard deviation may be provided to the peak detector 158.

In one embodiment, the dynamic cutoff calculator 156 may be configured to determine a cutoff to remove noise from a PSD value. The cutoff used to remove noise from a PSD value may be $\mu + m \cdot \sigma$, where μ and σ are the mean and standard deviation values obtained from the noise floor calculator 154 module, respectively, and m is a multiplier which may be adjusted by feedback. The dynamic cutoff may be provided to the peak detector 158.

In one embodiment, the peak detector 158 may be configured to determine and send a mask of the bins that were above the specified cutoff to bandwidth analyzer 160 and timing analyzer 162.

In one embodiment, the bandwidth analyzer 160 may be configured to compute sets of

contiguous bins above the cutoff using the mask received from the peak detector 158. For every set of inputs the bandwidth analyzer 160 receives, it looks for contiguous sets of bins that are above a threshold by looking for contiguous 1's (or another predefined indicator) in the mask it received from the peak detector 158.

5 In one embodiment, the timing analyzer 162 may be configured to track active and inactive intervals for each bin, separately. A bin may be considered active if it was considered part of a signal by the peak detector 158, and a bin may be considered inactive otherwise. Information about the amount of time bins were active and inactive (after a state change) may be provided as output from the timing analyzer 162. In one embodiment, the timing analyzer
10 162 may include a two-state state machine (not shown) that it maintains for every entry in the mask it receives from the peak detector 158. By way of non-limiting example, responsive to a value of 1 for a particular mask entry, the state machine for that entry enters or maintains an active state. Likewise, responsive to a value of 0 for a particular mask entry, the state machine for that entry enters or maintains an inactive state. Responsive to a state transition between
15 states for an entry, an output entry for the state from which the entry left may be created and sent to an associated pattern matcher 164.

 In one embodiment, the pattern matcher(s) 164 may be configured to determine if a particular signal is present or not based on comparing the received bandwidth and/or timing results received from the previous blocks against a known set of parameters for the signal. By
20 way of non-limiting example, an IEEE 802.11g signal using OFDM uses approximately 16.6 MHz of spectrum, so the pattern is matched against bands found by the bandwidth analyzer. Similarly, timing patterns such as Short Interframe Spacing (SIFS) inactive intervals and active transmission times required for packets based on various data rates are compared against inactive and active timing information found by the timing analyzer 162. On the other
25 hand, the ZigBee pattern matcher takes just the bandwidth information from the bandwidth analyzer 160 and compares it against the expected 2 MHz of spectrum occupancy for ZigBee.

 In one embodiment, upon detecting a match, the pattern matcher(s) 164 may be configured to computer a classification score of the classification for a band (s, f) using Equation 1:

- 10 -

$$score_{s,f} = \frac{count_{s,f}}{\max_{b \in B} count_{s,f}} \quad \text{Equation 1}$$

Where $count_{s,f}$ is the number of times band s , f has appeared since the last update, and $\max_{b \in B} count_{s,f}$ is the maximum number of times any band has appeared since the last update.

For example, if band₁ has appeared one time since the last update and band₂ has appeared five
 5 times since the last update, the classification score of band₁ is $1/5=0.2$, while the classification score of band₂ is $5/5 = 1.0$. Results may be stored for a set time period, t_0 , before being provided to the merger 133. By way of non-limiting example, the results may include signal types, frequencies, scores, counts for detected signals, and the like.

Notably, embodiments of the peak detector 158 and protocol detector 166 may be
 10 configured based on specific wireless protocols (*e.g.*, 802.11(b), ZigBee). For multiple protocols, multiple peak detector modules, protocol detector modules and/or components thereof, each configured for a signal class may be used. For example, a first pattern matcher 164 may be configured with band and timing information for 802.11g, and another pattern matcher 164 may be configured with bandwidth information for ZigBee (which occupies a
 15 specific 2 MHz spectrum). In other embodiments, a classification system including an EBD path may include a protocol detector for each protocol.

Embodiments of the merger 133 of the classification and detection system 130 are configured to take the classification results from both the EBD path 131 and the M-CBD path 132 and merge the results for a final classification. The classification results may include a
 20 center frequency, a signal classification, count and a classification score. Both EBD path 131 and M-CBD path 132 may also provide information about the scores of the classifications they make. Score levels are in the range (0:0; 1:0] where a value around 0.0 corresponds to a very low confidence in a classification and a value near 1.0 corresponds to a very high confidence in the classification.

25 In some embodiments, the merger 133 may be configured to apply pre-defined score weights to results from the two paths to make a final classification. Since the M-CBD path 132 provides higher accuracy of signal classification than the EBD path 131, the merger 133 may be configured to assign a higher weight to the results from M-CBD path 132 while merging the results.

- 11 -

The merger 133 may be initialized with classification score weights for the different classification sources (*i.e.*, M-CBD and EBD classification) and an update rate. The merger 133 may be configured to switch between two states, a merging state and an update state.

The merger 133 may be in a merging state between updates. During the merging state,
 5 it receives classification results from the different sources consisting of signal classification, center frequency, classification score, and count. Values from different sources and of different signal class are kept separate, but values from the same source and signal class are merged. If multiple sets of results are obtained from the same source while the merger 133 is in the merging state then the results are combined/merged. In various embodiments,
 10 classifications are merged based on their signal class and center frequency. For every new value received, if the signal class and center frequency match the signal class and center frequency of an existing entry then the entries are merged by updating the score to an average of the existing entry and the new entry, and the count is updated to be the sum of the existing count and the new count. By way of non-limiting example, if an entry for a Wi-Fi signal is
 15 found by the M-CBD path 132 located at 2.437 GHz (*i.e.*, the center frequency), having a score of 0.99, and has a count of “3” (because it has been seen 3 times already), and a new entry for a Wi-Fi signal is found by the M-CBD path 132 located at 2.437 GHz, having a score of 0.85, and a count of “1,” the merged entry has a score of $(3 * 0.99 + 1 * 0.85) / 4 = 0.955$ and a count of “4.” By way of another non-limiting example, if a new entry does not match an
 20 existing entry in terms of center frequency, then the new entry is added to the set of existing entries.

If the merger 133 is in update state, then the previously merged results from the sources are merged to a final classification using the results and source classification score weights specified by the user. For each signal class, the paths that provide classification
 25 results are counted and become the number of sources for that signal class. By way of non-limiting example, if only the EBD path 131 or the M-CBD path 132 provided entries for Wi-Fi signals then $\text{count}(\text{sources}(\text{Wi-Fi})) = 1$. However, if both paths provided entries for Wi-Fi signals then $\text{count}(\text{sources}(\text{Wi-Fi})) = 2$. For each center frequency, f , for that signal class, the final merged entry has a score computed as shown in Equation 2, below:

$$30 \quad \text{score}(f) = \frac{1}{c} \times \sum_{s \in \text{sources}(t)} w_s \times \text{score}_s(f) \quad \text{Equation 2}$$

- 12 -

In Equation 1, t is the signal type, $c = \text{count}(\text{sources}(t))$, $w_s = \text{weight}(s)$ (the source weight as specified by the user), and scores $(f) = 0$ if the frequency was not reported by the source, otherwise, it is the classification score value from the merged entry. By way of non-limiting example, if the following entries existed from EBD path 131 and M-CBD path 132, respectively: $\{(2.437 \text{ GHz}, 0.99), (2.438 \text{ GHz}, 0.8)\}$ and $\{(2.437 \text{ GHz}, 0.97)\}$, and if the source weights for the EBD path 131 and M-CBD path 132 are 0.75 and 1.0, respectively, then the final classification entries are: $(2.437 \text{ GHz}, (0.75 * 0.99 + 1.0 * 0.97)/2 = 0.85625)$, and $(2.438 \text{ GHz}, (0.75 * 0.8 + 1.0 * 0.0)/2 = 0.3)$.

Regarding the form of the classification results output by the merger 133, in various embodiments of the disclosure the entries for each signal class from both sources are provided as results. By way of non-limiting example using generic signal terms, if the M-CBD path 132 finds Wi-Fi bands $\{1, 2, 3, 4\}$ and ZigBee bands $\{1, 2\}$, and the EBD path 131 finds Wi-Fi bands $\{1, 5, 6\}$ and ZigBee bands $\{3, 4\}$, then the final merged results would be Wi-Fi bands $\{1', 2, 3, 4, 5, 6\}$ and ZigBee bands $\{1, 2, 3, 4\}$. For the Wi-Fi bands, 1' is a merged entry.

In some embodiments, the merger 133 may provide feedback control signals to the EBD path 131 based on the results of the M-CBD path 132. Parameters of the Peak Detector maybe modified/adjusted based on the feedback. In one embodiment, differences in the classification results reached by the EBD path 131 and M-CBD path 132 are tracked. If the differences are exceed a threshold then commands are sent from the merger 133 to the peak detector module of the EBD path 131 and a cutoff multiplier used in peak detection is adjusted. By way of non-limiting example, if the EBD path 131 reports detection and classification of eight (8) Wi-Fi bands, but the M-CBD path 132 reports detection and classification of three (3) Wi-Fi bands, a command is sent to the peak detector module to increase the noise floor cutoff parameter by adjusting the multiplier. However, if the M-CBD path 132 reports detection and classification of eight (8) Wi-Fi bands and the EBD path 131 reports detection and classification of three (3) Wi-Fi bands, then a command is sent to the peak detection module to lower the noise floor cutoff parameter by adjusting the multiplier. This allows the classification and detection system 130 to adapt at run time and be less susceptible to bad initialization parameters.

FIG. 2C shows classification logic of the M-CBD path 132, in accordance with the

- 13 -

present disclosure. The M-CBD may include the following modules: a data reducer 134, a Spectral Correlation Function (SCF)/ α -profile calculator 140 (merely referred to herein as an α -profile calculator), and a signal classifier 142.

Regarding the data reducer 134, it is now understood by the inventors of the present disclosure that a CBD typically cannot keep-up with the high-sample rates of SDRs such as the SDR 110. By way of non-limiting example, an SDR such as SDR 110 may generate 2 GB of data every 11 second at a 25 MHZ sample rate. It is also now understood that lowering the sample rate would result in poor signal detection and classification by the system.

Various embodiments of the M-CBD path 132 may maintain the real-time detection capabilities by, in part, incorporating a data reducer 134. Embodiments of a data reducer 134 module (“keep M in N”) are illustrated in FIGS. 2A and 2C. In these embodiments, the data reducer 134 module turns “on” the M-CBD path 132 periodically as opposed to continuously “on” like the EBD path 131. In one embodiment, the limited operation is achieved by forming blocks of data where one block has enough contiguous samples to calculate one SCF, forwarding the one block to the next module, and discarding the remaining blocks. The size of a block (c samples) is $N * L$ where N is the number of bins from the fast Fourier transform (FFT) and L is the number of contiguous FFTs used to compute one SCF. In some embodiments, the data reducer 134 may implement a data reduction algorithm to select the first, last, n th, or a random block from the available N blocks to send to the next module. In one embodiment, un-forwarded blocks may be discarded. Random selection may avoid a situation where a signal is always transmitted during an off-period of the M-CBD path 132 and so is always missed by the M-CBD path 132. Random selection may also hinder third-parties from evading detection by taking advantage of the “on” “off” periods of the M-CBD path 132.

FIG. 2C shows the data reducer 134 includes a data reducer 136 and a stream to vector 138 that is configured to collect the blocks of sampled data and provide the blocks to the α -profile calculator 140.

In some embodiments the computer 120 (FIG. 1) may include additional random-access-memory (RAM) to improve the operational speed of the EBD path 131 and/or the M-CBD path 132. The higher read/write speeds of the RAM enable the computer 120 to keep up with the high rates at which data may be generated by the SDR.

- 14 -

Embodiments of the α -profile calculator 140 may be configured to compute an SCF and an α -profile to be used for signal classification by the signal classifier 142. In one embodiment, the α -profile calculator may be configured to use a time smoothing process, such as Equation 3, below:

$$SCF(f, \alpha) = \frac{1}{L} \sum_{l=1}^L FFT_l[f] \times FFT_l^*[f - \alpha] \quad \text{Equation 3}$$

Where, $FFT_l[f]$ is the l th FFT of a signal at frequency f , and $FFT_l^*[f - \alpha]$ is the complex conjugate of the FFT of the signal at frequency bin, f shifted by α . The α -profile may be calculated according to Equation 4, below:

$$y = \max_{\alpha} SCF(f, \alpha) \quad \text{Equation 4}$$

The α -profile calculator 140 normalizes the computed α -profile by dividing all the entries with a maximum valued entry. The α -profile calculator 140 then passes the SCF and α -profile to multiple signal classifiers 142 for identification of the signal class. The forwarded α -profile is used to predict the signal class part of the SCF is used to estimate the center frequency of a detected signal.

Embodiments of the signal classifier 142 may be configured to classify the detected signals based on the SCF and α -profile provided by the α -profile calculator 140, and provide the classification results to the merger 133 of FIG. 2A. The signal classifier 142 for signal class c , receives an α -profile and uses it as a feature vector to determine if it belongs to class c (or not).

In one embodiment, the signal classifier makes the classification determination by using a previously trained one-class support vector machine (SVM) model. SVM are known to one of ordinary skill in the art of machine learning as a technique to find an optimal hyperplane separating different classes of data. A one-class version of SVM is trained with data from just one "class," learns the boundary of the class from the training data, and predicts if an input feature set belongs to the trained class (or not).

In another embodiment, the signal classifier 142 uses a multi-class classifier model trained with different signals and noise data. In this embodiment, to add support for a n th signal class the model is trained with previous data for 1, 2, ..., $n - 1$ classes and added data for class n .

- 15 -

The signal classifier 142 next computes the center frequency from the SCF if the prediction from the one-class SVM is 1 *i.e.*, the input feature vector is predicted to belong to the signal class c . In one embodiment, the computation is performed using the 0th column of the SCF which contains the magnitudes of the input FFTs averaging over L FFTs. In one
5 embodiment, this computation is carried out as follows: First, the 0th column is divided into blocks. Next, the signal classifier 142 sends two blocks that contain the most amount of the energy (or are above a threshold amount of energy). Next, the signal classifier 142 finds the location of a minimum value in a region bounded by the two blocks. This location is saved as the target center frequency location.

10 In one embodiment, to reduce computation, the signal classifier 142 stores the found center frequencies instead of immediately passing it to the merger 133. After a set time period s , the signal classifier 142 merges the stored results and outputs the merged information to the merger 133. While merging, for each different frequencies, f , the signal classifier 142 determines the number of times f has been detected in period s and uses the count to calculate
15 a classification score for f following a similar calculation as a protocol detector module of the EBD path 131. Like the EBD path 131, the signal classifier 142 may send streams of detected signal data (signal class, frequency, score, count) to the merger 133.

In various embodiments of the classification logic of the M-CBD path 132, the system may have multiple instances of a signal classifiers 142, each programmed/configured to detect
20 a particular signal class and to work in parallel. By way of non-limiting example, one module configured to detect Wi-Fi, one configured to detect ZigBee, etc.

It is specifically contemplated that the data reduction algorithm and the memory may be selected based on factors such as the sample rate of an SDR and the quantity of data it generates. These factors may necessitate different architecture based on different applications.
25 For example, a system may comprise multiple SDR, each SDR scanning different RF bands to detect different classes of wireless signals. The reverse engineering techniques described herein may be optimized for different classes of wireless signals and as such the architecture may be selected to accommodate a class of wireless signal.

FIG. 3 illustrates a classification and capture system (CCS) 300 according to an
30 embodiment of the present disclosure. The CCS 300 includes classification nodes 310, a

- 16 -

coordination server 320 and signal processing nodes 330. Each classification node 310 may include classification logic implemented, in one embodiment, in the manner(s) described with references to FIGS. 1, and 2A to 2C, above. Each classification node 310 may include classification logic optimized for a particular signal class. In some embodiments,
5 classification nodes 310 may be added to the CCS 300 for new or different signal classes, and thus, the CCS 300 is scalable.

Each classification node 310 may include a registration manager (not shown) that is configured to register the classification node 310 with other devices, including the coordination server 320. In one embodiment, registration indicates to the coordination server
10 320 that a classification node 310 is a resource available to the coordination server 320, including to receive job requests/commands from the coordination server 320. As part of registration, a classification node 310 may communicate to the coordination server 320 one or more of: the physical location of a classification node 310, identity of ports to receive updates and commands, RF spectrums with scan range and classification range, and the like. The
15 coordination server 320 may communicate to the classification node 310 initial values and patterns. In the case of energy-based detection, the initial patterns may comprise bandwidth and timing values (active/inactive, short interval spacing, etc.) for known signal classes, and the initial values may include initial cut-off values for the noise. In one embodiment, signal
20 information may be entered manually (e.g., by a user), and in another embodiment signal information may be entered automatically (e.g., using a predefined database or through automated blind signal analysis).

In one embodiment, the classification nodes 310 may include performance monitors configured to monitor for resource (i.e., CPU) usage of a host system as well as detect if processing overload occurs within a classification logic. Processing overload may happen
25 when classification logic is not able to process the stream of samples from the spectrum measurement source fast enough (*e.g.*, we are spending too much computation time on the MLBC path as described below). The performance monitor may be configured to send commands to classification logic, which is configured to make appropriate adjustments responsive to such commands. This enables the CCS 300 to automatically adjust to changes in
30 available resources.

Embodiments of the signal processing nodes 330 may be configured to have different features and functions, including recording a signal and demodulating a signal (or attempting to demodulate a signal). Each signal processing node 330 may comprise specialized hardware and software, relevant to, for example, one or more signal classes. By way of non-limiting example, a signal processing node 330 that specializes in recording signals (*e.g.*, for reverse engineering, demodulation, localization, etc.) may include an SDR for a specific RF band and memory architecture to record the signal. By way of another non-limiting example, a signal processing node 330 that specialize in demodulation of signals and data packet capture may include demodulation software, including software for demodulating specific signal classes.

10 In one embodiment, a specialty SDR such as USRP B210 may be used to demodulate a signal and capture data packets. In another embodiment, BBN's 802.11 demodulation software is used to demodulate the signal in software and capture data packets.

Each signal processing node 330 may include a registration manager (not shown) that is configured to register a signal processing node 330 with a coordination server 320. In one embodiment, the registration manager may send the signal processing node 330 a registration request that includes information about the capabilities of the signal processing node 330. Once registered, the coordination server 320 may direct a signal processing node 330 to perform additional processing on a signal, for example, responsive to capabilities of the signal processing node 330.

20 Each signal processing node 330 may include a controller (not shown) that receives commands and parameters from the coordination server 320, and controls the specific resources of the signal processing node 330 responsive to the commands/parameters.

Among the advantages of embodiments of classification nodes 310 is that they may perform auto-tuning for performance using a monitoring function. This allows the classification nodes 310 to automatically adjust based on available computational resources. This also allows a CCS 300 to be put in an operational state with minimal human intervention by automatically taking the necessary steps to determine its best configuration.

Embodiments of the coordination server 320 may include a classification node controller 321, a node database 322, a signal processing node controller 323, a classification result processor 324, a pattern database 325, and a processing feedback processor 326.

- 18 -

Various embodiments of the coordination server 320 may be configured to act as a centralized coordination point between the classification nodes 310 and the signal processing nodes 330.

Embodiments of the node database 322 may be configured to be used by the other component modules of the coordination server 320 to store and track registered nodes, node availability, node capability, etc.

Embodiments of the classification node controller 321 may be configured to manage registration/un-registration requests from classification nodes 310. Further, it may be configured to send commands to registered classification nodes 310 to monitor a specific frequency range and report the results to the coordination server 320. In some embodiments, the classification node controller 321 may be configured to update the pattern database 325 with new patterns and signals received, for example, from one or more classification nodes 310.

Embodiments of the signal processing node controller 323 may be configured to manage registration/un-registration requests, and work complete requests from the signal processing nodes 330. The signal processing node controller 323 may also be configured to send commands to signal processing nodes 330 responsive to requests/commands received from the classification result processor 324. In some embodiments, upon receipt of a work complete notification, the signal processing node controller 323 may be configured to inform the classification result processor 324 of the node status change.

Embodiments of the classification result processor 324 are configured to receive classification results from the classification nodes 310 and, responsive to a rules engine (not shown), determine whether to perform further processing of a signal and determine which signal processing node to assign a received signal for further processing. By way of non-limiting example, if the classification result processor 324 determines that a signal should be recorded for further investigation at a later time, it may send a request/command to a recording-type of signal processing node 330 to record a specific center frequency at a specific sample rate for a specific amount of time. Similarly, if the classification result processor determines that a demodulation attempt should be made on a signal, it may send a request/command to demodulation-type signal processing node to demodulate (or attempt to demodulate) the signal using demodulation software and to capture packet data from the

- 19 -

signal.

Embodiments of the pattern database 325 may be configured to store, manage and update patterns used classification - e.g., by detectors and SCF classifiers - as well as information about the patterns. In some embodiments, the pattern database 325 may store, manage and update information about types of known and unknown signals. Various
5 embodiments of the pattern database 325 may be updated automatically or manually.

Embodiments of the processing feedback processor 326 may be configured to receive feedback from signal and demodulation attempts by one or more signal processing nodes 330. Further, the processing feedback processing 326 may be configured to update the pattern
10 database 325 with patterns based on feedback received from one or more of the signal processing nodes 330.

Although not shown, the coordination server 320 may include one or more interfaces for users and external devices to access and/or communicate with the coordination server 320, including to retrieve classification results, detection results, load information, set
15 configurations, or issue commands. In one embodiment, a user may interact with an interface by way of a graphical user interface or a command line interface.

Embodiments of the coordination server 320 may include core logic (not shown) that interacts with the other modules to run the system as a whole and is configured to decide about system operation. The core logic may be configured to determine actions to take, such as node
20 assignment and additional processing to be taken (activate a signal processing node, alert a user to suspicious activity, etc.), responsive to information received from the other modules, by way of non-limiting example, classification results, detection results, node availability, and user configuration.

FIG. 4 illustrates an embodiment of a protocol reverse engineering system (PRES) 400,
25 in accordance with the present disclosure. The PRES 400 is configured to receive packet capture (P-CAP) files, for example, from a signal processing node 330 (FIG. 3) that has demodulated a signal and captured packet data. Embodiments of the PRES 400 may be implemented as one or more of the signal processing nodes 330 illustrated in FIG. 3, or as a separate application or tool.

30 Embodiments of the reverse engineering module 410 may be configured to receive the

- 20 -

P-CAP files and infer vocabulary and grammar the application layer protocol associated with the received P-CAP files. The reverse engineering module 410 may store the vocabulary and grammar in the database 420 for access by the traffic generator 430.

Embodiments of the traffic generator 430 may be configured to simulate
5 communication traffic for the wireless signals associated with the reverse engineered P-CAP files including over an Antenna. The traffic generator 430 may be programmed to perform simulations based on the vocabulary and grammar previously inferred. Dynamic vulnerability analysis using integrated fuzzing frameworks (mutation based or generation based), such as Sulley or Peach, may be used to generate optimized and specific fuzzing test cases that may
10 reveal software programming errors which can lead to software security vulnerabilities. By way of non-limiting example, the traffic generator 430 may be configured to generate malformed data packets to attempt to crash or disable a device that is operating over a specific wireless signal, and take over a device, as well as assess the robustness of an implementation.

In one embodiment, PRES 400 is implemented in software, for example using the
15 NetZob tool. In other embodiments, the modules of the PRES 400 may be implemented using tools such as ClusterFuzz and American Fuzzy Lop (AFL).

FIGS. 5A, 5B, and 5C show a classification and detection system 500, in accordance with an embodiment of the disclosure. Embodiments of the classification and detection system 500 may be, by way of non-limiting example, incorporated into a classification and
20 capture system, and a protocol reverse engineering system. The classification and detection system 500 includes different classification logic than the classification and detection system 130. In particular, the classification and detection system 500 may include an energy-based detection classification (EBDC) path 520 and a machine learning-based classification (MLBC) path 530 configured to receive RF measurements 510 from a spectrum measurement source
25 502. The classification results from the EBDC path 520 and the MLBC path 530 are merged by a merger 550. Merged classification results 552 may be provided to a server (e.g., in a classification or reverse engineering system) the noise floor calculator 522 of the EBDC path 520.

In various embodiments, the spectrum measurement source 502 may provide
30 measurement information (e.g., samples of RF signals) for classification. Non-limiting

examples of spectrum measurement sources include, but are not limited to, files and SDRs. The measurements may be processed according to the classification logic and provided, e.g., to a coordination server. The measurements may come from local (e.g., files, attached SDR via USB, etc.) or remote (e.g., sent over a network) sources.

5 FIGS. 5B and 5C show classification logic associated with the EBDC path 520 and MLBC path 530, respectively, in accordance with an embodiment of the disclosure. Some functional modules of the classification logic of the EBDC path 520 shown in FIG. 5B have been simplified for ease of description, but may include, by way of non-limiting example, one or more of the modules of the classification logic shown FIGS. 2B. Here, the EBDC path 520
10 is shown to include a noise floor calculator 522, a parameter estimator 524, and pattern matchers 526.

As shown in FIG. 5C, the MLBC path 530 includes a data reducer 532, feature calculator 534, a noise detector 536, a signal classifier 538, and an MLBC merger 542. The data reducer 532 and feature calculator 534 may be configured similar to the data reducer 134
15 and α -profile calculator 140 described with reference to FIG. 2C. Generally, the data reducer 532 is configured to limit the RF measurements 510 that pass to the rest of the modules of the MLBC path 530 to maintain the real-time capability of a classification process. Also generally, the feature calculator 534 may be configured to compute a pre-defined feature set from RF measurements 510, such as cyclostationarity, higher order cumulants, etc. that is
20 provided to the noise detector 536 and signal classifier 538. The noise detector 536 may use a pre-trained machine learning model to determine if the input is noise or not. The signal classifier 538 runs in parallel with the noise detector 536, and uses one or more signal class models 540, which may be one-class machine learning models, to determine a signal class. For example, one model may be for detecting WiFi, one model may be for detecting ZigBee, etc.
25 A one-class machine learning model is different than conventional models in that it is trained with data from just one class, “learns” the characteristics of the class from the training data, and predicts if an input feature set belongs to the trained class (or not). Although one of ordinary skill in the art may recognize many advantages or benefits to one-class models, one advantage of an one-class model is that, to detect class A signals, it may be trained with class
30 A data, but does not necessarily need to be trained with non-class A data like conventional

- 22 -

machine learning models. For example, to build a model for WiFi, the model does not necessarily need to be trained with noise data and data for other signals, *e.g.*, Bluetooth, ZigBee, etc., - it may only be trained with WiFi data. Advantageously, a representative non-WiFi dataset is not necessarily required, which may be otherwise difficult to create. Non-limiting examples of one-class models that may be used with the signal classifier 538 include one-class SVM, autoencoder neural network, combinations thereof, and the like. MLBC merger 542 may be configured to receive prediction results from both noise detector 536 and signal classifier 538, and to infer a final classification result to be sent to the Merger 550. If noise detector 536 predicts the input as a signal and none of the signal class models 540 recognize it, then an unknown signal type may be considered to be present and the MLBC merger 542 sends a result indicating the logic of the MLBC path 530 did not identify the signal type.

One of ordinary skill in the art will recognize that the embodiments described herein offer many advantages. A particular advantage includes automated detection, classification, capture and protocol reverse engineering of signals, all automatically. Embodiments may be used to automatically analyze RF emanations (signals) and wireless protocols from unknown “black-box” devices, and perform automated wireless protocol reverse engineering of such signals. Other advantages and benefits include an automated solution to: identify and study exposed surfaces of wireless systems, assess the robustness of wireless protocol implementations, reverse engineer wireless command and control protocols used by malicious actors (such as anticipating a botnet’s repertoire of nefarious activity), create protocol specifications for wireless IDS/IPS/Firewall or penetration testing, monitor packets from live wireless connections or captured wireless network traffic that use a wide variety of protocols, perform smart fuzzing of unknown protocols using the discovered protocols vocabulary and grammar and determine the unknown protocols vulnerabilities, reverse engineer standard protocols so actual implementations may be compared to the standard specifications. Other advantages and applications exist.

Memories used to store signals or data as described herein may be arranged inside and outside a computer. For example, in a network, the system memory may include (or be part of) a distributed storage system that provides both storage and file-system, such as network-

- 23 -

attached-storage (NAS), or a distributed storage system that provides only storage, such as a storage-area-network (SAN). In the case of NAS, it may include software capable of file management services, including, without limitation, FreeNAS™, NASLite™, and NexentaStor™. The NAS may contain one or more hard disks, arranged into logical, redundant storage containers or RAID arrays. The NAS may use one or more file-based protocols including, without limitation, Network File System (NFS), Windows NT™ File System (NTFS), File Allocation Table (FAT), Server Message Block/Common Internet File System (SMB/CIFS), or Apple Filing Protocol (AFP).

The information stored on a memory may be stored in a database. The particular architecture of the database may vary according to the specific type of data, mode of access of the data, or intended use of the data stored in the database; including, without limitation, a row-oriented data-store architecture, a column-based database management system, extensible-markup language, a knowledgebase, a frame database, or combinations thereof. A database management system (DBMS) may organize the storage of the data in the database, tailored for the specific requirements of the present system. The DBMS may use any number of query languages to access the database, including, without limitation, structured query language (SQL). In the case of SAN, embodiments of the invention may use any number of protocols to communicate between server and storage, including, without limitation, the SCSI protocol, HyperSCSI protocol, iSCSI protocol, ATA over Ethernet, Fibre channel Protocol, and Fibre Channel over Ethernet.

Various embodiments of the disclosure described with reference to FIGS. 1 through 4, above, may be implemented to facilitate user interaction via a user interface, such as a graphical user interface, invocable by an application program. A user interface may be understood to mean any hardware, software, or combination of hardware and software that allows a user to interact with a computer system. For the purposes of this discussion, a user interface will be understood to include one or more user interface objects. User interface objects may include display regions, user activatable regions, and the like.

A user interface may include input devices such as a keyboard and pointing device, commonly referred to as a mouse, trackball or touch pad. Other input devices may include a microphone, joystick, game pad, satellite dish, scanner, voice recognition device, keyboard,

- 24 -

touch screen, toggle switch, pushbutton, or the like. These and other input devices are often connected to a processing unit through a user input interface that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, a virtual port, game port or a universal serial bus (USB) type interface.

5 A display region is a region of a user interface which displays information to the user. A user activatable region is a region of a user interface, such as a button or a menu, which allows the user to take some action with respect to the user interface.

 A user interface may be invoked by an application program. When an application program invokes a user interface, it is typically for the purpose of interacting with a user. However, it is not necessary that an actual user ever interact with the user interface. It is also not necessary, for the purposes of this invention, that an interaction with the user interface be performed by an actual user. That is to say, it is foreseen that the user interface may have interaction with another program, such as a program created using macro programming language statements that simulate the actions of a user with respect to the user interface.

15 The graphical user interface may provide a user feedback. Based on the feedback, the user may select, input, and/or accept additional parameters (which includes user defined constraints) or to change parameter values. Parameters may also be added or changed automatically.

 A user may enter commands and parameters at a computer terminal via a user interface, including a graphical user interface such as is described herein

Embodiments described in this disclosure may be embodied wholly or partially in one or more computer program products supplied on any one of a variety of computer readable media. The computer program product(s) may be embodied in computer language statements of the types already described herein.

25 One of ordinary skill in the art will appreciate that “media,” “medium,” “computer-readable media,” or “computer readable medium” as used here, may include a diskette, a magnetic tape, a digital tape, a compact disc, an integrated circuit, a ROM, a CD, DVD, Blu-Ray, a cartridge, flash memory, PROM, a RAM, a memory stick or card, or any other non-destructive storage medium useable by computers, including those that are re-writable.

30 Although the enabling software might be “written on” a disc, “embodied in” an

- 25 -

integrated circuit, “carried over” a communications circuit, “stored in” a memory chip, or “loaded in” a cache memory, it will be appreciated that, for the purposes of this application, the software will be referred to simply as being “in” or “on” the computer readable medium. Thus, the terms “in” or “on” are intended to encompass the above mentioned and all equivalent
5 and possible ways in which software can be associated with a computer readable medium.

For the sake of simplicity, therefore, the term “computer program product” is thus used to refer to a computer readable medium, as defined above, which has on it any form of software to enable a computer system to operate according to any embodiment of the invention. Software applications may include software for facilitating interaction with
10 software modules, including user interface and application programming interfaces. Software may also be bundled, especially in a commercial context, to be built, compiled and/or installed on a local computer.

Although each operation illustrated by or in connection with figures and accompanying text recites acts performed in a particular order, embodiments of the present disclosure do not
15 necessarily need to operate in that recited order. One of ordinary skill in the art would recognize many variations, including performing acts in parallel, or in a different order.

Certain embodiments of the present disclosure were described above. It is to be understood that the features of the various embodiments described herein were not mutually
20 exclusive and can exist in various combinations and permutations, even if such combinations or permutations were not made express herein, without departing from the scope of the disclosure. In fact, variations, modifications, and other implementations of what was described herein will occur to those of ordinary skill in the art without departing from the scope of the disclosure. As such, the invention is not to be defined only by the preceding
25 illustrative description, but only by the claims which follow, and legal equivalents thereof.

- 26 -

CLAIMS

We Claim:

- 5 1. A computer-implemented wireless signal classification method, the method comprising:
receiving a first wireless signal classification, the first wireless signal classification based on
 blocks of radio frequency (RF) measurements of a wireless spectrum over a period of
 time;
10 receiving a second wireless signal classification, the second wireless signal classification based
 on part of the blocks of RF measurements;
 weighting the first wireless signal classification and weighting the second wireless signal
 classification; and
 merging the weighted first wireless signal classification and the weighted second wireless
15 signal classification to arrive at a classification result.
2. The method of claim 1 further comprising determining the first wireless signal
classification by:
receiving the blocks of RF measurements;
20 performing energy-based detection on the blocks of RF measurements; and
 classifying at least one wireless signal responsive to the energy-based detection and one or
 more predefined patterns.
3. The method of claim 2, further comprising determining the second wireless
25 signal classification by:
receiving the blocks of RF measurements;
performing feature-based detection on parts of the blocks of RF measurements; and
classifying at least one wireless signal responsive to the feature-based detection and one or
 more signal models.

30

- 27 -

4. The method of claim 3, wherein performing feature-based detection on parts of the blocks of RF measurements comprises:

selecting a block of the blocks of RF measurements; and

processing the selected block to emphasize one or more cyclostationary features.

5

5. The method of claim 4, further comprising determining one or more spectral correlation functions associated with the one or more cyclostationary features.

6. The method of claim 4, further comprising discarding un-selected blocks of RF
10 measurements.

7. The method of claim 1, further comprising:
determining that one or more differences between the first wireless signal classification and
the second wireless signal classification exceed one or more thresholds; and
15 updating a noise-floor associated with an energy-based detection responsive to the one or more
differences.

- 28 -

8. A system, the system comprising:
an energy-based detector configured to analyze an entire set of measurements and generate a first single classification result;
a cyclostationary-based detector configured to analyze less than the entire set of measurements
5 and generate a second signal classification result; and
a classification merger module configured to merge the first signal classification result and the second signal classification result.
9. The system of claim 8, wherein the cyclostationary-based detector comprises:
10 a data reducer configured to receive the entire set of measurements and discard a portion of the entire set of measurements;
a feature calculator configured to identify cyclostationary features of measurements received from the data reducer; and
a signal classifier configured to classify at least one signal responsive to the identified
15 cyclostationary features and at least one signal class model.
10. The system of claim 9, wherein the signal classifier comprises a data store having stored thereon signal class models, and further wherein the signal class models indicative of cyclostationary features of one or more predefined wireless signal protocols.
20
11. The system of claim 9 wherein the at least one signal class model is a one-class support vector machine model.
12. The system of claim 9, wherein the feature calculator is further configured to
25 determine a spectral correlation function for each identified cyclostationary feature.
13. The system of claim 9, wherein the cyclostationary-based detector further comprises a noise detector configured to identify noise among the identified cyclostationary features.
30

- 29 -

14. The system of claim 13, wherein the cyclostationary-based detector further comprises a merger configured to merge results of the noise detector and the signal classifier.

15. The system of claim 14, wherein merging the results of the noise detector and
5 the signal classifier comprises:

receiving a signal indicator or noise indicator from the noise detector;
receiving classification results from the signal classifier; and
generating the second signal classification result.

16. The system of claim 15, wherein the second signal classification result is
10 indicative of: at least one wireless signal protocol, an unknown wireless signal protocol, or noise.

17. The system of claim 15, further comprising:
15 at least one classifier, wherein the at least one classifier comprises the energy-based detector, the feature-based detector, and the merger; and
a coordination server communicatively coupled to the at least one classifier, wherein the coordination server is configured to:
receive the merged classification result and the set of measurements from the at least
20 one classifier; and
provide the set of measurements to a signal processor responsive to the merged classification result and a capability of the signal processor.

18. The system of claim 17, wherein the capability of the signal processor is signal
25 demodulation or signal recording.

19. The system of claim 17, wherein the coordination server is configured to register and unregister classifiers and processors.




FIGURE 1






FIGURE 2A



3/7

FIGURE 2B



CYCLOSTATIONARY-BASED DETECTION PATH 132

FIGURE 2C




FIGURE 3




FIGURE 4




FIGURE 5A




FIGURE 5B




FIGURE 5C

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US18/22639

A. CLASSIFICATION OF SUBJECT MATTER

IPC - H04B 17/20; G01R 23/16; H04W 12/12, 24/08 (2018.01)

CPC - H04L 63/1433; H04B 17/20; G01R 23/16; H04W 12/12, 24/08

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/0263247 A1 (BHATTAD, K. et al.) 18 October 2012; paragraphs [0038], [0043], [0045], [0054]-[0057].	1
---		----
Y		2-6
---		----
A		7
Y	US 2013/0288734 A1 (MODY, A. et al.) 31 October 2013; figure 4E; paragraphs [0029], [0081], [0109].	2-6
Y	LI, D. "Mixed Signal Detection and Parameter Estimation based on Second-Order Cyclostationary Features"; Wight State University. Publication [online]. 1 December 2015 [retrieved 23 April 2018]. Retrieved from the Internet: <URL: https://etd.ohiolink.edu/!etd.send_file?accession=wright1448386709&disposition=inline >; pp 1-39.	4-6
Y	US 2009/0154291 A1 (FERBER, R.) 18 June 2009; paragraph [0012].	6
A	US 6,175,634 B1 (GRAUMANN, D.) 16 January 2001; entire document.	1-7



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

23 April 2018 (23.04.2018)

Date of mailing of the international search report

02 JUL 2018

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 671 272 4300
PCT QSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US18/22639

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fee must be paid.

Group I: Claims 1-7 are directed towards a signal classification method merging weighted signal classifications to arrive at a first classification result

Group II: Claims 8-19 are directed towards a system of generating a first and second signal classification results via a cyclostationary detector and merging the results.

***-continued in supplemental box-**

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-7

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

-***-Continued from Box No. III Observations where unity of invention is lacking -***-

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

The special technical features of Group I include at least a computer-implemented wireless signal classification method, the method comprising: receiving a first wireless signal classification, the first wireless signal classification based on blocks of radio frequency (RF) measurements of a wireless spectrum over a period of time; receiving a second wireless signal classification, the second wireless signal classification based on part of the blocks of RF measurements; weighting the first wireless signal classification and weighting the second wireless signal classification, which are not present in Group II.

The special technical features of Group II include at least a system, the system comprising: an energy-based detector configured to analyze an entire set of measurements and generate a first signal classification result; a cyclostationary-based detector configured to analyze less than the entire set of measurements and generate a second signal classification result; and a classification merger module configured to merge the first signal classification result and the second signal classification result, which are not present in Group I.

The common technical features shared by Groups I-II are merging the first wireless signal classification and the second wireless signal classification to arrive at a classification result.

However, these common features are previously disclosed by US 2012/0263247 A1 to Bhattad, K. et al. (hereinafter "Bhattad"). Bhattad discloses merging the first wireless signal classification and the second wireless signal classification to arrive at a classification result (combining (merging) the classification of reference signal, or RS, tones (signal) of a wireless channel from a first base station and a different classification for the second RS tones from a second base station to derive a channel estimate (classification result); paragraphs [0014], [0054]-[0057]).

Since the common technical features are previously disclosed by the Bhattad reference, these common features are not special and so Groups I-II lack unity.