



- (51) **International Patent Classification:**
H04W 48/02 (2009.01) *H04W 76/00* (2009.01)
- (21) **International Application Number:**
PCT/US2010/036752
- (22) **International Filing Date:**
28 May 2010 (28.05.2010)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
12/473,951 28 May 2009 (28.05.2009) US
- (71) **Applicant (for all designated States except US):** MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** CHANDRA, Ranveer; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). MOSCIBRODA, Thomas; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). MURTY, Rohan N.; c/o Microsoft

Corporation, LCA - International Patents, One Microsoft Way, Redmond, 98052-6399 (US). **BAHL, Paramvir**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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,

[Continued on next page]

- (54) **Title:** CHANNEL DISCOVERY AND DISCONNECTION IN NETWORKS OVER WHITE SPACES AND OTHER PORTIONS OF THE SPECTRUM

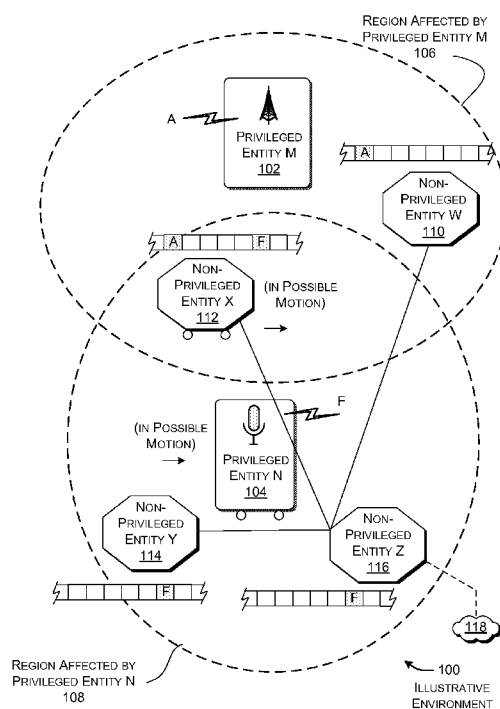


FIG. 1

(57) **Abstract:** Functionality is described for discovering a channel within an environment in which non-privileged entities have subordinate access rights to spectrum compared to privileged entities. The functionality operates by investigating spectrum units within the spectrum for the presence of the channel. In one case, the functionality operates by investigating the spectrum units in linear succession; in another case, the functionality advances in a staggered fashion over the available spectrum. Functionality is also described for handling disconnection by a node from a channel. The functionality allows the node to convey its disconnection status to other communication participants. In one case, various aspects of the functionality are implemented by performing analysis in the time domain.



LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

— *as to the applicant's entitlement to claim the priority of
the earlier application (Rule 4.17(iii))*

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted
a patent (Rule 4.17(ii))*

Published:

— *without international search report and to be republished
upon receipt of that report (Rule 48.2(g))*

CHANNEL DISCOVERY AND DISCONNECTION IN NETWORKS OVER WHITE SPACES AND OTHER PORTIONS OF THE SPECTRUM

BACKGROUND

[0001] An appropriate rule-making authority may occasionally reconsider the rules that govern the use of a particular portion of the wireless communication spectrum. The authority may initiate such reconsideration in view of advances in technology and changes in consumer needs. In general, historical considerations that were relevant at the time of enactment of the rules may no longer apply in full force. Moreover, new considerations may now apply which were not foreseen at the time of enactment of the rules.

[0002] A change in the rules may allow a tremendous expansion in the usefulness of a particular portion of the spectrum. However, such a change may also introduce significant technical challenges. For example, in the United States, the FCC recently modified the rules governing the use of TV spectrum. Historically, the rules reserve that portion of the spectrum for television transmissions and wireless microphones. Such licensed entities are referred to as privileged entities herein. The FCC ruling expands the use of this portion of the spectrum to other, non-licensed, communication devices. However, the FCC ruling states that the new devices are not permitted to interfere with the privileged entities. Such unlicensed entities are referred to as non-privileged entities herein because they have subordinate rights to the available spectrum compared to the privileged entities. By contrast, there are no such constraints governing the use of communication devices in ISM bands (used for Wi-Fi communication). It is a challenging task to develop devices which make efficient use of the unoccupied TV spectrum – which is commonly known as white spaces – under the constraints of the FCC ruling. Similar rule changes and associated challenges may apply to other rule-making jurisdictions and other portions of the spectrum.

[0003] In other cases, the rules governing a particular portion of the wireless spectrum have not changed. Here too, there remains potential for making more efficient use of the spectrum, e.g., by revisiting traditional paradigms that have governed the use of the spectrum.

30

SUMMARY

[0004] According to one illustrative implementation, functionality is described for discovering a channel that can be used to conduct communication among nodes. The functionality operates by investigating the spectrum units within an available spectrum

according to a specified search strategy. The spectrum units may correspond to underlying TV channels (in one illustrative and non-limiting case). The functionality investigates a spectrum unit by determining whether it is at least part of a channel. The functionality performs this task by determining whether a characteristic signal is obtained
5 over the spectrum unit that is associated with the channel, even though the spectrum unit may only partially overlap the channel.

[0005] If a channel is detected, the functionality further examines the characteristic signal obtained over the spectrum unit, e.g., by examining a characteristic pattern in the signal.

The functionality derives the width and center frequency of the channel based on the
10 characteristic signal. Once the width and center frequency of the channel are determined, nodes can communicate using this channel.

[0006] According to one illustrative feature, the functionality is applied in an environment in which the spectrum is shared by privileged entities and non-privileged entities, the non-privileged entities having subordinate rights to the spectrum compared to the privileged
15 entities. The functionality is used by the second entities to discover a channel; for example, a node can use the functionality to discover an access point device with which it may communicate.

[0007] According to one illustrative implementation, a channel used by one node (such as a node associated with an access point device) can have a variable width which is not
20 known to other nodes in advance of the discovery of this channel by the other nodes. In other words, a channel can encompass any number of underlying spectrum units (e.g., TV channels in one example). The center frequency of the channel is also not known to the other nodes. The functionality can discover the channel without examining all permissible channel permutations having different respective center frequencies and widths.

[0008] According to a linear search strategy, the functionality operates by examining each spectrum unit within the available spectrum in linear succession. According to a staggered search strategy, the functionality examines spectrum units within the available spectrum in staggered fashion, skipping over one or more spectrum units in a linear succession of spectrum units. The skipping is performed to investigate the available spectrum for the
25 presence of the channel on a class-width-by-class-width basis, starting with the largest class width first.
30

[0009] According to another illustrative implementation, functionality is described for handling disconnection from a channel. The functionality, when implemented by a first node, operates by detecting a presence of a privileged entity operating on a channel that is

being used by the first node to communicate with a second node (such as an access point device). The first node and the second node are associated with non-privileged entities, where, as stated, the non-privileged entities have subordinate access rights compared to the privileged entity. The functionality then operates by disconnecting the first node from the channel. In one example, disconnection is appropriate because the first node and the second node, being non-privileged entities, are no longer allowed to communicate on the channel in which a privileged entity suddenly begins transmitting (or is suddenly discovered to be transmitting). The functionality then entails sending notification information from the first node to the second node. The notification information informs the second node that the first node has lost its connection to the channel.

[0010] According to one illustrative feature, the functionality can notify the second node via a backup channel, e.g., using a chirping protocol or the like. In addition, the functionality can notify the second node via an alternative channel if a backup channel is determined to be unavailable.

[0011] According to another illustrative implementation, functionality is described for extracting information from a spectrum using time domain analysis. The functionality operates by investigating a plurality of spectrum units in the time domain to identify a characteristic signal. The functionality then derives information based on the characteristic signal, to provide derived information. The functionality then governs communication among at least two nodes based on the derived information.

[0012] According to one illustrative feature, the functionality can use time domain analysis to facilitate the discovery of a channel in the manner summarized above. For example, the functionality can use a characteristic pattern in the signal to derive at least the width of a channel. The functionality can then uses the identified channel to handle communication among nodes.

[0013] According to another illustrative feature, the functionality can use time domain analysis to handle disconnection from the channel in the manner summarized above. For example, the functionality can use time domain analysis to detect telltale chirping activity (or other activity) on a backup or alternative channel, indicating that a node has disconnected from its channel.

[0014] The above approach can be manifested in various types of systems, components, methods, computer readable media, data structures, articles of manufacture, and so on.

[0015] This Summary is provided to introduce a selection of concepts in a simplified form; these concepts are further described below in the Detailed Description. This

Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

- 5 [0016] Fig. 1 shows an illustrative environment in which privileged entities share a portion of a wireless communication spectrum with non-privileged entities.
- [0017] Fig. 2 shows an illustrative derivation of spectrum availability by taking a logical intersection of two spectrum maps.
- [0018] Fig. 3 shows an illustrative use of available spectrum to host channels of varying
 10 width.
- [0019] Fig. 4 shows an illustrative communication device for conducting communication within the environment of Fig. 1.
- [0020] Fig. 5 shows a high-level view of an illustrative communication management module for use in the communication device of Fig. 4.
- 15 [0021] Fig. 6 shows functionality for providing time domain analysis and frequency domain analysis for use in the communication device of Fig. 4.
- [0022] Fig. 7 is a graph which shows a high level view of a signal captured in the time domain; this graph serves as a vehicle for explaining various aspects of the operation of time domain analysis functionality shown in Fig. 6, as applied to the discovery of a
 20 channel.
- [0023] Fig. 8 shows an illustrative communication module for use in the communication device of Fig. 4.
- [0024] Fig. 9 shows an illustrative procedure that provides an overview of the operation of a channel assignment module of Fig. 5.
- 25 [0025] Fig. 10 shows an illustrative procedure that explains how the channel assignment module can be used to generate a suitability assessment for a channel.
- [0026] Fig. 11 shows an illustrative procedure that explains how the channel assignment module can convey a channel selection to a communication participant, and how the appropriateness of the channel selection can be subsequently evaluated.
- 30 [0027] Fig. 12 shows an illustrative procedure that provides an overview of time domain analysis that can be performed by the communication device of Fig. 4.
- [0028] Fig. 13 shows an illustrative procedure that explains how a channel discovery module (of Fig. 5) can be used to discover a channel using a linear search strategy.

[0029] Fig. 14 shows an illustrative procedure that explains how the channel discovery module can be used to discover a channel using a staggered search strategy.

[0030] Fig. 15 shows an illustrative procedure that explains how a disconnection management module (of Fig. 5) can be used to identify a loss of a channel, and to
5 communicate that loss of the channel to other communication participants.

[0031] Fig. 16 shows an illustrative procedure that explains how the disconnection management module can be used to detect notification information which alerts it that a communication participant has lost its channel; Fig. 16 also explains how the disconnection management module can initiate the selection of another channel in
10 response to the notification information.

[0032] Fig. 17 shows illustrative processing functionality that can be used to implement any aspect of the features shown in the foregoing drawings.

[0033] The same numbers are used throughout the disclosure and figures to reference like components and features. Series 100 numbers refer to features originally found in Fig. 1,
15 series 200 numbers refer to features originally found in Fig. 2, series 300 numbers refer to features originally found in Fig. 3, and so on.

DETAILED DESCRIPTION

[0034] This disclosure sets forth functionality for handling various aspects of communication in an environment in which privileged entities share spectrum with non-privileged entities. The non-privileged users have subordinate rights to the spectrum
20 compared to the privileged users, as will be explained in detail below. The functionality can also be used in other types of environments.

[0035] This disclosure is organized as follows. Section A describes illustrative systems for handling communication in the above-described type of environment that includes
25 privileged and non-privileged entities. Section B describes illustrative methods which explain the operation of the system of Section A. Section C describes illustrative processing functionality that can be used to implement any aspect of the features described in Sections A and B.

[0036] As a preliminary matter, some of the figures describe concepts in the context of
30 one or more structural components, variously referred to as functionality, modules, features, elements, etc. The various components shown in the figures can be implemented in any manner, for example, by software, hardware (e.g., discrete logic components, etc.), firmware, and so on, or any combination of these implementations. In one case, the illustrated separation of various components in the figures into distinct units may reflect

the use of corresponding distinct components in an actual implementation. Alternatively, or in addition, any single component illustrated in the figures may be implemented by plural actual components. Alternatively, or in addition, the depiction of any two or more separate components in the figures may reflect different functions performed by a single actual component. Fig. 17, to be discussed in turn, provides additional details regarding one illustrative implementation of the functions shown in the figures.

[0037] Other figures describe the concepts in flowchart form. In this form, certain operations are described as constituting distinct blocks performed in a certain order. Such implementations are illustrative and non-limiting. Certain blocks described herein can be grouped together and performed in a single operation, certain blocks can be broken apart into plural component blocks, and certain blocks can be performed in an order that differs from that which is illustrated herein (including a parallel manner of performing the blocks). The blocks shown in the flowcharts can be implemented by software, hardware (e.g., discrete logic components, etc.), firmware, manual processing, etc., or any combination of these implementations.

[0038] As to terminology, the phrase “configured to” encompasses any way that any kind of functionality can be constructed to perform an identified operation. The functionality can be configured to perform an operation using, for instance, software, hardware (e.g., discrete logic components, etc.), firmware etc., and/or any combination thereof.

[0039] The term “logic” encompasses any functionality for performing a task. For instance, each operation illustrated in the flowcharts corresponds to logic for performing that operation. An operation can be performed using, for instance, software, hardware (e.g., discrete logic components, etc.), firmware, etc., and/or any combination thereof.

A. Illustrative Systems

A.1. Illustrative Environment

[0040] Fig. 1 shows an illustrative environment 100 in which privileged entities and non-privileged entities share a portion of a wireless communication spectrum. As described above, a privileged entity is an entity which has superior rights to the spectrum compared to a non-privileged entity. Fig. 1 shows two representative privileged entities, namely privileged entity M 102 and privileged entity N 104. Privileged entity M 102 has a range of influence corresponding to region 106. Privileged entity N 104 has a range of influence corresponding to region 108. Generally, the range of influence refers to a geographic range of coverage in which entities can receive or are otherwise affected by the signals transmitted by the privileged entities (102, 104), as defined by any jurisdiction-specific

rules and/or other considerations. Fig. 1 shows four representative non-privileged entities, namely non-privileged entity W 110, non-privileged entity X 112, non-privileged entity Y 114, and non-privileged entity Z 116.

[0041] The environment 100 can correspond to any contextual setting governed by any
5 rule-making jurisdiction (or no rule-making jurisdiction). Further, the environment 100 can provide communication in any portion (or portions) of a wireless communication spectrum (e.g., an electromagnetic communication spectrum). To facilitate explanation, the following discussion presents the concrete illustrative example in which the environment 100 corresponds to an expansion of the use of the TV spectrum, as governed
10 in the United States by the FCC. The unoccupied TV spectrum is commonly known as white spaces. In the UHF spectrum, this portion includes, but is not limited to, 180 MHz of available bandwidth from channel 21 (512 MHz) to channel 51 (698 MHz), with the exception of channel 37. Additional TV spectrum is available in the VHF portion of the spectrum. To repeat, this example should be construed as representative and non-limiting.
15 The principles described herein are not limited to any particular portion of a wireless communication spectrum. Nor are the principles limited to any particular country or other rule-making authority.

[0042] In the context of white spaces, the privileged entities correspond to television transmitters and wireless microphones. These entities represent licensed users of the
20 spectrum; that is, as set forth by the FCC, these entities are the original intended users of the spectrum and the TV channels correspond to TV broadcast channels. The non-privileged entities correspond to any other users of the spectrum. As per the FCC ruling, a non-privileged entity cannot communicate on a portion of the spectrum currently being used by a privileged entity, if within a coverage area of that privileged entity. Further, if a
25 privileged entity begins to transmit on a portion of the spectrum being used by a non-privileged entity, it is appropriate that the non-privileged entity immediately vacate that portion of the spectrum. Because of these priority-based rules, the privileged entities are said to have superior rights to the spectrum compared to the non-privileged entities, and the non-privileged entities are said to have subordinate rights compared to the privileged
30 entities. Other rule-making jurisdictions may promulgate different rules governing access to the spectrum. The principles described herein are not limited to any particular rules governing use of the communication spectrum.

[0043] Fig. 1 shows the merely representative case in which non-privileged entity Z 116 serves as an access point device which provides access to a resource of any type. For

example, the resource may represent a network 118 of any type, such as a local area network, a wide area network (such as the Internet), a telecommunication network, etc. The other non-privileged entities (W 110, X 112, and Y 114) communicate with the non-privileged entity Z 116 in order to gain access to the network 118 (or other resource). In a common application, for instance, the non-privileged entities (W 110, X 112, and Y 114) are computer-type devices and the non-privileged entity Z 116 is a router or base station of any type that provides access to the network 118. In this context, the non-privileged entity Z 116 operates in the role of a master, while the non-privileged entities (W 110, X 112, and Y 114) each act in the role of a slave.

10 [0044] However, the principles described herein are not limited to the above-described relationship among non-privileged entities. In another case, any two or more non-privileged entities can communicate with each other using point-to-point communication. In another case, any two or more non-privileged entities can form a local network of any type, such as a mesh network or an ad hoc network. The non-privileged entities can use that local network to transact local transactions amongst themselves and/or to interact with a more encompassing network, such as network 118. Still other paradigms and protocols may apply to the interaction among non-privileged entities.

15 [0045] In the following discussion, the non-privileged entities can most generically be referred to as nodes. Thus, many of the principles can be framed in a general framework of a first node in communication with a second node. As explained above, the two nodes can have a master-slave relationship, a peer-to-peer relationship, and so on. The non-privileged entities can be implemented by communication devices of any kind. Fig. 4, to be discussed in turn, shows one such illustrative communication device.

20 [0046] On the topic of terminology, the term spectrum unit refers to any segment of available spectrum. For example, the TV spectrum is divided into a number of channels, which are referred to as spectrum units herein. The segments of the spectrum are referred to as spectrum units, rather than channels, because the term channel is used in another context, as will be clarified in the course of the ensuing discussion. As one example, the TV spectrum is commonly partitioned into 6 MHz-wide spectrum segments because 6 MHz is the width of an individual TV broadcast transmission. In Europe, the channels are 8 MHz wide.

25 [0047] The use of privileged and non-privileged entities in the same portion of the wireless spectrum contributes to three salient characteristics of the environment 100:

spectrum variation; spectrum fragmentation; and spectrum temporality. These features are explored in greater detail below.

[0048] Spectrum variation means that the number and type of privileged entities that affect a particular region varies from location to location. For example, assume that
5 privileged entity M 102 is using a spectrum unit that is generically labeled as “A,” while privileged entity N 104 is using a spectrum unit that is generically labeled as “F.” For instance, spectrum units A and F may correspond to respective channels in the TV spectrum. All entities within region 106 are affected by privileged entity M 102, transmitting on spectrum unit A. All entities within region 108 are affected by privileged
10 entity N 104, transmitting on spectrum unit F. As such, non-privileged entity W 110 cannot use spectrum unit A due to the presence of privileged entity M 102, but is not constrained by privileged entity N 104. Non-privileged entity Y 114 and non-privileged entity Z 116 cannot use spectrum unit F due to the presence of privileged entity N 104, but are not constrained by privileged entity M 102. And non-privileged entity X 112 cannot
15 use spectrum units A and F due to the presence of both privileged entity M 102 and privileged entity N 104. The following explanation will set forth additional considerations that restrict the spectrum that is availability to any non-privileged entity.

[0049] Spectrum fragmentation means that the presence of privileged entities breaks up the spectrum that is available for transmission to the non-privileged entities. For example,
20 consider the spectrum as “seen by” non-privileged entity X 112. Non-privileged entity X 112 detects that it can transmit on neither spectrum unit A nor spectrum unit F. This means that the available portions of the spectrum are fragmented by the unavailability of spectrum unit A and spectrum unit F. That is, the totality of available spectrum no longer forms a contiguous band, but is partitioned into chunks separated from each other by
25 spectrum units used by privileged entities (in this case, spectrum unit A and spectrum unit F). Here, there are four spectrum units available between spectrum unit A and spectrum unit F (corresponding to generically labeled spectrum units B, C, D, and E, as shown in Fig. 2, to be discussed). But generally, the available spectrum as “seen by” by any non-privileged entity can include any dispersal of available and unavailable spectrum units.

30 [0050] Spectrum temporality means that the presence of privileged entities that affect a particular region may change over time. For example, a privileged entity can become active or inactive at any particular time without warning to the non-privileged entities. For example, a user can switch on or off a wireless microphone at any time without warning. Further, temporality may ensue from the fact that any of the privileged entities and/or non-

privileged entities can be in motion. For example, Fig. 1 indicates that privileged entity N 104 is in possible motion and non-privileged entity X 112 is in possible position.

[0051] With the above characteristics in mind, consider the example in which non-privileged entity W 110 and non-privileged entity Y 114 seek to communicate with non-privileged entity Z 116. In this case, non-privileged entity W 110 is affected by privileged entity M 102 operating on spectrum unit A, while non-privileged entity Y 114 and non-privileged entity Z 116 are affected by privileged entity N 104 operating on spectrum unit F. To ensure overall interference-free interaction, it is desirable that the interaction between non-privileged entities W 110, Y 114, and Z 116 takes place on a channel that is free to all non-privileged entities (W 110, Y 114, and Z 116).

[0052] Fig. 2 illustrates the above concept in graphical form, where the identified reference numbers refer back to Fig. 1. A first column of spectrum units corresponds to a first spectrum map. This spectrum map represents spectrum units that are available to non-privileged entity W 110. A second column of spectrum units corresponds to a second spectrum map. This spectrum map represents spectrum units that are available to non-privileged entity Y 114. A third column of spectrum units corresponds to a third spectrum map. This spectrum map represents spectrum units that are available to non-privileged entity Z 116. In one case, any non-privileged entity can generate its spectrum map by scanning the spectrum for the presence of privileged entities. A non-privileged entity can also generate its spectrum map using other mechanisms. For example, a non-privileged entity can obtain information regarding privileged entities that constrain its operation from a database of any type, or from any other node or nodes, or from any other source (or from any combination of sources). For example, a non-privileged entity can assess its location using any mechanism (e.g., a GPS mechanism), and then consult such a geo-location database to determine privileged entities (and associated spectrum units used by these privileged entities) in its vicinity. That is, such a geo-location database may identify operating constraints on a location-by-location basis.

[0053] A channel assignment module (to be described in detail below) takes the logical intersection (e.g., logical OR) of the three spectrum maps to form an indication of spectrum that is available to non-privileged entities W 110, Y 114, and Z 116. In this case, the channel assignment module determines that the available spectrum includes spectrum units B, C, D, E and G, etc., as represented by the fourth column in Fig. 2. This scenario represents a simplified example; in general, the channel assignment module can

combine spectrum maps with respect to a larger group of communication participants, respectively having more complex spectrum maps.

[0054] The channel assignment module next operates by selecting a channel from within the available spectrum on which to conduct communication. In doing so, the channel
5 assignment module has multiple candidate channels to select from. To convey this point, consider the graphical illustration of Fig. 3. Here, the channel assignment module has determined that spectrum units labeled with “x” are not free to use, e.g., because they are currently being used by privileged entities. The remaining spectrum defines available spectrum for use in conducting communication among non-privileged entities. The
10 channel assignment module can select a channel from any portion of the available spectrum.

[0055] By way of terminology, as explained above, a spectrum unit refers to an underlying portion of the spectrum, such as a pre-existed channel of the TV spectrum, each having a width of 6 MHz. One such representative spectrum unit is spectrum unit 302. Privileged
15 entities communicate over these spectrum units. A channel refers to a portion of the available spectrum that the non-privileged entities use to conduct communication. A channel is associated with a center frequency and a width, as discussed below. In other words, while both spectrum units and channels serve the functional role of communication channels, the following explanation distinguishes between the two for reasons of clarity.

[0056] In one case, all of the channels have the same width. In another case, the channels can have varying channel widths. In other words, the channel widths need not all be the same. For example, in one case, the channel assignment module selects a channel which overlaps only a single spectrum unit. In another case, the channel assignment module
20 selects a channel which overlaps two or more spectrum units.

[0057] For instance, consider the examples shown in Fig. 3, which shows illustrative channels 304 of varying width. A first channel has a width of 20MHz and spans five spectrum units. A second channel has a width of 5 MHz and spans one spectrum unit. A third channel has a width of 10 MHz and spans three spectrum units. These examples are representative and non-limiting. In general, a channel can have any width that is
25 accommodated by a portion of the available spectrum and can overlap any number of spectrum units.

[0058] Each channel is characterized by a width, W , and a center frequency, F . The width defines the frequency span of a channel, e.g., 5 MHz, 10 MHz, 20 MHz, etc. The center frequency defines the center point in the span of frequencies encompassed by the channel.

In one case, the channel assignment module places the channel so that it coincides with a center frequency of an underlying spectrum unit, although other approaches may adopt a different placement rule.

A.2. Illustrative Communication Device

- 5 [0059] Advancing to Fig. 4, this figure shows an overview of a representative communication device 400 that can be used to implement any one of the non-privileged entities shown in Fig. 1. For example, the communication device can be used to implement the non-privileged entity W 110 (which, in one case, operates in a slave mode) or the non-privileged entity Z (which, in one case, operates in a master mode). The
- 10 communication device 400 can correspond to, or can be integrated with, any type of device, or more generally, electrical processing functionality, such as a personal desktop computing device, a laptop computing device, a personal digital assistant (PDA) type computing device, a mobile phone type computing device, a game console device, a set-top box device, a router type device, a server type device, and so on.
- 15 [0060] The communication device 400 includes a wireless interaction module 402 for communicating with other devices via wireless (e.g., radio) communication. The wireless interaction module 402, in turn, can include a scanning module 404 and a main communication module 406. As part of a channel assignment procedure, the scanning module 404 can scan the spectrum (e.g., the TV spectrum) to discover the presence of
- 20 available channels and the quality of the available channels. As part of a channel discovery procedure, the scanning module 404 can scan the appropriate spectrum (e.g., the TV spectrum) for the presence of an assigned operational channel, e.g., a channel associated with an access point device. As part of a disconnection handling procedure, the scanning module 404 can scan the spectrum for the presence of notification information.
- 25 The notification information informs the communication device 400 that another communication device with which it is communicating has moved to another channel.
- [0061] The main communication module 406 enables the communication device 400 to transmit and receive data packets and non-data messages over a channel to one or more other communication devices. The main communication module 406 can be tuned to a
- 30 channel having a specified center frequency and width. Fig. 8, to be discussed in turn, provides additional information regarding the operation of the main communication module 406. Any mechanism and method of wireless communication can be employed by the main communication module 406.

[0062] A number of modules process the information received from or sent to the wireless interaction module 402. For example, an incumbent detection module 408 detects the presence of privileged entities based on a signal provided to it by the scanning module 404. In one implementation, the incumbent detection module 408 performs this operation
5 by detecting telltale patterns in the signal. The telltale patterns are indicative of the transmissions generated by privileged entities, e.g., television transmissions and wireless microphone transmissions. In one case, the incumbent detection module 408 can perform its analysis by operating on a version of the signal which has been transformed from the time domain to the frequency domain. (In addition, or alternatively, a non-privileged
10 entity can determine the presence of privileged entities using other mechanisms, such as by consulting a geo-location database, as described above.)

[0063] A communication management module 410 governs all aspects of the communication performed by the communication device 400. For example, advancing momentarily to Fig. 5, the communication management module 410 can include a channel
15 assignment module 502, a channel discovery module 504, and a disconnection management module 506. The channel assignment module 502 selects an appropriate channel for use in communicating between one or more communication devices. The channel discovery module 504 discovers the presence of a channel that has already been assigned (e.g., a channel associated with an access point device). The disconnection
20 management module 506 performs various tasks upon detecting that a previously assigned channel has been lost by one or more participants of a communication session (e.g., because the channel has been “taken over” by a privileged entity, or for some other reason). Section B (below) provides a detailed description of the illustrative operation of the channel assignment module 502 (with reference to Figs. 9-11), the channel discovery
25 module 504 (with reference to Figs. 13 and 14), and the disconnection management module 506 (with reference to Figs. 15 and 16). The “. . .” notation in Fig. 5 indicates that the communication management module 506 may encompass additional components, not illustrated in this drawing.

[0064] Returning to Fig. 4, an interface module 412 is used to interactively couple any
30 type of processing functionality 414 with the other components of the communication device 400. For example, the interface module 412 may represent, in part, driver functionality used to send information to the wireless interaction module 402 and to receive information from the wireless interaction module 402. The processing functionality 414 may represent, for instance, any type of application module (not shown)

provided by a computing device or the like. For example, the processing functionality 414 may represent a web browser program provided by a laptop computer device that sends data to, and receives data from, the wireless interaction module 402.

[0065] Advancing to Fig. 6, this figure shows a high-level view of processing that can be performed by the communication device 400 based on the signal received from the scanning module 404. A time-to-frequency conversion module 602 converts the signal from a time domain representation to a frequency domain representation, e.g., using a Fast Fourier Transform (FFT) technique or the like. The communication device 400 includes frequency domain (FD) functionality 604 for performing analysis on the signal in the frequency domain, as well as time domain (TD) functionality 606 for performing analysis on the signal in the original time domain. For example, the incumbent detection module 408 can use the FD functionality 604 for performing analysis in the frequency domain. The channel discovery module 504 and the disconnection management module 506 can use the TD functionality 606 for performing analysis in the time domain (as will be described later). The use of the TD functionality 606 for performing analysis in the time domain is, in some respects, more efficient than the FD functionality 604, e.g., because the TD functionality 606 does not require conversion of the signal into the frequency domain and analysis of the signal in the frequency domain.

[0066] Consider, for example, the use of the TD functionality 606 by the channel discovery module 504. In one possible strategy, the channel discovery module 504 can determine the width and center frequency of the channel by investigating each possible permutation of width and center frequency until it discovers a channel. That is, the channel discovery module 504 can use the main communication module 406 to successively tune to each possible center frequency and channel width and then “listen” on this candidate channel for the presence of beacon sent by an access point device or the like. This approach may be characterized as a one-to-one matching strategy, in which the channel discovery module 504 identifies a hypothetical “location” of a channel and then advances to a particular center frequency and channel width to determine whether the channel is actually present at that location.

[0067] The approach described above may involve processing a potentially large number of permutations. In one example, there are a total of thirty 5 MHz channels, twenty eight 10 MHz channels, and twenty six 20 MHz channels. This defines a relatively large number of candidate channels to investigate.

[0068] The channel discovery module 504 can apply a more efficient approach to detect a channel that is agnostic, at the outset, to the particular characteristics of the channel. In this approach, the channel discovery module 504 investigates the available spectrum by investigating spectrum units in successive fashion according to some search strategy
5 (discussed in further detail below). At any given time, the channel discovery module 504 receives a signal from the scanning module 404 representing samples taken from a particular spectrum unit, e.g., a particular TV channel. Assume that the channel discovery module 504 detects a characteristic signal over this particular spectrum unit which reveals the presence of a channel. At this juncture, the channel discovery module 504 can draw
10 the limited conclusion that the spectrum unit is at least part of a channel. One possibility is that the detected channel is entirely contained within the spectrum unit. But the channel can have a variable width which is not known *a priori*. So another possibility is that the channel extends over additional spectrum units.

[0069] In any case, through this process, the channel discovery module 504 has at least
15 detected the presence of a channel without having to “try out” different channel widths and center frequencies in the manner described above. In other words, the channel discovery module 504 can methodically advance through the available spectrum on the basis of a single-width investigation unit, such as an investigation unit that corresponds to the size of an individual TV spectrum unit. And this entire analysis can be performed in
20 the time domain based on the signal provided by the scanning module 404.

[0070] After having detected the presence of the channel, the channel discovery module 504 can derive additional information regarding the characteristics of the detected channel. For example, the channel discovery module 504 can infer at least the width of the channel using time domain analysis. In this approach, the channel discovery module 504 detects a
25 telltale pattern (or patterns) in the signal received by the scanning module 404 over a particular spectrum unit. The channel discovery module 504 can then derive at least the channel width of the channel based on the characteristics of the telltale pattern.

[0071] Consider, for example, the signal of Fig. 7. The scanning module 404 provides raw samples of an RF signal, each sample associated with an (I, Q) pair. Fig. 7 shows a
30 signal formed by plotting the amplitude of the samples $(\sqrt{I^2 + Q^2})$ as a function of time. The signal includes one or more characteristic patterns therein which correlate with the width of the channel. For instance, in one merely illustrative case, the signal may reveal the transmission of a data packet 702 followed by the transmission of an acknowledgment

message 704 having a width 706. An interval 708 separates the end of the data packet 702 from the start of the acknowledgement message 704. In one case, the width 706 of the acknowledgement message 704 and/or the width of the interval 708 may be proportion to an overall width of the channel being used. Based on this correlation, the channel
5 discovery module 504 can infer the width of the channel by measuring the width 706 of the acknowledgment message 704 and/or the width of the interval 708. Section B will provide additional details regarding two search strategies that the channel discovery module 504 can use to discovery a channel using time domain analysis. After detecting the characteristics of the channel in the manner described above, the channel discovery
10 module 504 can tune to the channel using the main communication module 406 and decode its beacon signal.

[0072] The example presented above applies to merely one particular illustrative protocol. In other protocols, other features of the signal may correlate to the characteristics (e.g., width and/or center frequency) of the channel. In general, note that the time domain
15 analysis is efficient because it allows the channel discovery module 504 to determine the characteristics of the channel by investigating a single spectrum unit subsumed by the channel, e.g., without having to probe all possible permutations of channels.

[0073] The channel discovery module 504 can detect the characteristics of the signal in the time domain using a sliding window 710. That is, the channel discovery module 504
20 performs a moving average of the samples encompassed by the sliding window 710, and then performs detection based on the result of the moving average. This operation helps ensure that the channel discovery module 504 accurately detects patterns within the signal – that is, by not being influenced by errant signal events of short duration. The channel discovery module 504 can set the width of the sliding window 710 to be smaller than the
25 width of the interval 708 and the width 706 of the acknowledgment message 704. This width allows the moving average to accurately detect the end of the data packet 702, the start of the acknowledgment message 704, the end of the acknowledgment message 704, etc. For example, the channel discovery module 504 can detect the end of a packet when the moving average drops below a prescribed threshold. The channel discovery module
30 504 can detect the start of a packet when the moving average rises about a prescribed threshold.

[0074] Consider next the use of the TD functionality 606 by the disconnection management module 506. The disconnection management module 506 can examine the signal generated by the scanning module 404 to determine a telltale signal associated with

notification information, such as a chirp-type signal or the like. The chirp-type signal is sent by another communication device to notify the communication device 400 that it has lost its channel (e.g., because a privileged entity has “taken over” the channel). In one implementation, the telltale signal associated with the notification information can also
5 encode the identity of the communication device which sent it. For example, the length of the telltale signal can be used to encode the identity of the sender. By detecting the identity of the sender, the disconnection management module 506 can determine, in the time domain, whether it has a prior communication relationship with the sender. If there is no communicative relationship, the disconnection management module 506 need not
10 spend further time analyzing the signal.

[0075] The communication device 400 can make yet further use of the time analysis performed by the TD functionality 606. The examples set forth above are representative, non-exhaustive, and non-limiting.

[0076] Advancing to Fig. 8, this figure shows one illustrative implementation of the main
15 communication module 406, as introduced in Fig. 4. The main communication module 406 transmits and receives data packets and non-data messages over a channel defined by a width W and a center frequency F . The main communication module 406 includes communication functionality 802 and a frequency translation module 804.

[0077] The communication functionality 802 formats information to be sent by the main
20 communication module 406 and interprets information received by the main communication module 406. In one case, the communication functionality 802 operates in a frequency domain that may not coincide with the frequency domain of signals sent (and received) over the air. To accommodate this scenario, the main communication module 406 can include the frequency translation module 804. On a transmission path, the
25 frequency translation module 804 either upconverts or downconverts the signals received by the communication functionality 802 to an appropriate frequency for transmission over the air. On a reception path, the frequency translation module 804 either upconverts or downconverts the signals received over the air to match the appropriate frequency expectations of the communication functionality 802.

[0078] The main communication module 406 also tunes to a particular channel defined by
30 a particular channel width W and center frequency F . The main communication module 406 can perform this operation by adjusting one or more reference clocks and one or more associated phase locked loops (PLLs), etc. Changes in the channel width may affect various other characteristics of the operation of the main communication module 406.

Hence, the main communication module 406 can also select appropriate parameter values as a function of the channel width, e.g., by a loading an appropriate parameter table or the like as a function of a selected channel width. One representative technology for adjusting channels widths is described in U.S. Application Serial No. 12/163,187, entitled

- 5 “Adapting Channel Width for Improving the Performance of Wireless Networks,” filed on June 27, 2008, naming the inventors of Paramvir Bahl, Ranveer Chandra, Ratul Mahajan, Thomas Moscibroda, and Ramya Raghavendra.

[0079] In one particular representative implementation, aspects of the main communication module 406 can build upon pre-established functionality used for Wi-Fi
10 functionality, operative in the ISM bands. In this case, the frequency translation module 804 can convert between signals appropriate to Wi-Fi communication to signals appropriate to TV spectrum communication. In another implementation, the main communication module 406 can provide appropriate functionality that is originally designed to operate in the TV spectrum (or other target spectrum) without adapting Wi-Fi
15 functionality or other existing communication functionality.

B. Illustrative Processes

[0080] Figs. 9-16 show procedures that explain the operation of the environment 100 and the communication device 400 in flowchart form. Certain principles underlying the operation of the environment 100 and communication device 400 have already been
20 described in Section A; as such, certain operations will be addressed in summary fashion in this section.

B.1. Channel Assignment

[0081] Starting with Fig. 9, this figure shows a procedure 900 for use by the channel assignment module 502 in selecting a channel from a collection of candidate channels
25 within the available spectrum. This procedure 900 can be performed by any first node which is communication with a second node. However, to simplify and facilitate discussion, the procedure 900 is explained in the context of an access point device (such as non-privileged entity Z 116 of Fig. 1) which assumes the role of a master with respect to other communication devices with which it communicates.

30 [0082] To begin with, arrow 902 indicates that the access point device can initiate the procedure 900 in response to various triggering events. One event corresponds to the loss of an existing channel. A channel may be lost due to the appearance of a privileged entity that begins operating on that channel. In one case, the access point device may directly detect the loss of the channel. In another case, another communication device detects the

loss of the channel and sends notification information to the access point device to alert it of the loss of the channel.

[0083] Another event corresponds to the deterioration of the channel, e.g., due to an increase of traffic on that channel, etc. This event may not equate to the outright loss of the channel; rather, it indicates that the channel has become undesirable relative to other channels (which have potentially better quality). In one implementation, the access point device can periodically examine the quality of its selected channel relative to other available channels, and initiate a switch if another channel has better performance.

[0084] Another event corresponds to the start up of the access point device. Still other events may trigger the initiation of the procedure 900.

[0085] In a first general phase 904 of operation, the access point device performs availability analysis. The access point device performs availability analysis to determine what portions of the spectrum that it is entitled to use. In one case, the access point device performs availability analysis to determine those portions of the spectrum that are not currently being used by privileged entities.

[0086] In a second general phase 906 of operation, the access point device performs suitability analysis to select an available channel from a group of candidate channels within the available spectrum (where the available spectrum has been defined by the availability analysis performed in the general phase 904). The access point device performs the suitability analysis by examining the suitability of the candidate channels and selecting the channel deemed most desirable.

[0087] Now examining each operation in turn within the procedure 900, in block 908, the access point device generates local availability information. The local availability information indicates the portions of the spectrum which are available to the access point device (that is, from the “perspective” of the access point device). The access point device can perform this operation by using the scanning module 404 and incumbent detection module 408 to detect the presence of privileged entities within the spectrum. The portions of the spectrum which are not occupied by privileged entities define local availability information from the perspective of the access point device. The access point device can also determine its spectrum map using any alternative mechanism discussed above.

[0088] In block 910, the access point device receives availability information from other communication devices (e.g., nodes) with which it has a communication relationship. Each of these other communication devices can form availability information in the manner described above, e.g., using its own scanning module 404 and incumbent detection

module 408, or through some alternative mechanism discussed above. Each of the other communication devices can forward its availability information to the access point device in any manner, e.g., as a message in a main channel, a message in a backup channel, etc.

However, when the access point device first starts up, it may not have any established relation with other communication devices. In this case, the access point device can select a channel without obtaining availability information from other communication devices.

[0089] In block 912, the access point device identifies spectrum that is available to all participants of a communication session. It performs this operation by taking the logical intersection of the availability information obtained in blocks 908 and 910. For example, the local availability information for each communication device can be expressed as a spectrum map, which, in turn, may take the form of a vector of 1's and 0's. A value of 0 may indicate that a spectrum unit is not available and a value of 1 may indicate that it is available. The access point can form a logical intersection of these spectrum maps to generate a final spectrum map which indicates the spectrum units which are available to all communication participants.

[0090] In block 914, the access point commences the suitability phase of its operation by computing a suitability assessment for each candidate channel within the available spectrum (determined in the availability analysis phase).

[0091] In block 916, the access point selects a channel on which to conduct

communication. In one case, the access point selects a channel having a suitability assessment deemed to be most desirable.

[0092] Fig. 10 shows a procedure 1000 which explains, in one illustrative and non-limiting case, how the channel assignment module 502 can generate a suitability assessment for each candidate channel.

[0093] In block 1002, the channel assignment module 502 can identify a candidate channel to investigate. Recall, based on the discussion of Fig. 3, that a channel can have any width and can span any number of underlying spectrum units (e.g., TV channels).

[0094] In block 1004, assume that the candidate channel under analysis spans two or more spectrum units (e.g., TV channels). In one case, the channel assignment module 502 proceeds by performing a suitability assessment for each individual spectrum unit. In one case, the channel assignment module 502 can compute the suitability assessment for a particular spectrum unit as:

$$\rho_n(c) = \max\left(1 - A_c^n, \frac{1}{B_c^n + 1}\right) \quad (1).$$

[0095] Here, $\rho_n(c)$ represents the suitability assessment of spectrum unit c from the perspective of a node n (e.g., the access point device). A_c^n is a measure of airtime utilization of spectrum unit c from the perspective of a node n . B_c^n is a number of other nodes or entities (e.g., other access point devices) competing with node n with respect to spectrum unit c . That is, this defines how many other entities that node n must contend with. Equation (1) indicates that the suitability assessment $\rho_n(c)$ for a particular spectrum unit is formed by taking the maximum of a first value (based on A_c^n) and a second value (based on B_c^n).

[0096] Less formally stated, $\rho_n(c)$ represents an expected share of spectrum unit c that node n will receive if the spectrum unit c is contained within a channel defined by center frequency F and width W . Equation (1) indicates that, at any instant in time, the probability that a node n will be able to transmit on the spectrum unit c is at least the residual airtime $1 - A_c^n$. However, on each spectrum unit, a node n can also expect to get its “fair share” of the airtime when contending with other access point devices (as reflected by the value $1/(B_c^n + 1)$). The channel assignment module 502 can take the maximum of these two values as an estimate of the probability that a node n will be able to use the spectrum unit c on each transmission opportunity.

[0097] The channel assignment module 502 can generate the airtime utilization measure A_c^n by sampling the utilization of spectrum unit c using its scanning module 404. The channel assignment module 502 can define utilization based on any criterion or plural criteria. For example, the channel assignment module 502 can assess utilization based on the amount of information being transmitted over the spectrum unit in a given time interval. Alternatively, or in addition, the channel assignment module 502 can assess utilization based on the amount of time that a particular communication device is granted access to the spectrum unit, and so on. The channel assignment module 502 can provide an indication of the number of other nodes B_c^n that are using the spectrum unit by detecting signals being transmitted over the spectrum unit by respective access point devices. The channel assignment module 502 can perform aspects of these measurements using the time domain analysis approach described above in Section A.

[0098] The channel assignment module 502 performs the above analysis for each spectrum unit c associated with a channel under investigation.

[0099] In block 1006, the channel assignment module 502 next forms a suitability assessment for the overall channel by combining the suitability assessments of its

constituent spectrum units. In one approach, the channel assignment module 502 can generate the suitability assessment for a channel as:

$$MCham_n(F, W) = \frac{W}{SF} \prod_{c \in (F, W)} \rho_n(c) \quad (2).$$

[00100] Here, $MCham_n(F, W)$ is the suitability assessment for the candidate channel as a whole from the perspective of node n . The candidate channel has center frequency F and
5 width W . SF is a scaling factor, e.g., in one case 5 MHz.

[00101] Less formally stated, since $\rho_n(c)$ represents the expected share of a TV spectrum unit c , the product of these shares across each TV spectrum in the channel (F, W) gives the expected share for the entire channel. This value is scaled by the optimal capacity of the probed channel, e.g., $W = 5$ MHz in one case. That is, the approach uses a 5 MHz channel
10 as a reference point because it fits into one single TV spectrum unit. In general, the $MCham_n(F, W)$ suitability assessment for any candidate channel is related to the current demands placed on the channel, ultimately reflecting the A_c^n and B_c^n measurements associated with its constituent spectrum units.

[00102] Consider two examples. In one case, assume that there is no background
15 interference or other access point devices occupying any portion of the channel (F, W) . Here, $Mcham_n(F, W)$ reduces to the optimal channel capacity. That is, $MCham_n(F, W) = 1$ for $W = 5$ MHz, 2 for $W = 10$ MHz, and 4 for $W = 20$ MHz.

[00103] In another example, consider a channel defined by $(F, W = 20 \text{ MHz})$. Out of the
20 5 TV spectrum units that are spanned by the channel, assume that three have no background interference, one has one competing access point device and airtime utilization of 0.9, and one has one competing access point device with airtime utilization 0.2. $MCham_n(F, 20 \text{ MHz}) = 4 \cdot 0.5 \cdot 0.8 = 1.6$. That is, the suitability assessment predicts a throughput on this channel that is equivalent to roughly 1.6 times an empty 5 MHz channel.

25 [00104] After computing the $MCham_n(F, W)$ suitability assessment for all candidate channels, the channel assignment module 502 can select the channel having the most desirable (e.g., best) suitability assessment.

[00105] Fig. 11 shows a procedure 1100 for conveying a channel selection to other communication participants and evaluating the performance of the channel selection.

30 [00106] In block 1102, the access point device communicates its channel selection to other communication devices with which it has a communicative relationship. The access

point device can convey the channel selection using a beacon signal or other type of message. The communication devices can receive the channel selection on a backup channel, or using a discovery procedure (described below), etc. Upon receipt of the channel selection, the communication devices can use the channel when interacting with the access point device.

[00107] In block 1104, the access point device can evaluate the performance of the selected channel. The access point device can perform this evaluation in different ways. In one case, the channel assessment module 502 of the access point device can periodically scan other available channels and compute the above-described

10 $MCham_n(F, W)$ suitability assessment for these other channels. The access point device can then compare the suitability assessment associated with the other channels with the channel that has been selected and is being used.

[00108] In one case, the access point device can determine the $MCham_n(F, W)$ suitability assessment for the selected channel based solely on its own measurements. In other words, the access point device can compute the $MCham_n(F, W)$ suitability assessment based on measurements made from the “perspective” of the access point device itself. In another case, the access point device can generate the $MCham_n(F, W)$ suitability assessment for the selected channel by taking into consideration measurements made by other communication devices.

20 [00109] For example, assume that, with reference to Fig. 1, non-privileged entity Z 116 is an access point device which communicates with non-privileged entity W 110 and non-privileged entity X 112 over a channel ϑ that has been selected as per the procedures of Figs. 9 and 10. The non-privileged entity Z 116 can compute the efficiency of the selected channel ϑ based solely on its own computation of a $MCham_n(F, W)$ suitability assessment for the selected channel ϑ , in relation to $MCham_n(F, W)$ suitability assessments for other available channels. Alternatively, the non-privileged entity Z can also receive an $MCham_n(F, W)$ suitability assessment generated by non-privileged entity W 110 for the selected channel ϑ and an $MCham_n(F, W)$ suitability assessment generated by non-privileged entity X 112 for the selected channel ϑ . The non-privileged entity Z
 30 can then compute the final $MCham_n(F, W)$ suitability assessment for the selected channel ϑ by averaging the three suitability assessments provided by non-privileged entities Z 116, W 110, and X 112.

[00110] In another case, the access point device can perform the averaging procedure with respect to a plurality of available channels, including the selected channel. In this approach, each communication participant can compute the $MCham_n(F, W)$ suitability assessment for each candidate channel and transmit a set of $MCham_n(F, W)$ suitability assessments to the access point device. The access point device can then generate an average suitability assessment for each available channel and select the most appropriate channel based on average suitability assessments. For that matter, the procedure shown in Fig. 10 (in which an initial channel is selected) can also be modified to use the above-described averaging procedure. That is, the access point device can take the suitability assessments generated by other communication devices into account when selecting an initial channel.

[00111] In yet another implementation, the other communication devices can send their raw measurements (e.g., their A_c^n and B_c^n measurements) to the access point device, and the access point device can compute the $MCham_n(F, W)$ suitability assessments on behalf of the other communication devices (instead of requiring the other communication devices to generate the $MCham_n(F, W)$ suitability assessments). Still other implementations are possible.

[00112] In block 1106, the access point device takes appropriate corrective actions in response to the computations made in block 1104. For example, the access point device can decide to switch back to a previously-used channel (e.g., a channel in use before switching to a new channel). This is assuming that the previously-used channel was abandoned for quality reasons, not because it was “taken over” by a privileged entity.

B.2. Time Domain Analysis

[00113] Fig. 12 is a procedure 1200 that provides an overview of time domain processing that can be performed by the TD functionality 606 of any communication device. In one case, the channel discovery module 504 uses this analysis to detect a channel associated with an access point device. In another case, the disconnection management module 506 uses this analysis to detect notification information which alerts it to the fact that a communication device has lost its channel. Still other applications are possible.

[00114] In block 1202, a communication device investigates spectrum units in the time domain for the occurrence of a characteristic signal. As explained above, one such characteristic signal may be indicative of the presence of a communication channel that overlaps one of the spectrum units. The characteristic signal may further include a

characteristic pattern that corresponds to the duration of certain messages, or the interval between different events, and so on. Another such characteristic signal corresponds to a signal that notifies the communication device of the loss of a channel, such as a chirping-type signal.

5 [00115] In block 1204, Fig. 12 addresses the case in which the characteristic signal corresponds to the presence of a channel. In this case, the channel discovery module 504 has successively detected the presence of a channel by examining an individual spectrum unit, even though the boundaries of the channel may not be defined by the spectrum unit itself (that is, for instance, the channel may extend over one or more additional spectrum
10 units).

[00116] In block 1206, the communication device derives information based on the characteristic signal detected in block 1202. In one case, the communication device derives at least the width of the channel based on the characteristic signal detected in block 1202. In another case, the communication device forms a conclusion that a
15 communication participant has lost its communication channel based on the characteristic pattern detected in block 1202.

[00117] In block 1208, the communication device governs some aspect of communication based on the information that has been derived in block 1204. For example, block 1206 may entail using a channel that has been detected in block 1204 to
20 conduct communication among devices. Or block 1206 may entail invoking a procedure to establish another channel if it has been detected that a previous channel has been lost.

B.3. Channel Discovery

[00118] Fig. 13 shows a procedure 1300 for discovering a channel using a linear search strategy. In one representative and non-limiting case, this procedure may be performed by
25 a communication device that is attempting to establish communication with an access point. In the case of Fig. 13 (and in the case of Fig. 14), the channel discovery module 504 can detect the presence of a channel from the “vantage point” of any individual spectrum unit that is part of the channel. Thus, the channel discovery module 504 can detect the presence of the channel by performing a scan of the available spectrum using a
30 single-sized investigation unit in the time domain, rather than having to investigate different channel permutations using the one-to-one matching approach described above (in Section A).

[00119] In block 1302, the channel discovery module 504 of a communication device sets a spectrum unit index n to the integer value 1.

[00120] In block 1304, the channel discovery module 504 advances to the spectrum unit n in a linear sequence of spectrum units. At the outset, $n = 1$, so the channel discovery module 504 examines the first spectrum unit in the sequence of spectrum units.

[00121] In block 1306, the channel discovery module 504 determines whether the spectrum unit n includes a channel, e.g., by determining whether a characteristic signal is detected over this spectrum unit. Assume at this juncture, that a determination is made that the spectrum unit does not include the channel.

[00122] In block 1308, the channel discovery module 504 determines whether the spectrum unit n is the last spectrum unit available for analysis. If so, and if a channel has not been detected, then the procedure terminates with the conclusion that no channel has been found. If the spectrum unit n is not the last spectrum unit, then, in block 1308, n is incremented by 1 and the above-described procedure is repeated with respect to the next spectrum unit n .

[00123] Assume that the channel discovery module 504 eventually discovers a channel being transmitted over at least part of a spectrum unit. If so, in block 1310, the channel discovery module 504 examines a characteristic pattern (or patterns) associated with the characteristic signal obtained from the spectrum unit. Namely, the channel discovery module 504 computes the width W of the channel based on the characteristic pattern. The width of the channel can be generated in the manner described above, e.g., by deriving the width from the length of an interval between a data packet and an acknowledgement message and/or the duration of the acknowledgment message itself.

[00124] Assume that the channel width W is detected based on analysis performed at a sampled frequency F_s . In general, the center frequency F of the channel occurs within an error bound $\pm E$ of F_s , where E corresponds to $W/2$. In the linear search strategy, the channel discovery module 504 progresses in orderly successive fashion from lower frequencies to higher frequencies. Hence, the center frequency F of the channel is conclusively given as $F_s + E$. Hence, in the linear search strategy, the channel discovery module 504 can compute both W and F from time domain analysis without decoding the signal.

[00125] After detecting the characteristics of the channel in the above-described manner, the channel discovery module 504 can tune to the channel and decode its beacon signal.

[00126] In the linear mode of discovery, the expected number of iterations until a channel is discovered is $N_{su}/2$, where N_{su} is the number of spectrum units to be scanned, e.g., the number of TV channels to be scanned. In the worst case, the number of iterations is N_{su} .

[00127] Fig. 14 shows a procedure 1400 for discovering a channel using a staggered search strategy. Again, in one scenario, this procedure may be performed by a communication device that is attempting to establish communication with an access point device. In this case, instead of investigating each spectrum unit in linear succession, the channel discovery module 504 examines the spectrum units on a width-class-by-width-class basis. That is, the channel discovery module 504 operates by first investigating the spectrum for the presence of a channel that has the largest possible channel width. If a channel of this nature is not found, the channel discovery module 504 investigates the spectrum for the presence a channel having the next-largest channel width. This process continues until the remaining unexamined spectrum units are investigated for the presence of a channel having the smallest possible channel width. Each possible channel width defines a channel width class.

[00128] In block 1402, the channel discovery module 504 sets a channel width index to correspond to a maximum possible channel width, such as 20 MHz (in one merely illustrative example).

[00129] In block 1404, the channel discovery module 504 sets an index m to an integer value of 1. The index m is used to sequence through a group of spectrum units associated with a particular width class.

[00130] In block 1406, the channel discovery module 504 advances to spectrum unit m in the selected width class. This operation may involve skipping over one or more spectrum units. For example, assume that the width class corresponds to a width of 20 MHz. In this case, the channel discovery module 504 advances through the available spectrum by skipping over four spectrum units at a time. The channel discovery module 504 does not investigate the spectrum units that it skips over until another iteration (that is, if that other iteration is eventually performed).

[00131] In block 1408, the channel discovery module 504 determines whether the spectrum unit that it has advanced to has already been investigated in a prior investigation (in a prior iteration). If so, then there is no need to investigate it again, and the procedure 1400 operates by skipping over this spectrum unit and advancing to the next spectrum unit in the width class.

[00132] In block 1410, presuming the spectrum unit has not yet been examined, the channel discovery module 504 determines whether the spectrum unit contains a channel, e.g., by determining whether it contains a characteristic signal associated with a channel. Assume, at this juncture, that the spectrum unit does not contain the channel.

[00133] In block 1412, the channel discovery module 504 determines whether the spectrum unit m is the last spectrum unit in the width class that is currently being examined, e.g., to begin with, the 20 MHz width class. If not, in block 1414, the channel discovery module 504 increases the index m to $m + 1$, upon which it repeats the above-described operation. For example, for the case of the 20 MHz channel class, if the channel discovery module 504 has just examined spectrum unit x , it next examines spectrum unit $x + 5$, because it skips over four spectrum units at each iteration for this width class.

[00134] Alternatively, presume that, in block 1410, the channel discovery module 504 determines that it has advanced to the last spectrum unit in a width class. If so, in block 1416, the channel discovery module 504 next asks whether the channel unit that it has reached is the overall last spectrum unit to be examined. If so, then the channel discovery module 504 terminates the procedure 1400, reaching the conclusion that a channel has not been found.

[00135] Alternatively, if the spectrum unit is not the overall last spectrum unit, then, in block 1418, the channel discovery module 504 decrements the class width index by 1. This prompts the channel discovery module 504 to repeat the above-described procedure for the next smallest channel width, such as 10 MHz. In the case of 10 MHz, the channel discovery module 504 sequences through the spectrum units by skipping over two spectrum units at a time. As stated above, if the channel discovery module 504 determines that it has already examined a spectrum unit in a prior iteration (e.g., for a prior class width), there is no need to reexamine it again. The channel discovery module 504 moves on to the next spectrum unit for that class width.

[00136] The above procedure advances by scanning through the spectrum units in succession, on a width-class-by-width-class basis. If a channel is not discovered, the channel discovery module 504 eventually reaches the smallest possible channel width, e.g., 5 MHz in one example.

[00137] Block 1420 is eventually invoked when the channel discovery module 504 eventually detects a channel that overlaps one of the spectrum units. The channel discovery module 504 can then infer the channel width based on time domain analysis of the characteristic pattern (or patterns) of the detected characteristic signal, as described above with respect to Fig. 7.

[00138] However, in the case of the staggered search strategy, the channel discovery module 504 has moved through the spectrum units in a staggered fashion, rather than a linear fashion as in the case of Fig. 13. This means that the channel discovery module 504

cannot conclusively determine the center frequency of the channel in the manner described above with respect to the linear mode of operation. Instead, the channel discovery module 504 can examine the spectrum in the vicinity of the detected channel to decode the beacon being transmitted over the channel. From this information, the channel discovery module 504 obtains the center frequency of the channel and other appropriate information. In other words, the channel discovery module 504 can narrow down the presence of the channel within a small range. The channel discovery module 504 can successively advance through different center frequencies within this range to discover the actual center frequency of the channel.

10 [00139] The expected discovery time of the staggered version of the discovery procedure can be shown to be $(1/N_W)(N_{su} + 2^{N_W-1} + (N_W - 1)/2)$. The expected number of iterations is $(N_{su} + 4 + 1)/4$. Here, N_{su} is the number of spectrum units to scan, and N_W is the number of width class options to process (e.g., $N_W = 3$ in the example shown in Fig. 3). Both algorithms have a worst-case discovery time of N_W .

15 B.4. Disconnection Management

[00140] Advancing to Fig. 15, this figure shows a procedure 1500 performed by the disconnection management module 506 to convey the loss of a channel. For example, a communication device may use this procedure 1500 to alert its access point device that it has lost its channel. One reason a communication device may lose its channel is because it is suddenly discovered that a privileged entity is transmitting using the channel (on any spectrum unit that overlaps the channel). This may be due to the fact that the privileged entity has just started transmitting. Or the privileged entity may have just moved into the vicinity of the communication device, and so on. In any case, the privileged entity “takes over” the channel.

25 [00141] In block 1502, the communication device detects that a privileged entity is now transmitting on the channel that the communication device is using to communicate with other communication participants. The communication device can make this determination using its scanning module 404 in conjunction with the incumbent detection module 408, or using some other mechanism. The communication device can discriminate a transmission that originates from a privileged entity from other communication based on a telltale signature associated with the transmission.

30 [00142] In block 1504, the communication device immediately disconnects from the channel it is using upon detecting a privileged entity.

[00143] In block 1506, the communication module sends notification information to its other communication participants, such as its access point device, to alert those devices of the loss of the channel. In one case, the communication module can use the main channel to send this message to the other participants, e.g., immediately before it abandons this
5 main channel. In another case, the communication module can use a backup channel to send this message to the other participants. The access point device can, in advance, inform the communication device of the existence of this channel as part of its beacon message, or as part of another message. There is a risk that the backup channel is itself “taken over” by a privileged device and is therefore unavailable for use. In this case, the
10 communication device can transmit the notification message over any other alternative channel.

[00144] In one case, the communication device can transmit the notification information using a chirping-type signal that the recipient device can readily interpret as a message which indicates loss of a channel. In one case, the notification information can also
15 encode the identity of the communication device which is sending the notification information. For example, the communication device can encode the identity of the sender by modulating the length of the signal used to transmit the notification information.

[00145] Fig. 16 describes a procedure 1600 for detecting notification information sent by another communication device, and for taking action based on the notification
20 information. For example, although not so limited, an access point device can perform the procedure 1600 upon receiving notification information from one of its communication devices.

[00146] In block 1602, the access point device detects notification information from a communication device. In one case, the access point device performs this operation by
25 periodically scanning its backup channel, and also alternative channels. It can perform this task using the scanning module 404 in conjunction with the type of time domain analysis described above. In this manner, the detection operation does not unduly interfere with other communication tasks performed by the access point device. Once the access point device detects the notification information (e.g., a chirping-type signal), it can tune
30 the main communication module 406 to the backup channel and decode the contents of the chirp signal.

[00147] In block 1604, the access point device invokes the channel assignment procedure described above to assign another channel for use in conducting communication with its

communication devices. The other communication devices can communicate their spectrum maps and the like to the access point device via the backup channel.

[00148] In block 1606, the access point device can convey the new channel selection to its communication devices, e.g., within a beacon signal transmitted by the access point
5 device or using some other message transmitted by the access point device.

C. Representative Processing Functionality

[00149] Fig. 17 sets forth illustrative electrical processing functionality 1700 that can be used to implement any aspect of the functions described above. With reference to Fig. 4, for instance, the type of processing functionality 1700 shown in Fig. 17 can be used to
10 implement any aspect of the communication device 400. In one case, the processing functionality 1700 may correspond to any type of computing device that includes one or more processing devices.

[00150] The processing functionality 1700 can include volatile and non-volatile memory, such as RAM 1702 and ROM 1704, as well as one or more processing devices 1706. The
15 processing functionality 1700 also optionally includes various media devices 1708, such as a hard disk module, an optical disk module, and so forth. The processing functionality 1700 can perform various operations identified above when the processing device(s) 1706 executes instructions that are maintained by memory (e.g., RAM 1702, ROM 1704, or elsewhere). More generally, instructions and other information can be stored on any
20 computer readable medium 1710, including, but not limited to, static memory storage devices, magnetic storage devices, optical storage devices, and so on. The term computer readable medium also encompasses plural storage devices. The term computer readable medium also encompasses signals transmitted from a first location to a second location, e.g., via wire, cable, wireless transmission, etc.

[00151] The processing functionality 1700 also includes an input/output module 1712 for
25 receiving various inputs from a user (via input modules 1714), and for providing various outputs to the user (via output modules). One particular output mechanism may include a presentation module 1716 and an associated graphical user interface (GUI) 1718. The processing functionality 1700 can also include one or more network interfaces 1720 for
30 exchanging data with other devices via one or more communication conduits 1722. One or more communication buses 1724 communicatively couple the above-described components together.

[00152] Although the functionality has been described in the illustrative context of any environment that is shared between privileged and non-privileged entities, aspects of the functionality can also be applied in other environments that do not have this characteristic.

[00153] More generally, although the subject matter has been described in language
5 specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

CLAIMS

1. A method (1200), using electronic processing functionality, for extracting information from spectrum, comprising:
 - investigating (1202) a plurality of spectrum units in a time domain to identify a characteristic signal;
 - deriving (1206) information based on the characteristic signal to provide derived information; and
 - governing (1208) a communication among at least two nodes based on the derived information.
2. The method of claim 1, wherein the spectrum is used by both privileged entities and non-privileged entities, the non-privileged entities having subordinate rights to the spectrum compared to the privileged entities, said at least two nodes corresponding to two respective non-privileged entities.
3. The method of claim 1, wherein the derived information corresponds to a width of a channel that is being used to conduct communication, and wherein said governing comprises using the channel to conduct communication between said at least two nodes.
4. The method of claim 3, wherein the derived information also corresponds to a center frequency of the channel.
5. The method of claim 1, wherein the derived information corresponds to an indication that a node has lost a channel on which it was communicating, and wherein said governing comprising initiating a procedure to reestablish communication with the node.
6. The method of claim 1, wherein said deriving comprises detecting a width of a packet within the characteristic signal and providing the derived information based on the width of the packet.
7. The method of claim 1, wherein said deriving comprises detecting an interval between events in the characteristic signal and providing the derived information based on the interval between events.
8. The method of claim 1, wherein said deriving comprises detecting information in the characteristic signal that is indicative of an identity of an entity which has sent the characteristic signal, and providing the derived information based on the information that is detected.
9. A computer readable medium (1710) for storing computer readable instructions, the computer readable instructions providing a discovery module (504) when executed by one or more processing devices (1706), the computer readable instructions comprising:

logic (1304, 1406) configured to advance to a spectrum unit within the available spectrum;

logic (1306, 1410) configured to determine whether the spectrum unit is at least part of a channel, the channel being said to be discovered if it is determined that the spectrum unit is at least part of the channel; and

said logic (1304, 1406) configured to advance and said logic (1306, 1410) configured to determine repeating their respective advancing and determining operations until the channel is discovered.

10. The computer readable medium of claim 9, further comprising logic configured to derive, if a channel is discovered, at least a width of the channel based on a characteristic signal detected over the spectrum unit.

11. The computer readable medium of claim 9, wherein said logic configured to advance and said logic configured to determine operate by examining each spectrum unit within the available spectrum in linear succession.

12. The computer readable medium of claim 9, wherein said logic configured to advance and logic configured to determine operate by examining spectrum units within the available spectrum in staggered fashion, skipping over one or more spectrum units in a linear succession of spectrum units, wherein said skipping is performed to investigate the available spectrum for the presence of the channel on a class-width-by-class-width basis.

13. A disconnection management module (506), comprising:

logic (1502) configured to detect a presence of a privileged entity operating on a channel that is being used by a first node to communicate with a second node, the first node and the second node being associated with non-privileged entities, the non-privileged entities having subordinate access rights compared to the privileged entity;

logic (1504) configured to disconnect the first node from the channel; and

logic (1506) configured to send notification information to a second node that indicates that the first node has disconnected from the channel.

14. The disconnection management module of claim 13, wherein said logic configured to send notification information is configured to use a backup channel to send the notification information to the second node.

15. The disconnection management module of claim 13, wherein said logic configured to send notification information is configured to use an alternative channel to send the notification information to the second node if a backup channel is not available.

1/12

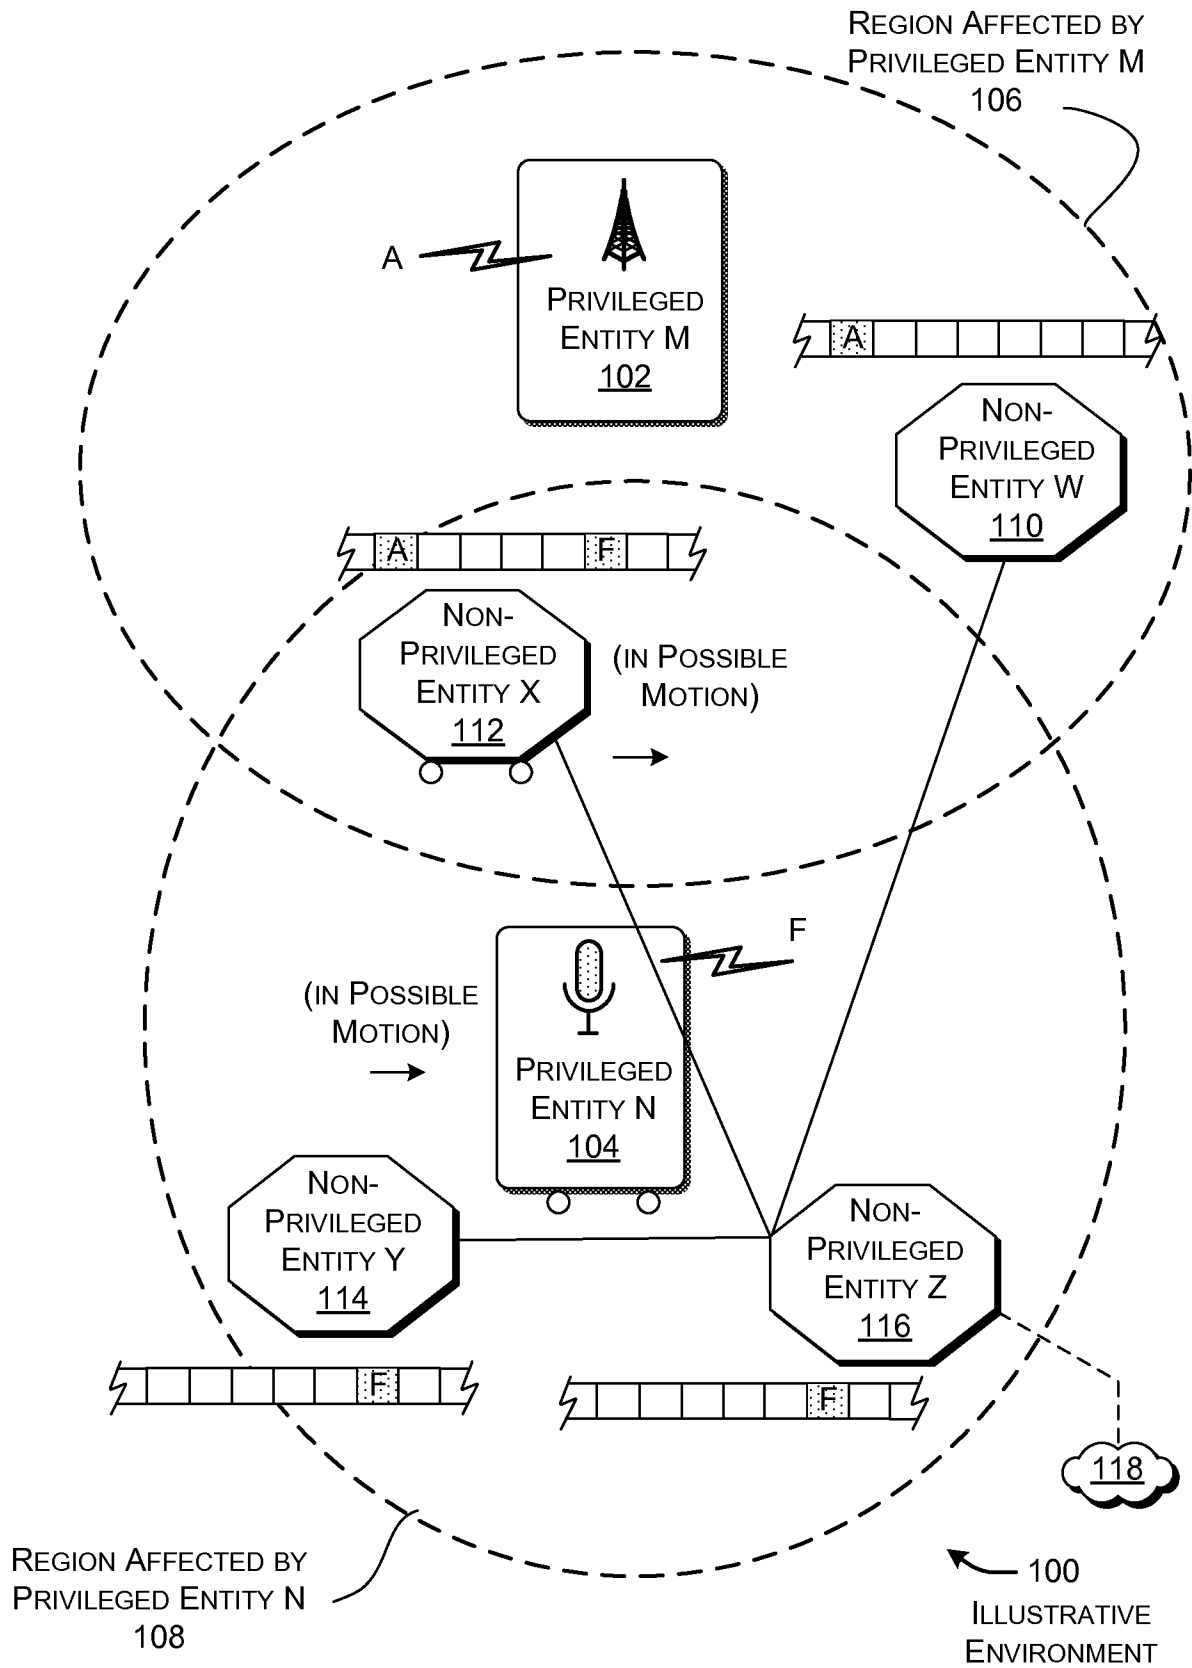
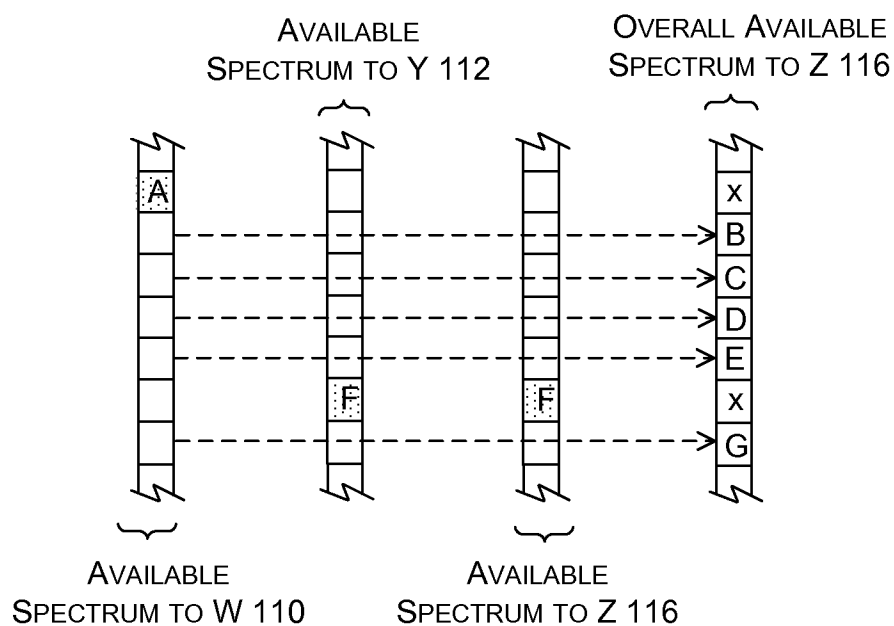

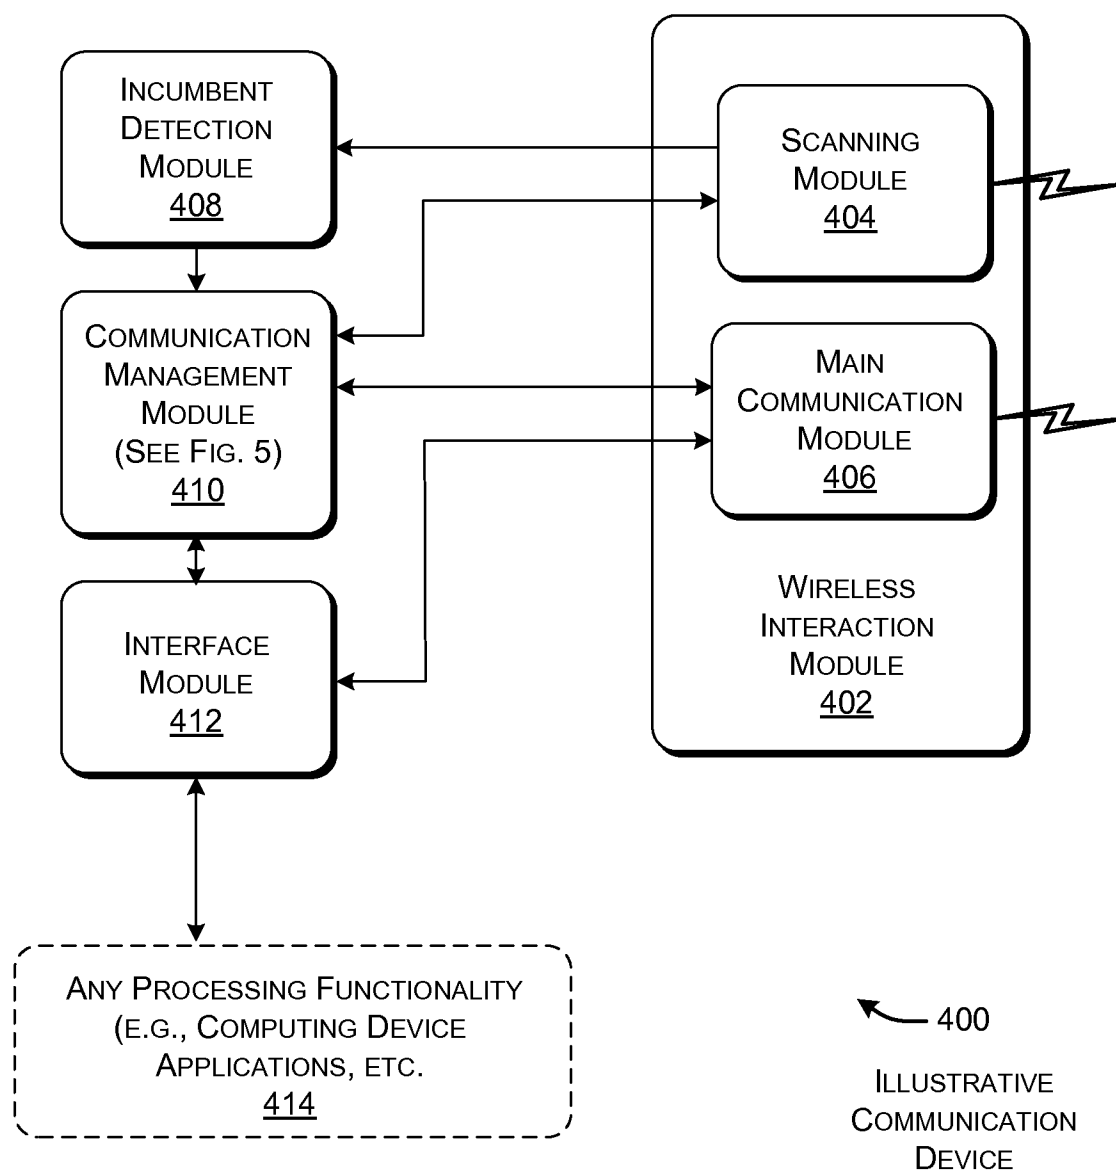


FIG. 1

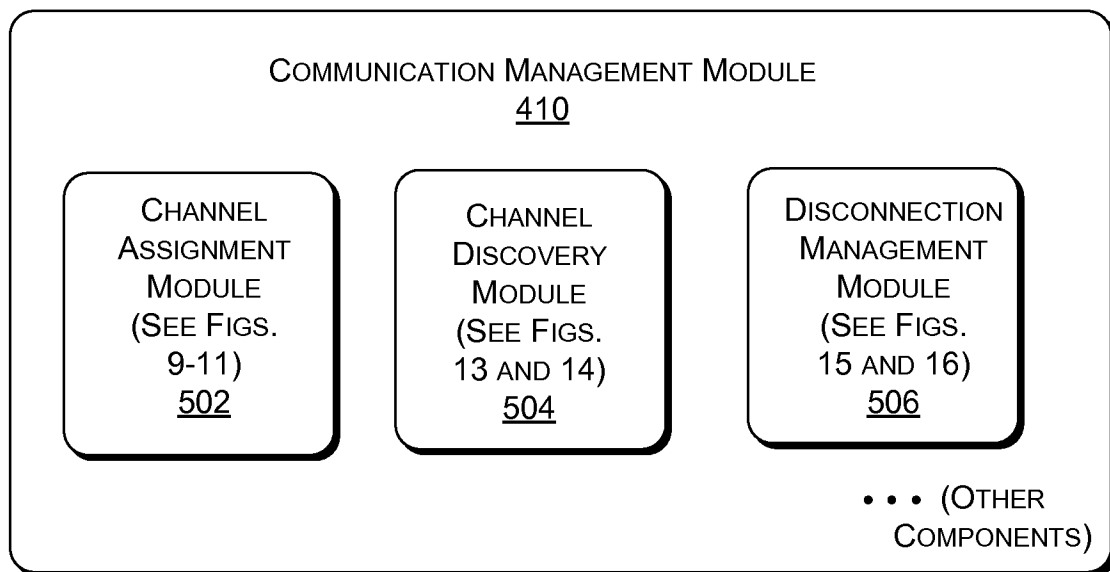
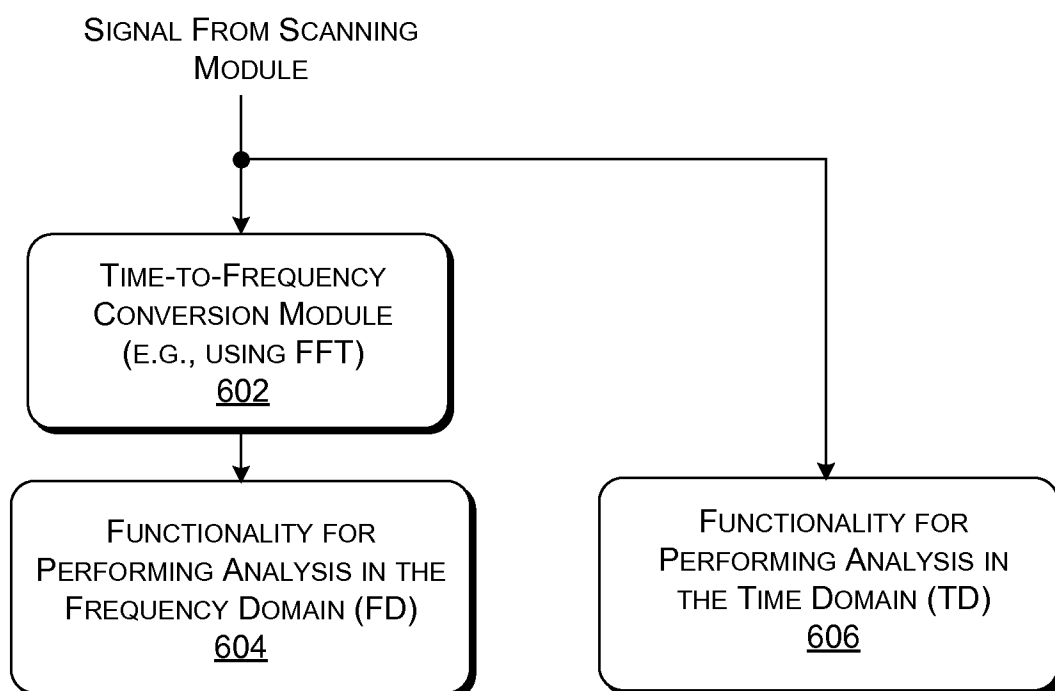
2/12

**FIG. 2****FIG. 3**


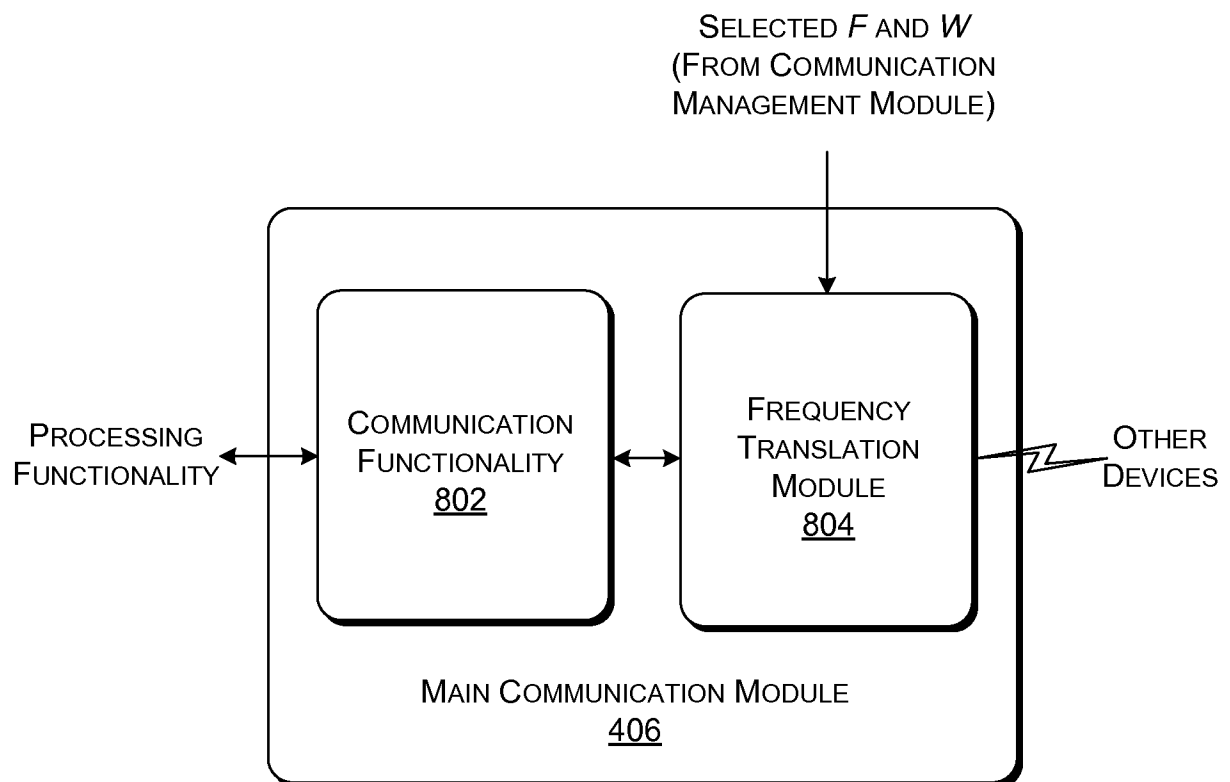
3/12

**FIG. 4**

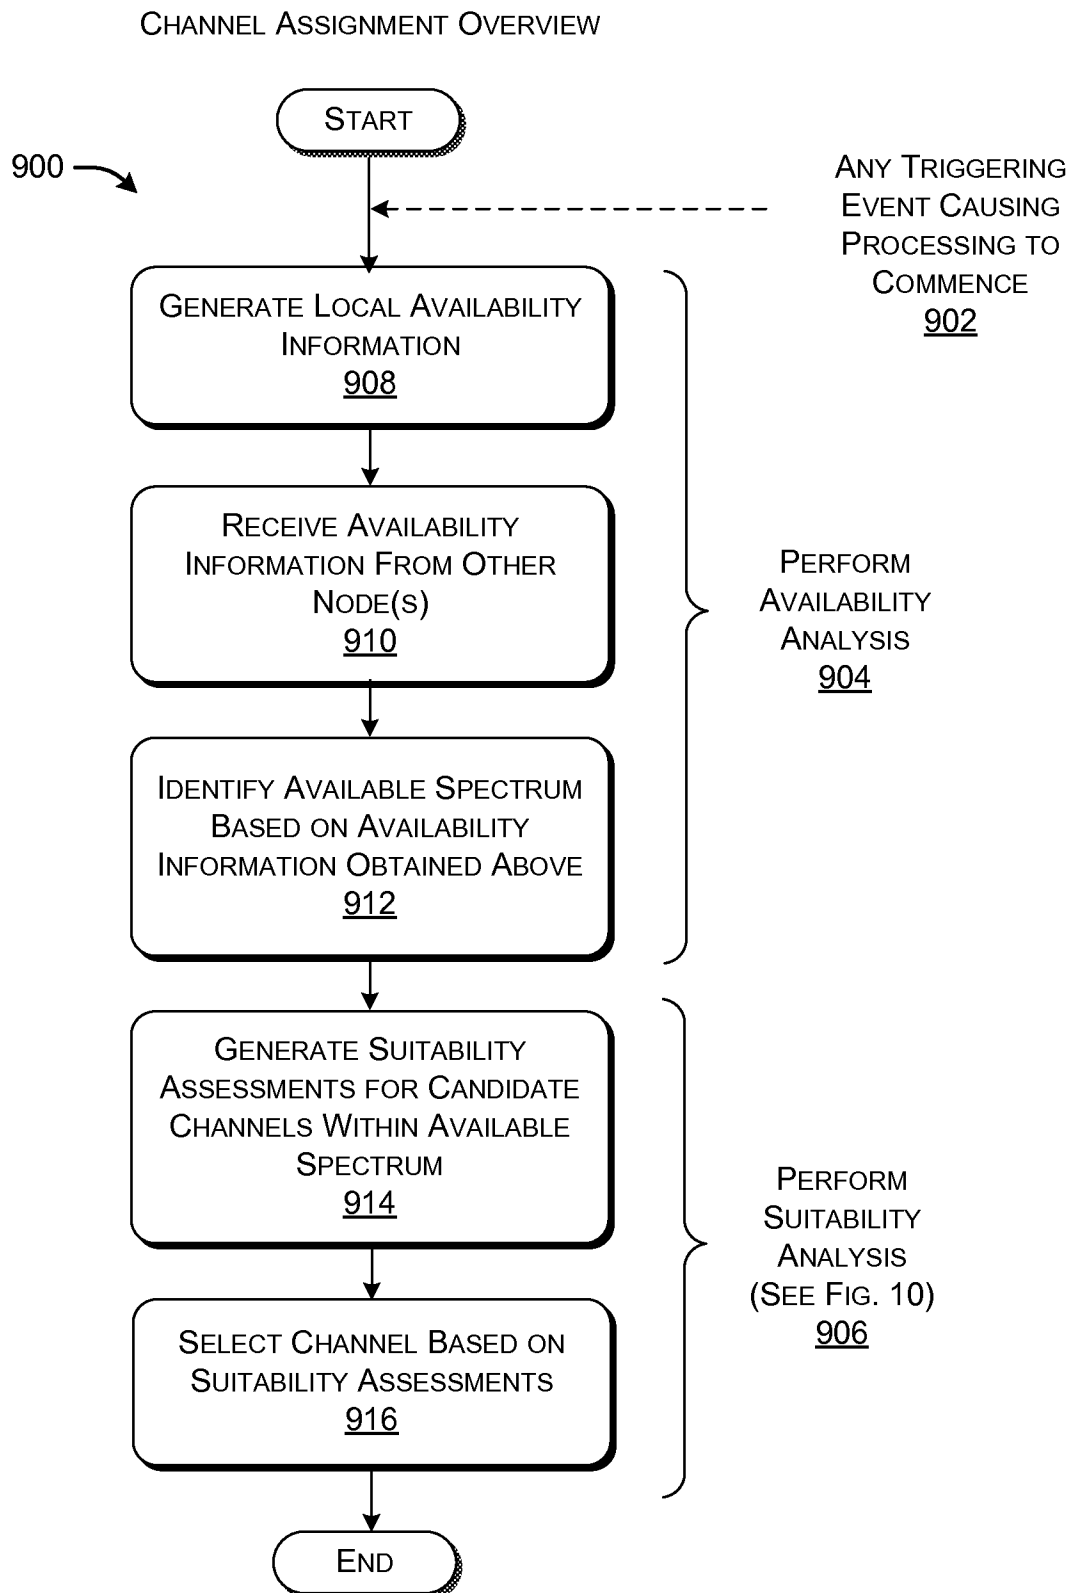
4/12

**FIG. 5****FIG. 6**

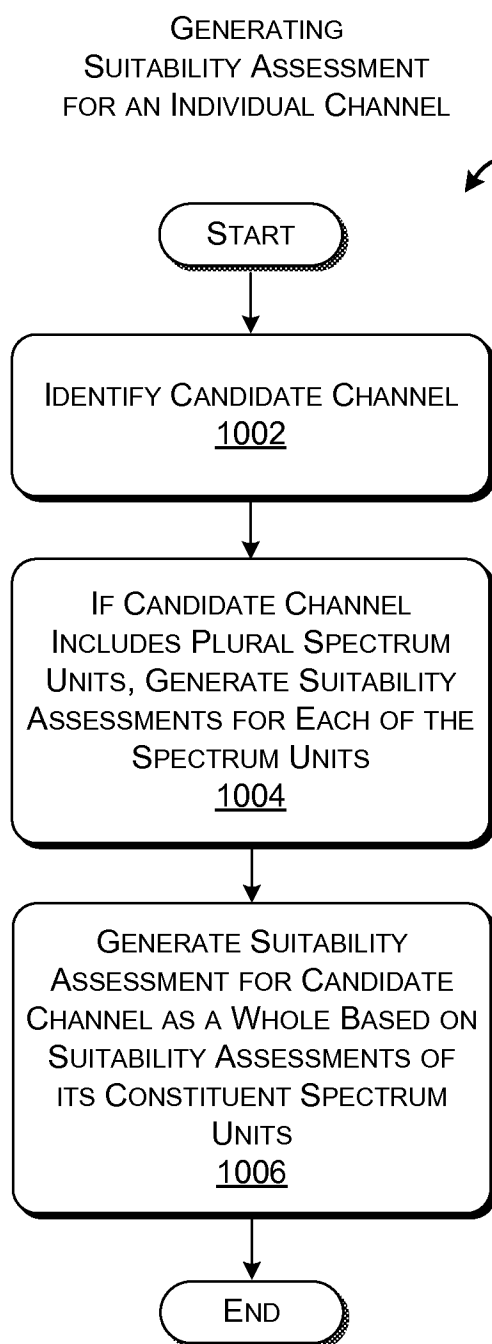
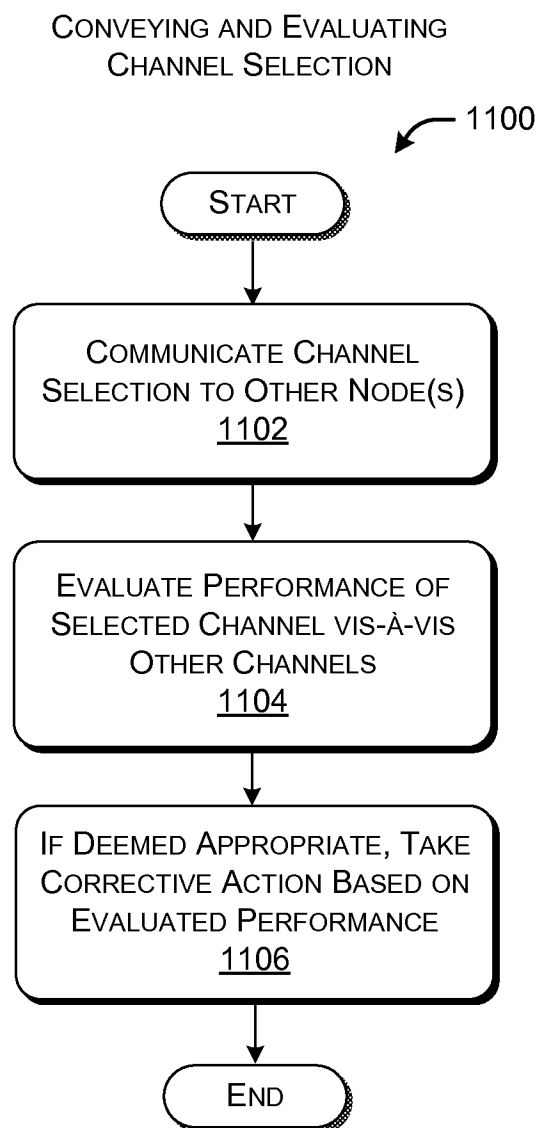
5/12

**FIG. 7****FIG. 8**

6/12

**FIG. 9**

7/12

**FIG. 10****FIG. 11**

8/12

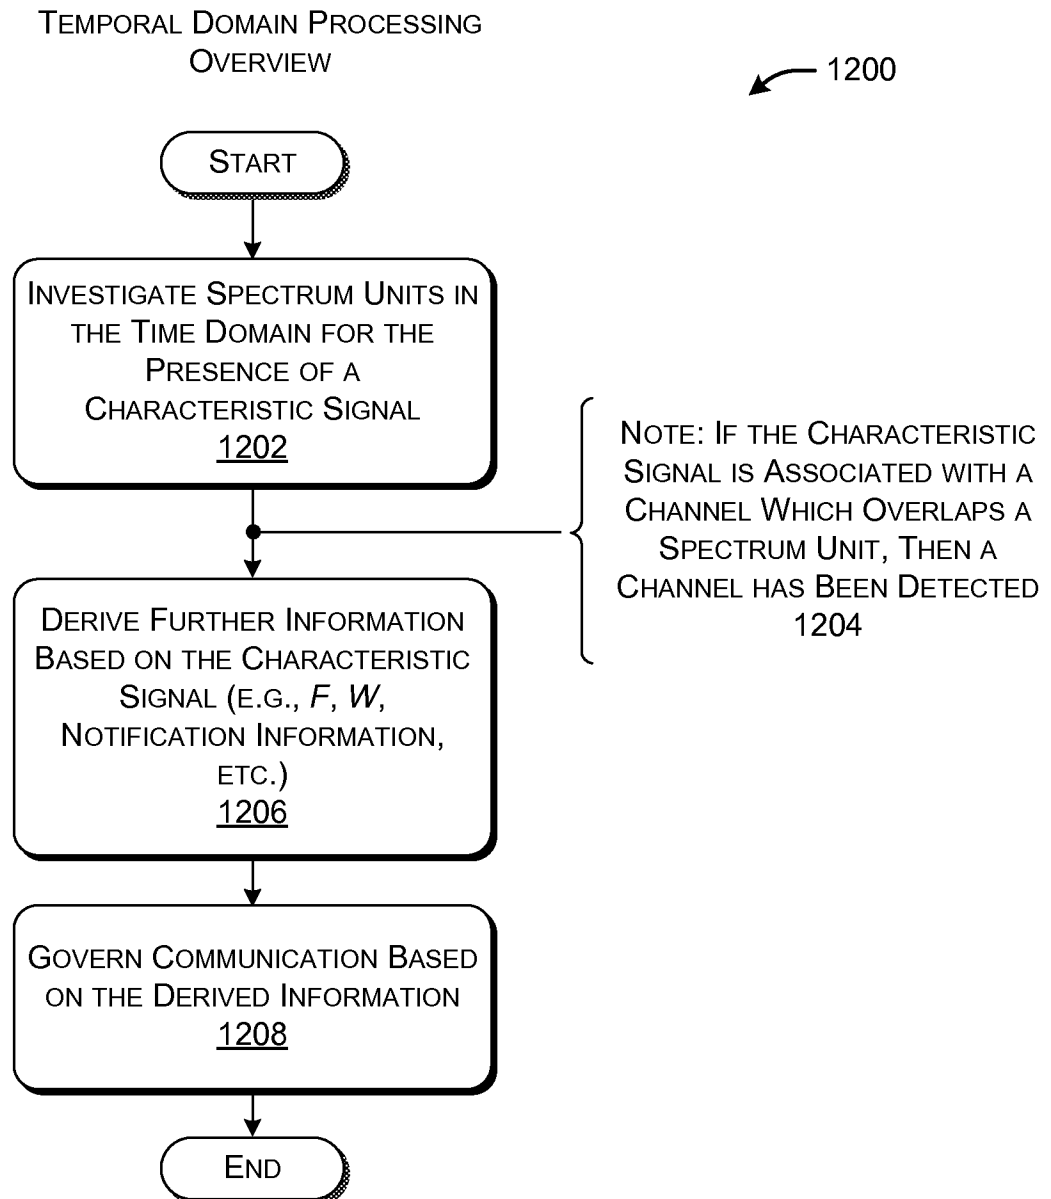


FIG. 12

9/12

LINEAR CHANNEL DISCOVERY

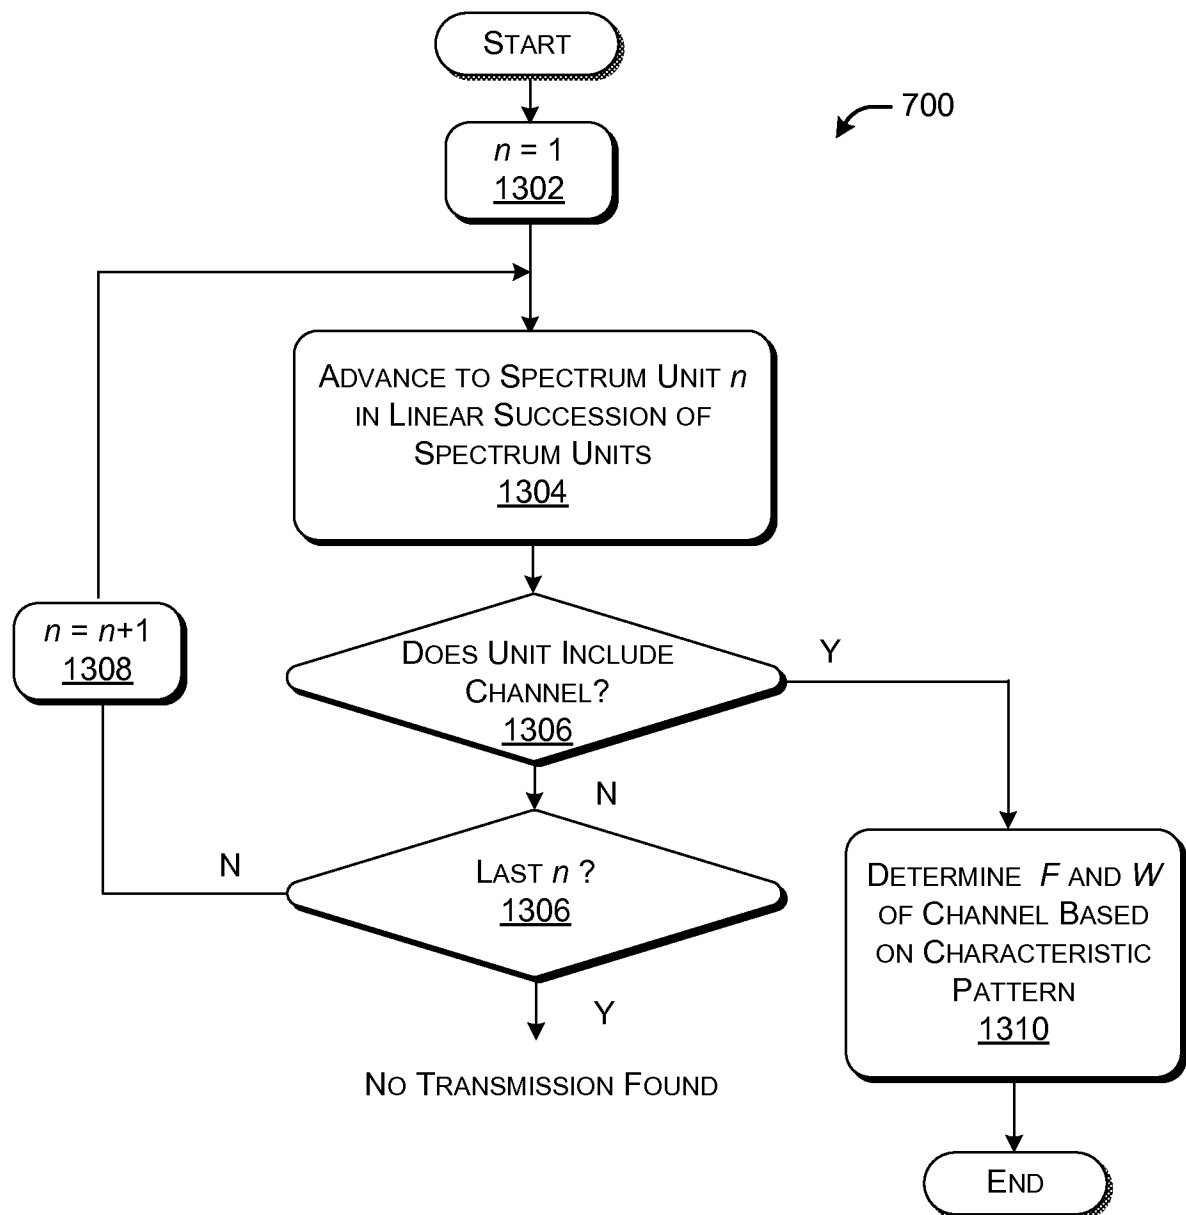


FIG. 13

10/12

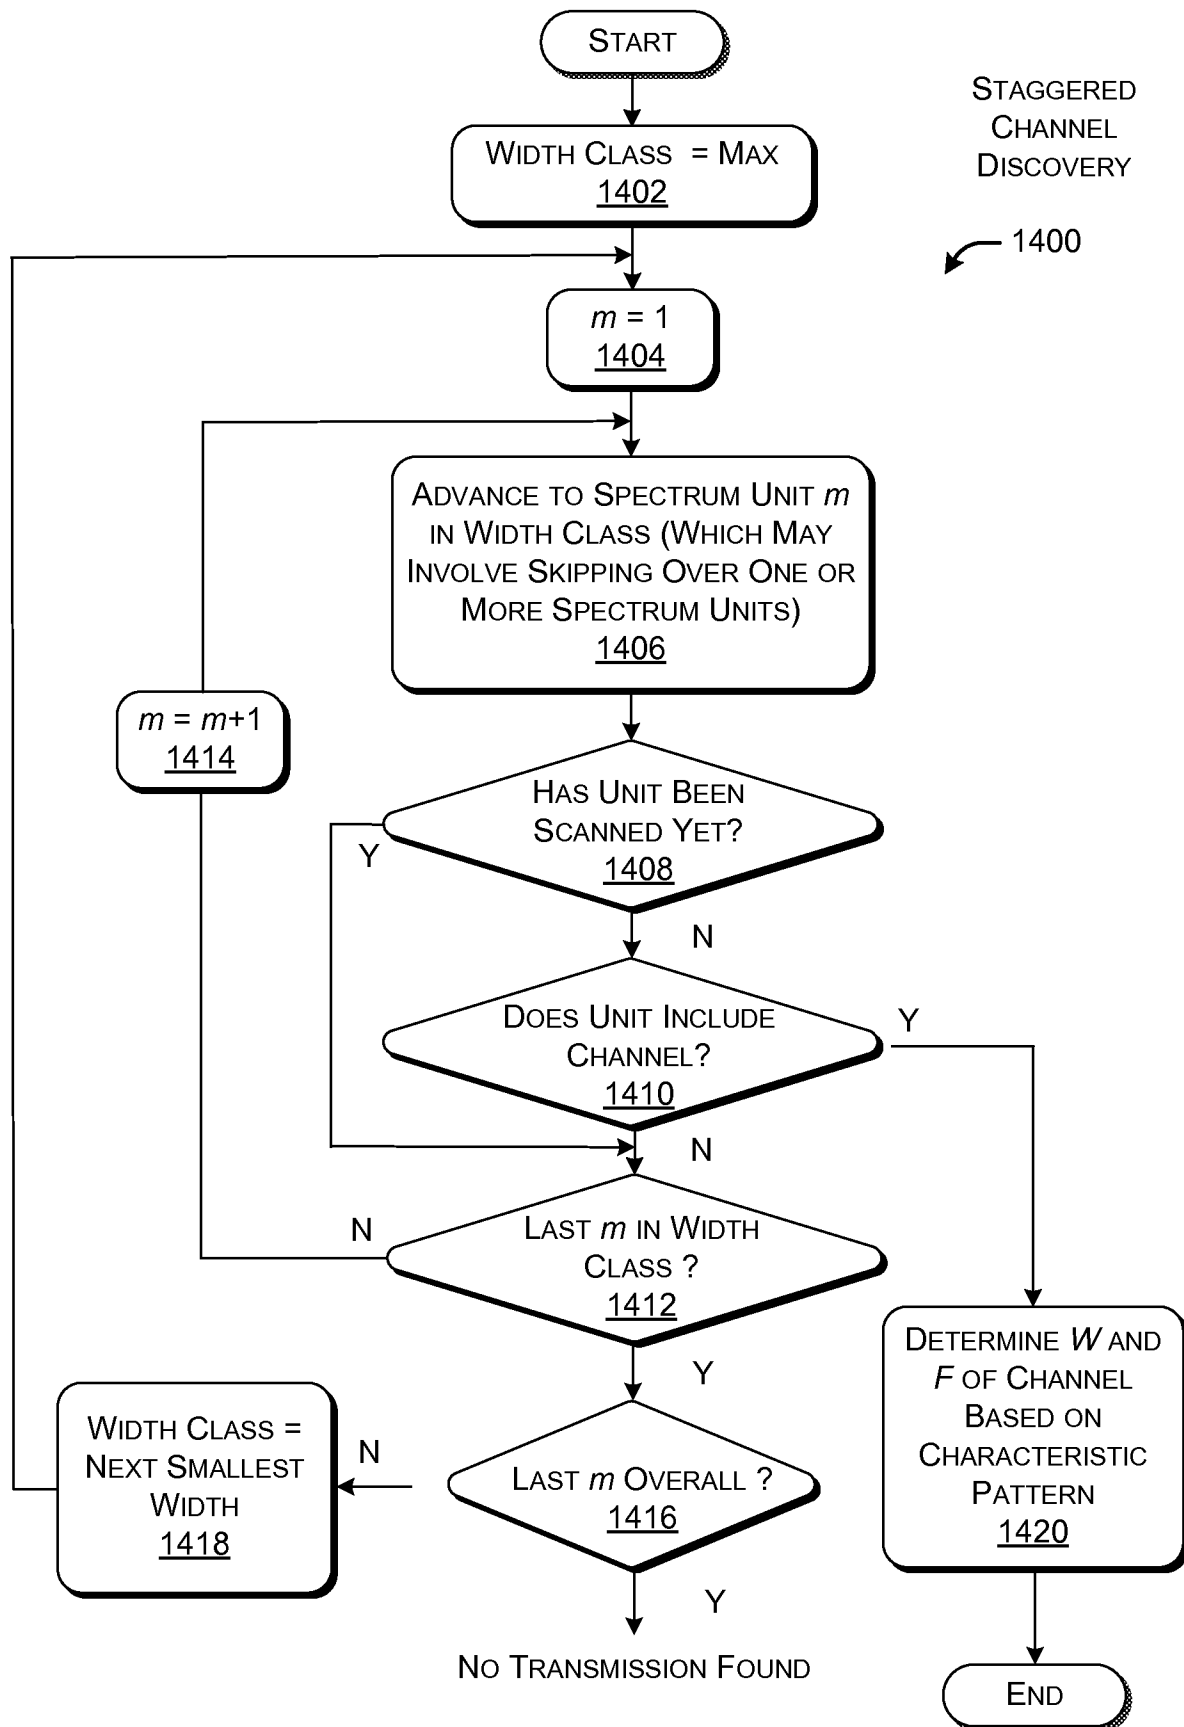
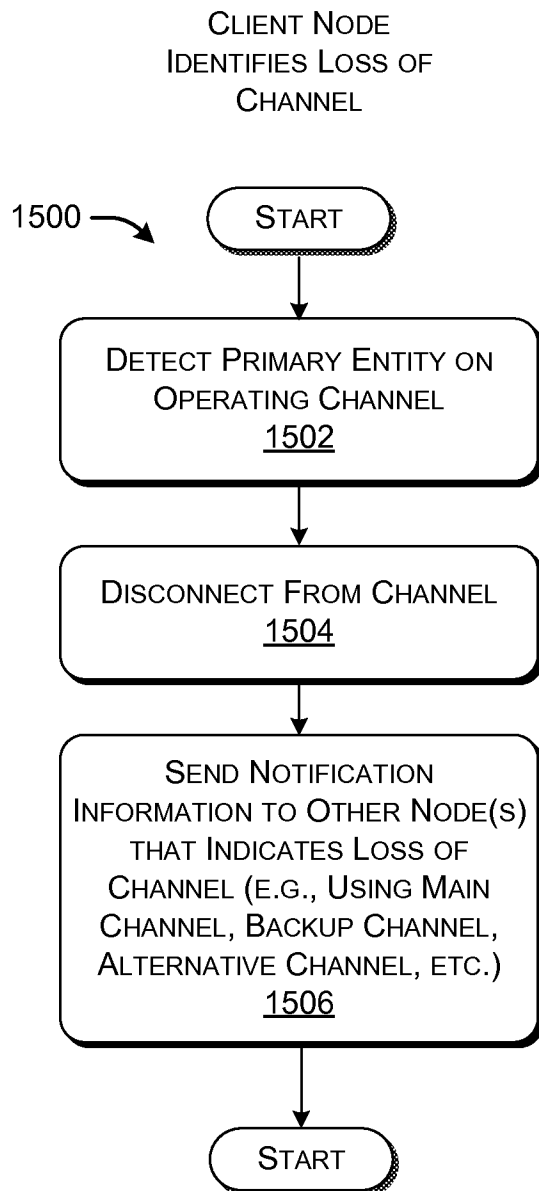
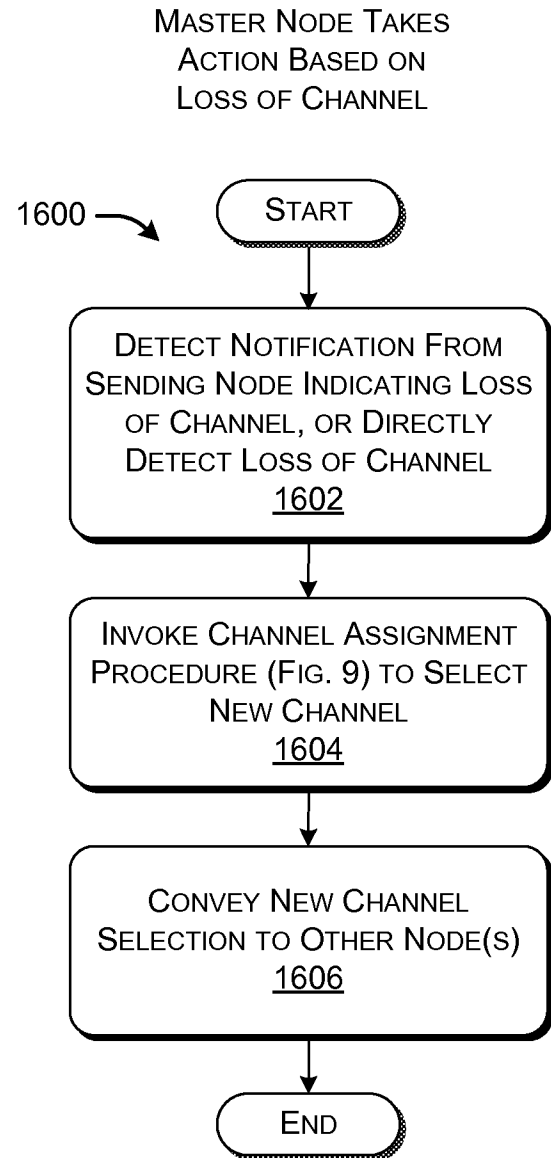
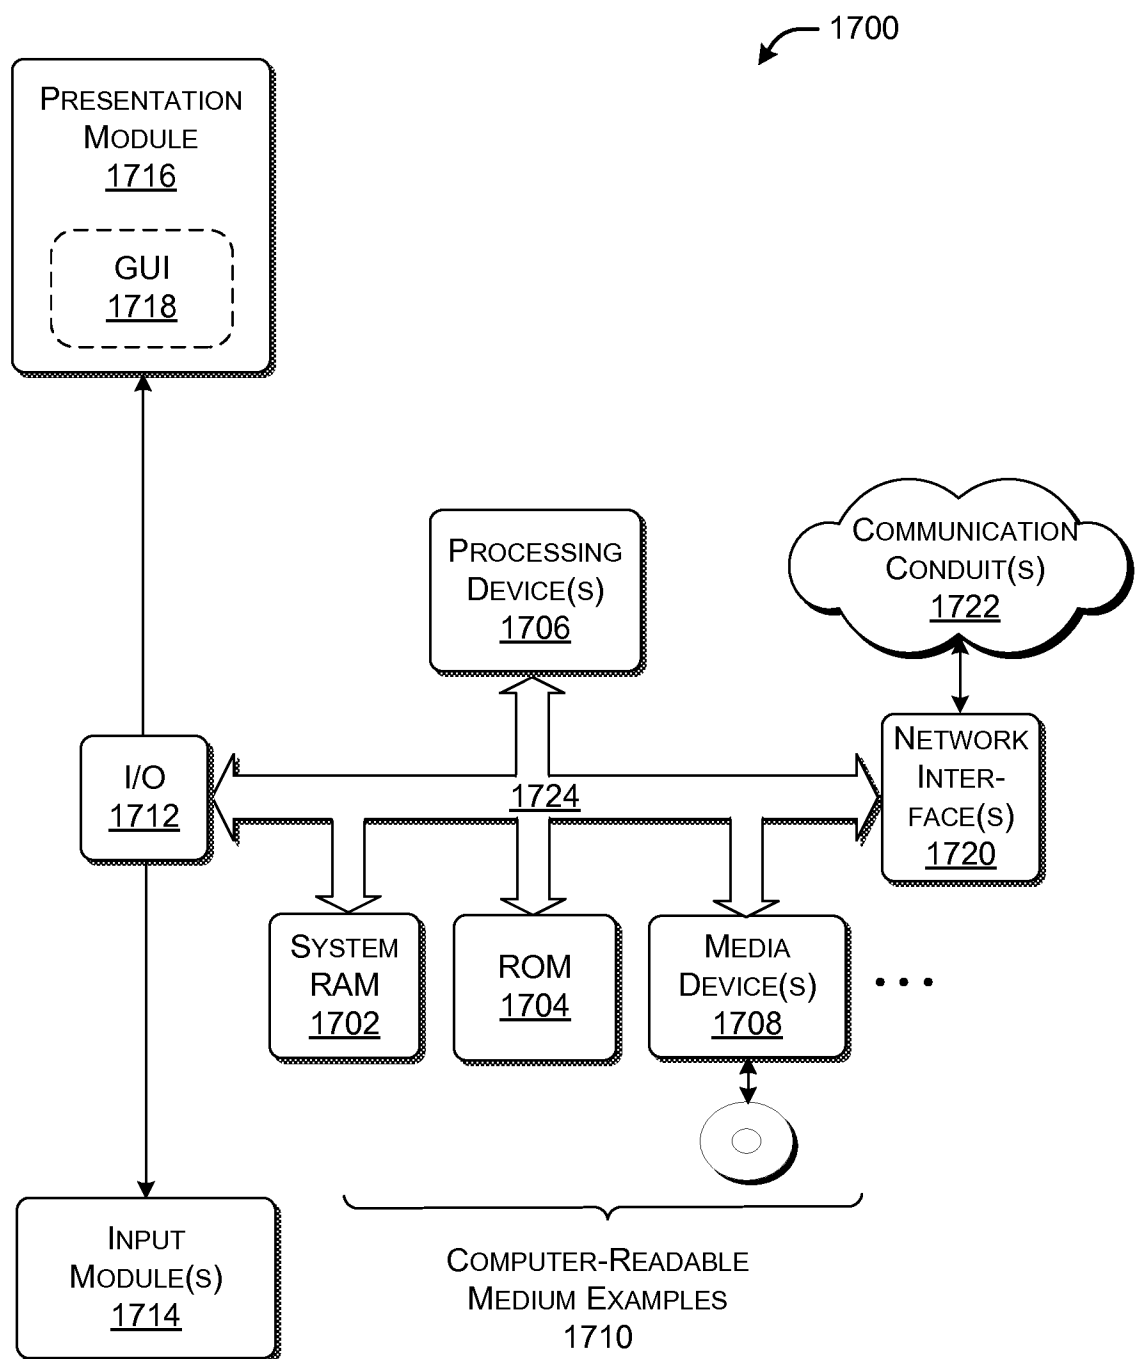


FIG. 14

11/12

**FIG. 15****FIG. 16**

12/12

**FIG. 17**