

(51) International Patent Classification:
H04W 16/14 (2009.01) *H04B 7/26* (2006.01)(21) International Application Number:
PCT/US2011/063340(22) International Filing Date:
5 December 2011 (05.12.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/964,492 9 December 2010 (09.12.2010) US(71) Applicant (for all designated States except US): **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).(72) Inventors: **HASSAN, Amer A.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **REED, Danny Allen**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **GRIFFIN, Trenholme J.**; 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, QA, RO, RS, RU, RW, 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, RW, 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, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) Title: COGNITIVE USE OF MULTIPLE REGULATORY DOMAINS




FIG. 3

(57) Abstract: A wireless communication system that concurrently communicates information in multiple regulatory domains to facilitate audio/video media streaming and other high bandwidth operations. One domain may be licensed and the other may be unlicensed. Transmission in the licensed domain may occur in white space in the domain, and the amount of information transmitted in that domain may be limited by regulations. The amount of information conveyed in the licensed domain may also depend on channel conditions in either or both of the domains. As a result, the relative amount of information transmitted in each domain may vary dynamically. The system includes a transmitter that dynamically determines weighting coefficients applied to each of a plurality of channels to set power levels in both domains to achieve a desired metric for the overall communication. A corresponding receiver assembles the substreams into a stream that can then be displayed or otherwise processed.



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))*

COGNITIVE USE OF MULTIPLE REGULATORY DOMAINS

BACKGROUND

[0001] There are many instances in today's information age in which information is communicated wirelessly. Computing devices widely are equipped with radios for the wireless exchange of information with other computing devices. Wireless communication enables a wide range of functionalities to be available through a computing device, without a requirement that the computing device be physically connected to those computing devices or even to a network through which those computing devices can be reached. Rather, mobile computing devices are known to be wirelessly connected to a network, such as the Internet, through which information can be exchanged with other computing devices. Computing devices may also be directly connected to each other or to peripheral devices, such as printers and display devices, through wireless connections. Even for devices that are primarily intended for use in a single location, a wireless connection may be provided to avoid the need to route wires in that location.

[0002] In addition, wireless communication is used for purposes beyond interconnecting computing devices. Wireless communication allows audio and video content to be broadcast to multiple receivers. For instance, television signals are often communicated wirelessly. Similarly, audio radio stations also broadcast signals wirelessly. Moreover, wireless communications play an important role in military and public safety related communications. As yet a further example, wireless communication also enables mobile phone services.

[0003] To avoid interference between many entities trying to communicate wirelessly, multiple techniques may be used. One technique separates users by frequency such that transmissions by users are made at different frequencies. A device may have a radio tuned to receive communications having a particular frequency characteristic, while excluding other communications.

[0004] To accommodate the many entities that desire to use wireless communication, the available frequency spectrum that is usable for wireless communication is divided into domains. These domains are established by government regulators such that different frequencies may be used for different purposes in different countries. However, in all countries it is generally the case that regulations establishing a domain define who is permitted to use the domain and the purposes for which the corresponding frequency band can be used. For example, in the United States, a domain is established for transmission of digital television signals. A separate domain has been

established for Industrial, Scientific and Medical (ISM) radios and another domain. Another domain, called the Unlicensed National Information Infrastructure (UNII) domain, is used for wireless communication in accordance with a widely used standard, the 802.11 standard, for connecting computers wirelessly to local area networks.

5 **[0005]** Each domain may be assigned a range of frequencies such that, even within the domain, multiple users may communicate using different frequencies. As part of the regulation establishing a domain, a mechanism may be defined by which multiple users share the spectrum allocated for the domain without interfering with one another. In some instances, a domain is licensed, meaning that a specific frequency is assigned to a specific entity, at least within some limited geographic area.

10 **[0006]** In other instances, though the domain is unlicensed, limitations may be imposed on devices intended to communicate in the domain to avoid interference. For example, maximum power output limitations may be imposed to reduce the likelihood that transmission from one user's device will interfere with another user's device, which is not intended to receive the transmission. As another example, a device attempting to transmit

15 in a particular domain may be required to sense whether a particular set of frequencies is in use prior to using those frequencies.

20 **[0007]** Despite the establishment of multiple domains, and techniques within each domain to avoid interference, there is a shortage of available spectrum for some uses. Unlicensed computer-to-computer communications and computer-to-peripheral communications are two examples of such uses. To make more frequencies available for such uses, it has been proposed to use unassigned bands of the licensed spectrum, referred to as "white space," which may exist for any of numerous reasons. For example, when frequency bands of the licensed spectrum are assigned to organizations, typically they are not assigned exactly consecutively; instead, there may be gaps between the licensed

25 portions to prevent two entities from interfering with one another.

30 **[0008]** In particular, it has been proposed to use white space in the digital television domain for computer (e.g. computer-to-computer, computer-to-peripheral) communications because this spectrum has the advantage of being close in frequency to the frequencies used for computer communications. Also, there may be white space in a geographic region because there are fewer licensed TV stations in that region than the digital TV domain can accommodate. However, because the amount of available white space varies from location to location, such use of the digital TV spectrum is done cognitively, such that a computer determines whether a frequency is already in use, before

using it. This determination may be made by a computing device sensing signals to find an unused channel or by consulting a database of licensed channels applicable in the location.

SUMMARY

5 **[0009]** Computer operations that require large amounts of bandwidth for wireless transmission are facilitated by equipping a computing device with a system that can transmit information by concurrently using multiple regulatory domains. The information may be an audio/video stream or other type of information that similarly consumes a large bandwidth. One of the domains, for example, may include the digital TV domain.

10 Another of the domains may be a UNII domain or an ISM domain.

[0010] A computing device may dynamically determine the amount of the information transmitted in each domain. This determination may be based at least in part on regulatory considerations. The amount of information transmitted in the licensed domain may be limited in part by regulation, which may limit which subchannels in
15 domain that can be used and/or the maximum power level in each subchannel.

Accordingly, the transmitter may include a mechanism to identify subchannels available for use in the licensed domain.

[0011] Alternatively or additionally, the determination of the amount of information transmitted in each domain may be made based at least in part on conditions
20 in the available subchannels in either or both of the regulatory domains. Based on the permissible power level in the regulated domain and sensed noise, a supportable bit rate may be determined for each subchannel. Appropriate error control coding and modulation schemes may be selected to provide this bit rate. The remaining portion of the stream for transmission may then be transmitted in an unregulated domain.

25 **[0012]** As the availability or usability of subchannels, channels changes, the allocation of relative amounts of information for transmission in each regulatory domain may be updated.

[0013] Accordingly, in some embodiments, transmission in multiple domains may be achieved by separating a stream of data into multiple substreams. Each substream may
30 be separately encoded using an error control encoder with a coding rate set based on channel conditions. The encoded streams may then be modulated to yield frequency components that may be transmitted at power levels, applicable to the respective domains in which each substream is transmitted.

[0014] A device receiving such a transmission may decode the incoming substreams separately. The substreams may then be combined into a single information stream that may be used to render a display, provide input to a speaker, or otherwise be processed.

5 [0015] The foregoing is a non-limiting summary of the invention, which is defined by the attached claims.

BRIEF DESCRIPTION OF DRAWINGS

[0016] The accompanying drawings are not intended to be drawn to scale. In the drawings, each identical or nearly identical component that is illustrated in various figures is represented by a like numeral. For purposes of clarity, not every component may be
10 labeled in every drawing. In the drawings:

[0017] FIG. 1 shows an exemplary system for the communication of information by concurrently using multiple regulatory domains, in accordance with some embodiments of the present disclosure.

15 [0018] FIG. 2 shows an illustrative architecture for a computing device that may be operated to transmit information in at least two regulatory domains, in accordance with some embodiments of the present disclosure.

[0019] FIG. 3 shows a block diagram of an exemplary system for transmission of information using multiple regulatory domains that may be implemented on a computing
20 device, in accordance with some embodiments of the present disclosure.

[0020] FIG. 4 shows a block diagram of an exemplary system for the reception of information from at least two regulatory domains that may be implemented on a computing device, in accordance with some embodiments of the present disclosure.

[0021] FIG. 5a and FIG. 5b show an illustrative process for the concurrent
25 transmission of information over at least two regulatory domains, in accordance with some embodiments of the present disclosure.

[0022] FIG. 6 shows an illustrative process for the reception of information over at least two regulatory domains, in accordance with some embodiments of the present disclosure.

30 [0023] FIG. 7 is a block diagram generally illustrating an example of a computer system that may be used in implementing aspects of the present disclosure.

DETAILED DESCRIPTION

[0024] The inventors have recognized and appreciated that a computing device may appear to be more responsive and generally behave consistently with user's expectations by transmitting concurrently in multiple regulatory domains. Concurrent use
5 of multiple regulatory domains may make available more bandwidth for wireless communication, and increase the speed at which operations involving wireless communications can be performed. The amount of information transmitted in each regulatory domain may be adjusted dynamically based on regulations and/or channel conditions. The allocation may be specified by a relatively small number of parameters,
10 such as a weighting applicable to the frequency components used to transmit each substream and/or a coding rate for an error control code.

[0025] Such an approach is well suited for use in transmitting a media (e.g., audio or video) stream for display on a peripheral device, such as a television. Using such a technique, a computing device is likely to be able to wirelessly transmit a media stream at
15 a sufficiently high bandwidth, for instance, to provide high-resolution images at a sufficiently high frame rate to yield a high quality audio/video presentation. Thus, a computing device may be configured to act as a media controller for streaming audio/video information to display devices, such as televisions or devices with audio speakers.

[0026] FIG. 1 shows an exemplary system 100 that may be operated to transmit information in at least two regulatory domains. The exemplary system 100 comprises a computing device 110, a television 111, a digital TV broadcaster 112, and a wireless access point 140 connected to a server 141, which is in turn connected to a regulatory database 142 via a network 130. The computing device 110 is connected wirelessly
20 through a channel 118 to the wireless access point 140, and may wirelessly transmit information to the television 111 using channels 115 and 116. In this example, channels 115 and 116 may be in different regulatory domains. The digital TV broadcaster 112 transmits information to the television 111 using channel 117.

[0027] The example system 100 may be implemented in any number of ways, and
30 may comprise alternative and/or additional elements (e.g., peripheral devices, computers, spectrum users). The computing device 110 may be any communication-enabled computing device such as a laptop, personal digital assistant, media server, or media controller and may communicate with any system component (e.g., television, wireless access point) using one or multiple wireless channels or wired connections.

[0028] The television 111 may be capable of wired and wireless communication, may be equipped with a radio, and may communicate with other devices including servers, handheld computing devices, and broadcasters. The network 130 may comprise, for example, the Internet, a LAN, a WAN and/or any other wired or wireless network, or combination thereof.

[0029] Wireless communications may be implemented using many different wireless communication technologies. For instance, Wireless Local Area Networks (WLANs) based on IEEE 802.11 standards, Wi-Fi infrastructure mode or Wi-Fi Direct connections, Wireless Personal Area Networks (WPANS), Wireless Wide Area Networks (WWANs), and Wireless Regional Area Networks (WRANs) may be employed. Technologies conforming to IEEE 802.11 standards, such as the 802.11a,b,g, and n standards may be employed.

[0030] The computing device 110 may stream media wirelessly to another device, such as a computer and/or a peripheral device. In the system 100, for example, the device 110 may wirelessly transmit information for real-time display on the digital television 111. The information may include a data stream representing streamed media of any suitable type. For instance, it may comprise audio content, video content, or some combination thereof.

[0031] The device 110 may stream data over multiple regulatory domains. This may be useful in the case that there is not enough bandwidth on available channels in any one domain to transmit the data stream at a rate to meet a user's expectations. For example, for a media stream transmitted for real-time presentation, a lack of bandwidth may reduce the data rate, and may result in quality degradation of the presented media. Audio may sound "noisy" or video may appear "grainy" or "jerky."

[0032] In the example system 100, the computing device 110 may transmit over channel 115, comprising a set of frequencies in an unlicensed regulatory domain, and channel 116 comprising an additional set of one or more frequencies in a different domain. The device 110 may determine which frequencies are included in channels 115 and 116. For instance, the channel 116 may comprise a set of one or more frequencies in the white space of a digital TV domain or in the white space of another licensed domain. Alternatively and/or additionally, the channel 116 may include frequencies in an unlicensed domain. For example, frequencies in the ISM or UNII domains—distinct from the frequencies comprising channel 115—may be included.

[0033] In order to transmit information over multiple domains, the device 110 may dynamically determine how much of the information should be transmitted over each domain. This determination may be based on any of numerous criteria. For instance, regulatory criteria may be considered. In licensed domains, such information may
5 comprise a list of assigned subchannels within certain geographic areas, and/or may comprise white spaces in these areas. Furthermore, regulations may impose limits on maximum and/or average power outputs for frequencies in licensed and unlicensed domains alike. Transmissions over channel 116, for example, may be required to be at a power level low enough that they do not interfere with transmissions of the DTV
10 broadcaster 112 over the licensed domain channel 117. Another criterion for determining how to split information across multiple channels may be channel conditions such as the presence of noise and/or other sources of interference.

[0034] The computing device 110 may use regulations and/or channel conditions to determine which frequencies to include in channels 115 and 116, to decide how to
15 divide the data stream to be sent out into multiple substreams, and to assign substreams to channels. The computing device 110 may have the regulatory requirements readily available (e.g., stored in memory) or may obtain them from an external source. For instance, computing device 110 may connect to a regulatory database 142, which may contain the appropriate regulations such as the subchannels available for use in the
20 licensed domain along with the associated power limits. In order to connect to the database 142, the device may use the wireless connection 118 to a wireless access point 140, which is in turn connected via a wired connection to the server 141. The server 141, may access the regulatory database 142 through a network 130. Other means for the device 110 to connect to the regulatory database 142 may be additionally, or alternatively defined, as
25 embodiments of the invention are not limited in this respect.

[0035] The computing device 110 may further sense and/or otherwise become aware of channel conditions to determine how much information may be sent over each channel. In the example system 100, the computing device 110 may sense signal power levels within channel 116 before transmitting over that channel.

[0036] Regardless of what information is collected or how it is collected, the
30 computing device 110 may use this information to decompose the data for transmission into portions that can be transmitted over channels, which may be in different regulatory domains.

[0037] For example, the device 110 may then transmit data over channels 115 and 116 to the digital television 111. The device 110 may employ any suitable appropriate error correcting and modulation schemes for this purpose. Because channel conditions, channel availability, and other constraints may vary over time, the device 110 may dynamically alter the channels used for transmission as necessary. Other aspects of the system may be varied as well, including the error correction and modulation schemes, coding rate, and power outputs in each frequency.

[0038] FIG. 2 illustrates, at a high level, an architecture for computing device 210 that may be operated to transmit and/or receive information in at least two regulatory domains. In the example of FIG. 2, computing device 210 includes two radios, radio 250 and radio 254. Each of the radios may be adapted to send and receive wireless communications. Radio 250, for example, may be used for wireless communication in a first domain. Radio 254, for example, may be used for wireless communication in a second domain. Though, it should be appreciated that, in some embodiments, a radio may support transmission of information in multiple domains such that a single radio may be used to support concurrent communication in multiple domains.

[0039] In the example of FIG. 2, an application 220 is illustrated. Application 220 may generate information for wireless transmission or may process information received wirelessly. In the embodiment illustrated, that information may be an audio/video stream, which may contain data representing audio content and/or video content, and possible other information, such as control information.

[0040] As a specific example, application 220 may be a media controller, configuring computing device 210 to provide a user interface through which a user may select audio/video content to be streamed to a display device, such as a television. Though, it should be appreciated that a device receiving and presenting an audio/video stream may have a similar architecture. In such an embodiment, application 220, rather than being a media controller, may be computer-executable components that receive and render the audio/video stream.

[0041] In the example of FIG. 2, application 220 interfaces with operating system 230. In some embodiments, operating system 230 may be a general purpose operating system, such as the WINDOWS® operating system. Such a configuration may be desirable when computing device 210 executes applications other than application 220. Though, in embodiments in which computing device 210 is configured specifically for

presentation of audio/video information, operating system 230 may have more limited functionality.

[0042] Regardless of whether operating system 230 is a special purpose or general purpose operating system, in the embodiment illustrated, a function of operating system 230 is to provide services that facilitate wireless transmission and reception of information processed by application 220. For transmission, operating system 230 may receive a request from application 220 to establish a connection with a nearby device. Thereafter, operating system 230 may receive from application 220 a stream of data representing audio/video content to be transmitted over that connection.

[0043] In embodiments in which computing device 210 is receiving and presenting audio/video data, operating system 230 may respond to a request, received wirelessly, from another device to establish a connection. Operating system 230 may then provide data received over that connection to application 220 for processing, including presentation of the information.

[0044] Such a connection may be formed using techniques as are known in the art. In the illustrated example, that connection may be a direct, device-to-device connection. Though any suitable technique may be used.

[0045] In the embodiment illustrated, a specific transport used for a connection may be selected by operating system 230. Accordingly, a transport control layer 260 may be included within operating system 230 to establish a transport for a connection. For example, a wireless link established in either the UNII or the ISM domain may be used as a transport. Alternatively, a wireless link established in a licensed domain, such as the Digital TV domain, may be used as a transport. Though, in the embodiment illustrated, a link established by concurrent transmissions in multiple regulatory domains may act as the transport.

[0046] Transport control layer 260 may receive information generated by application 220 defining the characteristics of a transport required for a connection. For example, application 220, upon requesting a connection may indicate a requested bit rate for the connection. Transport control layer 260 may select one or more transports that collectively provide sufficient bandwidth to support the requested bit rate. In an embodiment in which the transport is established by concurrent transmissions in multiple regulatory domains, transport control layer 260 may select one or more subchannels in each domain, and associated parameters of communication on those subchannels, to achieve the desired bit rate.

[0047] Transport control layer 260 may also respond to received information.

That received information may define a link requested by another device. Such a link may span multiple regulatory domains. Regardless of whether the link is requested by another device or an application 220 executing on computing device 210, when data is received in
5 the separate domains used to form the link, transport control layer 260 may combine data received in each of the domains for processing on computing device 210.

[0048] To send and receive data, transport control layer 260 may interact with one or more radios, of which radios 250 and 254 are illustrated. Radio 250 may be controlled through software, represented as driver 240 in FIG. 2. Here, driver 240 includes an
10 interface 242 through which operating system 230 may issue commands to driver 240 and through which driver 240 may report status and notify operating system 230 of received data. Interface 242 may be implemented in any suitable way, including according to a known standard. An example of such a known standard is called NDIS, but that standard is not critical to the invention.

15 **[0049]** Interface 242 may support a number of commands in a format that does not depend on the construction of radio 250. These commands may include commands to configure radio 250 for transmission at certain frequencies or to use certain modulation schemes or error control coding for symbols to be transmitted. Additionally, through
interface 242, driver 240 may receive data for transmission by radio 250. Accordingly,
20 interface 242 provides a mechanism through which transport control layer 260 may control radio 250 to transmit data representing a substream or to process data received from radio 250 representing a substream.

[0050] Regardless of the specific commands, driver 240 may translate the commands, in the standardized format of interface 242, into specific control signals that
25 are applied to radio 250. Additionally, driver 240 may be programmed to perform certain low level functions associated with a wireless connection. For example, upon receipt of a packet, driver 240 may check that the packet is properly formatted. If the packet is properly formatted, driver 240 may control radio 250 to generate an acknowledgement. Conversely, if the packet is not properly formatted, driver 240 may control radio 250 to
30 transmit a negative acknowledgement.

[0051] Though driver 240, and in some instances radio 250, may automatically perform low level functions associated with establishing and maintaining a wireless connection, higher level functions may be performed under control of operating system 230 or applications 220. In some embodiments, an application 220 or operating system

230 may provide a user interface such that ultimate control of wireless communication is provided by a user of computing device 210.

[0052] In the embodiment illustrated in FIG. 2, computing device 210 includes a second radio 254. While radio 250 may be used, for example, for transmitting and receiving frequency components in a first regulatory domain, radio 254 may be used for transmission and reception in a second regulatory domain.

[0053] Radio 254 is incorporated into computing device 210 with generally the same architecture as radio 250. Radio 254 is associated with a driver 244 that provides a mechanism for operating system 230 to control radio 254. Driver 244 has an interface 246 through which operating system 230 may send commands to driver 244 and driver 244 may provide status to operating system 230. Interface 246, like interface 244, maybe a standardized interface such that operating system 230 may communicate with driver 244 using a similar set of commands as are used to control driver 240.

[0054] FIG. 3 shows a block diagram of some of the components of an illustrative system 300, for transmission of information in multiple regulatory domains, implementing techniques described herein. The system 300 comprises an application 301 that may generate or process information for wireless transmission, a policy engine 302 that may determine how to transmit the information over multiple domains, and a network adapter 310 configured to accept a bit stream and prepare it for subsequent wireless transmission by a radio subsystem (320). The radio subsystem may transmit the bit stream using multiple regulatory domains. In the system 300, for instance, a first regulatory domain may comprise a first domain channel 321, and a second regulatory domain may comprise a second domain channel 322. In some instances the first domain channel 321 may be the same as the second domain channel 322, and in other instances the channels may be different.

[0055] The system 300 may be implemented in any number of ways. For instance, the system 300 or any of its components may be implemented on a computing device such as the device 210 shown in FIG. 2 or a plurality of computers. In some embodiments, the system 300 may include a copy of a regulatory database internal to such a computing device. Though, in the embodiment illustrated in FIG. 3, the regulatory database 303 may be an external source of information about channel availability and regulations governing channel use. Communications among any subset of components (e.g., between the policy engine 302 and the regulatory database 303 or between the encoder 312 and modulator 314) may be realized using any suitable means such as wired connections, wireless

connections, and networks. The radio subsystem 320 may transmit over one domain or over multiple domains using one or multiple channels. Each domain may be a licensed domain (e.g., digital TV band) or an unlicensed domain (e.g., UNII and ISM bands).

[0056] The application 301 may be any suitable content-generating application. It may be an application for generating information for wireless transmission such as the application 220 shown in FIG. 2. The content may be represented as a media stream such an audio/visual stream or any other data stream comprising, for instance, an ASCII text file or compiled bytecode. The application 301 may provide the generated content to a policy engine 302 for further processing in preparation for subsequent wireless transmission.

[0057] The policy engine 302 may obtain information which it may use to dynamically decide how to separate the stream of information from the application 301 into multiple substreams that may be transmitted in multiple domains. The policy engine 302 may determine parameters relating to transmission of these bit streams, such as the number of bit substreams to be used, which domains and channels each substream may be transmitted over, and/or what coding rate may be used for transmission over each channel. The number of substreams may be any suitable number, and, for example, maybe two substreams as shown by the two parallel tracks in FIG. 3 (and the associated two parallel tracks shown in FIG. 4, that will be discussed below). The policy engine 302 may determine a coding rate for each channel based on channel conditions, which may be either dynamically or previously determined by the policy engine 302 or another system component, such as the radio subsystem 320.

[0058] The policy engine 302 may use the regulatory database 303 to check for availability of channels in each licensed domain considered. For instance, the policy engine 302 may obtain a list of available white space channels in the digital TV domain. The policy engine 302 may use this information to select frequencies in the digital TV domain that can be used for transmission of a portion of the data stream.

[0059] In some embodiments, information on the availability of white space channels may change from time-to-time. Policy engine 302 may include a mechanism to account for such changes. Such a mechanism may entail repeating the selection of frequencies, to use for transmission, from time-to-time based on changes in availability of white space channels or channel conditions or other factors.

[0060] Policy engine 302 may obtain information with changes in the available white space channels in any suitable way. For example, the policy engine 302 may

receive a specific time period over which each channel is available. The time periods may differ from one channel to the next, and this information may be updated periodically 304.

[0061] The policy engine 302 may further obtain regulatory limitations for transmission over channels in licensed and unlicensed domains. These limitations may include limitations on the maximum transmission power in each channel/frequency and limitations on the average transmission power across all frequencies in a channel and/or multiple channels. This data may be obtained from regulatory database 303, may be preprogrammed into policy engine 302, stored as a configuration parameter of policy engine 302, and/or provided in any other suitable way.

[0062] After determining the number of substreams to split the application 301 bit stream into, the policy engine 302 may forward the data stream to the network adapter 310 along with any associated control information. Examples of control information include what channels to transmit each substream over and the associated coding rates. The coding rate may be selected based on channel conditions. Control information may further include instructions for the type of error control coding and modulation schemes to be used prepare each data substream for transmission.

[0063] The network adapter 310 comprises a multiplexer 311, two error control (EC) encoders (312, 313), two modulators (314, 315) and two weighters (316, 317). Many alternative implementations are possible. For instance, the network adapter 310 may comprise one or a plurality of encoders, modulators, and weighters. Additionally, or alternatively, the adapter 310 may include other components for implementing various signal processing tasks associated to source and channel coding. The network adapter 310 may forward any information from the policy engine 302 to any of its components. In particular, the network adapter 310 may forward the application data stream to the multiplexer 311, which may split the data stream into multiple bit substreams and forward each to an error control (EC) encoder. In the example network adapter 310, the multiplexer 311 splits the application data stream into two substreams and sends a first substream to a first EC encoder 312 and a second substream to a second EC encoder 313. The multiplexer 311 may send the coding rates associated to each substream to the appropriate EC encoders.

[0064] The EC encoders 312 and 313 may encode the data into symbols according to a selected error-control code (ECC) prior to transmission. Error control codes aid in identifying data corruption and, despite its presence, allow for effective communication. Each ECC has its own rules of construction that define valid symbols. Transmitting data in

a signal as a limited set of valid symbols may allow a receiver of the signal to identify that an error has occurred during decoding, when received data does not match a valid symbol. A receiver may then determine the transmitted data, with high likelihood, by matching the corrupted data to a valid symbol, which may be decoded to output the original data. Any
5 suitable ECC may be used by the EC encoders 312 and 313. Examples include block codes, convolutional codes, concatenated codes, low-density parity check (LDPC) codes. The EC encoders 312, 313 may use the coding rate received from the multiplexer 311 or any other system component to select a specific code and parameter of that code to achieve the desired coding rate. After encoding the data substreams, the EC coders 312
10 and 313 may forward the encoded substreams to the modulators 314 and 315.

[0065] The modulators 314 and 315 may transform the encoded substreams to yield frequency components for controlling signal transmission in the channels that were associated to each substream by the policy encoder 302. The modulators 314 and 315 may use any of numerous modulation techniques. For instance, digital modulation schemes
15 such as phase-shift keying (PSK), frequency-shift keying (FSK), amplitude-shift keying (ASK), quadrature amplitude modulation (QAM) may be employed. Alternatively, analog modulation methods such as double-sideband modulation (DSB), frequency modulation (FM), and phase modulation (PM) may be used. However, many alternate implementations of the coding and modulation steps are possible. For instance, the coding
20 and modulation schemes may be done jointly rather than in series.

[0066] The modulation may be done in any suitable components. In some embodiments, modulators 314 and 315 may include digital and/or analog components. As one example, each modulator may perform digital processing to compute a Fast Fourier Transform or (FFT) of a data substream. The output of such a circuit may be used to
25 determine frequency components to be transmitted to represent the encoded substream. Though, in other embodiments, other components alternatively or additionally may be used.

[0067] The frequency components produced by the modulators 314 and 315 for first and second bit substreams may be weighted by the weighters 316 and 317
30 respectively. The weighting may be done in the frequency (i.e., Fourier) domain or, equivalently, in the time domain.

[0068] The weighting may be accomplished by multiplying the frequency components produced by the modulators 314 and 315 by a time-variant (i.e., time-dependent) function. The time-variant function may be a set of time-dependent weights.

There may be a single weight assigned to each frequency value. The weights may depend on the regulatory limitations obtained by the policy engine 302, on channel conditions and any of numerous other factors. For instance, the weights may represent upper bounds on the total transmission power in each frequency or in a channel. In some embodiments, the weights may limit the transmission power in a set of frequencies to zero, meaning that no signal is to be sent out over this set of frequencies. Additionally, or alternatively, the weights may be such that their application to the frequency components effectively limits the average power of the transmitted signal, wherein the average is computed over the channel frequencies or a subset thereof.

[0069] In some instances, the same set of weights may be applied by the weighters 316 and 317 and a different set of weights may be used in other instances. In some embodiments, the weights may vary as a function of time and/or frequency. For instance, the weight associated to frequency f_1 and the weight associated to frequency f_2 may differ from one another. In another example, the weight associated to frequency f_1 at time t_1 may differ from the weight associated to frequency f_1 at time t_2 . Note that the weighting step may be performed at different times. For instance, it may be performed prior to modulation by the modulators 314 and 315.

[0070] Policy engine 302 may compute the weights to achieve one or more purposes. The weights may be set, for example, to ensure that transmit power in each domain is consistent with protocols and regulations applicable to that domain. As a simple example, the weights may be set to ensure that transmitted power in any regulatory domain is at or below power limits for unlicensed use of that spectrum. Additionally and/or alternatively, the weights may be set so as to increase the total data rate of transmission. In accordance with a communication theory called Shannon's Water-Filling Theorem allocating more transmit power to more reliable subchannels can improve data transfer rates. This allocation may be made based on the maximum power that can be transmitted in a domain based on regulations. Alternatively, the theories can be applied across all regulatory domains to ensure total transmit power is below a selected value. That theorem, or other suitable approaches, may be used to allocate power to frequency components by setting weights as a function of frequency.

[0071] The frequency components produced by the modulator 314 and weighted by the weighter 316 using a first time-variant function may be sent to the radio subsystem 320 for transmission on a first regulatory domain channel 321, along with any necessary control information. Similarly, the frequency components produced by the modulator 315

and weighted by the weighter 317 using a second time-variant function may be sent to the radio subsystem 320 for transmission on a second regulatory domain channel 322.

[0072] The radio subsystem 320 may transmit over one or multiple channels using any suitable means. For instance, the radio subsystem may comprise at least one radio
5 frequency (RF) transmitter (not shown) and at least one antenna (not shown) so that the encoded, modulated, and weighted substreams may be passed to the RF transmitter, applied to the antenna and wirelessly transmitted over the first domain channel 321 and the second domain channel 322.

[0073] Complementary operations may be performed as a receiver to receive the
10 data stream. FIG. 4 shows a block diagram of some of the components of a system 400, for the reception of information from at least two regulatory domains, in accordance with some embodiments of the present disclosure. The system 400 comprises a radio subsystem 420 for receiving data streams across multiple domains. A network adapter 410 accepts bit substreams received by the radio subsystem 420 and prepares a data stream for
15 subsequent use by the application 401.

[0074] The system 400 may be implemented in any number of ways. For example, the system 400 or any of its components may be implemented on a computing device such as device 210 shown in FIG. 2 or a plurality of computers. Communications among any subset of components (e.g., between the application 401 and the network adapter 410 or
20 between the decoder 413 and the demodulator 415) may be realized using any suitable means such as wired connections, wireless connections, and networks. The radio subsystem 420 may receive over one domain or over multiple domains using one or multiple channels. Each domain may be a licensed domain (e.g., digital TV) or an unlicensed domain (e.g., UNII, ISM). For instance, the first domain channel 321 may be in
25 UNII or ISM domain and the second domain channel 322 may be in a digital TV domain or vice versa. Alternatively, both channels 321 and 322 may both be in an unlicensed domain.

[0075] The radio subsystem 420 and network adapter 410 may be configured to receive one or multiple data substreams. For instance, the radio subsystem 420 may be
30 made aware of specific channels over which to receive substreams and the network adapter 410 may be made aware of what ECC and modulation scheme were employed to encode and modulate the data substream prior to transmission so that an appropriate demodulation and EC decoding scheme may be used to demodulate and decode the received substream.

[0076] The radio subsystem 420 may receive a substream on channel 321 and another substream on channel 322, and may pass the received substreams to the network adapter 410. The network adapter 410 comprises two demodulators (414, 415), two EC decoders (412, 413) and a combiner 411. The received substreams, comprising frequency components, may be passed to demodulators 414 and 415. For instance, the frequency components received on channel 321 may be passed to the demodulator 414 and the frequency components received on channel 322 may be passed to the demodulator 415. Many alternative implementations are possible, and these components may be implemented using technology similar to that used to implement system 300, though complementary operations may be performed. For example, demodulators 414 and 415 may be implemented as inverse FFTs. In some embodiments one demodulator may be used to process multiple streams. In other embodiments, the demodulators may be part of the radio subsystem 420 rather than the network adapter 410.

[0077] The frequency components may be sent from the demodulators 414 and 415 to EC decoders 412 and 413, respectively. The EC decoders may decode the frequency components into bit substreams that may subsequently be combined into a single data stream by the combiner 411. The EC decoders may further correct any errors in the received data to increase the likelihood of accurately recovering the transmitted data. The combiner 411 may send the combined bit substreams to an application 401.

[0078] The application 401 may be any suitable application for receiving and processing content. It may be an application for receiving and processing information wirelessly, such as the application 220 shown in FIG. 2. It may have similar structure to the application 301, or may be substantially different. In some embodiments, a system 300 and system 400 may be implemented in the same computing device such that the computing device can both transmit and receive data. The received content may be a media stream such an audio/visual stream or any other data stream such as an ASCII text file or byte code. The application 401 may process the received content in any number of ways. For example, the application 401 may render a display device (e.g., the screen of the computing device 110 or the television 111 both shown in FIG. 1) by using a received audio/video stream. Additionally, or alternatively, the application 401 may forward the data stream to another application or manipulate the data in any suitable way.

[0079] FIG. 5a and FIG. 5b together show an illustrative process 500 for the transmission of information over at least two regulatory domains, in accordance with some embodiments of the present disclosure. In this example, the process 500 may attempt to

transmit data over an unlicensed domain and/or white space of a licensed domain, but other alternative combinations of regulatory domains may be employed. This process may be performed by a system 300 as described in connection with FIG. 3 or any suitable components.

5 **[0080]** The process begins in act 501 with an application (e.g., the application 301 in FIG. 3) generating a request for bandwidth. The amount of bandwidth requested may be determined by the application, and may depend on the amount of data that needs to be transmitted and the time interval within which the transmission needs to occur.

10 **[0081]** In act 502, the process may receive a listing of available white space subchannels and associated regulations for transmitting across these subchannels. This information may, for instance, be obtained by policy engine 302 from a regulatory database 303 as shown in FIG. 3, or be obtained by any other suitable means. This information may further include a time interval $[t_i \ t_f]$, for each channel, over which each of the subchannels will be available. As one example, an external source providing access to
15 a regulatory database 303 may indicate that a channel is available for use as a white space channel, and may indicate a time t_f until which that channel is available. If use of the subchannel were required after time t_f , the process may once again request a list of available white space subchannels to determine whether that channel is still available. In this way, the value t_f indicates a time at which policy engine 302 may redetermine which
20 white space channels are available for use. Though, it should be appreciated that other mechanisms may be used. For example, the ending time t_f need not be expressly specified. As one possible alternative, policy engine 302 may recheck from time to time whether the allocated white space subchannels remain available. As another alternative, a server, such as server 141 (FIG. 1), or other component providing access to regulatory
25 information may periodically report on the continued availability of white space channels. In such an embodiment, policy engine 302 may monitor such reports and, if a report confirming the continued availability of a white space subchannel is not received within an expected interval, policy engine 302 may cease use of the white space subchannel.

30 **[0082]** Transmission restrictions may be used to determine an appropriate power level in each subchannel. For instance, maximum transmission power levels in each subchannel may be set in act 503. Similarly, the average transmission power level across all frequencies in a subchannel may be set accordingly. These power levels may be determined based on regulations, standards, and/or other suitable criteria. To enforce

determined power levels, weights may be computed by the process 500 and provided to other system components such as the weighters 316 and 317 illustrated in FIG. 3.

[0083] In act 504, the conditions in the available white space subchannels may be determined using any suitable means. For instance, the conditions may be sensed by passively listening on the subchannels or actively probing the channels, this passive listening, for example may be performed within radio subsystem 320 (FIG. 3) or other suitable component in a computing device executing process 500. Alternatively, the channel conditions may be obtained from another suitable source, such as a nearby computing device, an access point, or a base station.

[0084] The process continues in act 505, by computing the capacity of each white space subchannel, that is determining a supported transmission bit rate using conventional communication and information theory techniques. This calculation may depend on a number of factors including the bandwidth, maximum transmission power level, and amount of noise in each subchannel. The maximum transmission power level for each of the white space subchannels may be determined in any suitable way. As one example, the maximum transmission power level may be set based on regulations limiting transmission power in a specific domain. Though, other criteria may alternatively and/or additionally be employed. For example, the maximum transmission power level may be set based on capabilities of radio subsystem 320 or may be limited to preserve battery life in a portable electronic device. In some embodiments, the maximum transmission power level in each of the white space subchannels may be the same. Though, in other embodiments, the maximum transmission power level in each of the white space subchannels may be selected based on conditions of the subchannel. For example, Shannon's Water-Filling Theorem may be applied to allocate an available amount of transmission power among subchannels by preferentially allocating power to subchannels with lower noise levels to improve the overall transmission data rate. Other factors may be taken into account as well. For instance, the capacity of neighboring channels may be taken into account. The supported bit rate may be computed as a solution to an optimization problem such as, for example, attempting to maximize the data rate in each subchannel.

[0085] In act 506, the process may select an appropriate EC coding and modulation scheme along with an associated coding rate for each subchannel. These choices may depend on the capacity of each subchannel and may be made in any suitable manner. In preferred embodiments, these choices are made so as to achieve the smallest possible bit error rate while transmitting at or near the capacity of each subchannel.

Though, in some embodiments, other criteria may be used. As one example, the coding rate of an error correcting code and a modulation scheme may be selected to increase overall channel throughput. Multiple subchannels in the same regulatory domain may either employ the same EC coding and modulation schemes or different ones. As a result of acts 501—506, the process 500 may have obtained at least one subchannel in the white space of a licensed domain and various parameters (e.g., coding rate, power level) that may be used for controlling transmission over the at least one subchannel.

[0086] As a result of these computations a data rate that can be transmitted through the white space subchannels can be obtained. It should be appreciated that other computations, alternatively or additionally, may be performed to arrive at an estimate of the data rate supportable in identified white space subchannels. Regardless of how the estimate of the data rate that can be supported by the available white space subchannels, there may not be enough bandwidth to transmit all the data in the white space subchannels obtained according to acts 501—506 of the process 500. The process 500 may obtain the rest of the required bandwidth in an unlicensed domain. Accordingly, in act 507, the process may select a number of unlicensed subchannels for transmission. The process may further obtain regulatory information controlling the maximum power levels for transmission over these subchannels, sense or otherwise obtain subchannel conditions, determine the maximum bit rate in the subchannels, and select EC coding and modulation schemes along with a coding rate for each subchannel. As a result, the process 500 may have obtained at least one subchannel in an unlicensed domain (e.g., UNII band) and various parameters (e.g., coding rate, power level) that may be used for controlling transmission over the at least one subchannel.

[0087] Accordingly, FIG. 5 illustrates a process 500 in which transmission of a data stream is split between an unlicensed domain and a licensed domain. It should be appreciated that, though not expressly illustrated in FIG. 5A, other operating states are possible. For example, the entire data stream may be transmitted in the white space subchannels in the licensed domain. Such an operating state may arise, for example, in geographic areas where there is a low density of digital television stations. Conversely, an operating state may arise in which all of the data is transmitted in the unlicensed spectrum. Such an operating state may arise, for example, in a geographic area with a high density of digital TV stations or in an environment where there is a source of interference spanning a large portion of the digital TV spectrum. Moreover, an operating state may exist where the combined bandwidth available for transmission in all of the domains used may be

insufficient to achieve the bandwidth requested at block 501. In that scenario, lossy or lossless compression techniques may be employed to reduce the required bandwidth. Alternatively, the data transmission rate may be throttled, possibly reducing the quality of multimedia content contained in the stream if used for real time presentation of multimedia content. Though, regardless of the specific mode of operation, the process may proceed to block 508 where the system is configured for transmitting over whatever subchannels have been identified.

[0088] In act 508, the process 500 may inform a receiver of the subchannels to be used for transmission and obtains a data stream from an application to transmit in act 509.

The receiver may be a specific receiver or the process 500 may broadcast the subchannels to be employed for transmission. This information may be exchanged with a remote computing device or presentation device, such as the television 111 (FIG. 1), during a pairing ceremony or other suitable process by which a transmitter and a receiver may coordinate action. The data stream may be any suitable data stream, and may be, for example a media stream (e.g., audio/video), or a stream produced by the application 301 in the example system of FIG. 3. The process 500 may next partition the data stream for transmission across the subchannels selected in acts 501-507. In act 509, the process may partition the data stream so that all available bandwidth in the identified white space channels is used. In other words, usage of the licensed domain subchannels (e.g., in the digital TV white space) may be maximized. The manner of partitioning data into multiple substreams may depend on the bit rates in the at least one white space subchannel as determined in act 504. The remaining data may be transmitted over the selected unlicensed subchannels.

[0089] Each data substream to be transmitted over a subchannel may be encoded using an EC encoder. Each subchannel may use a different encoder or multiple channels may use the same EC encoder. In act 510, the process 500 determines whether or not each substream is in a licensed domain. All unlicensed-domain substreams are encoded using a first EC coder in act 511, and all licensed-domain substreams are encoded using a second EC coder in act 512. Any suitable EC coder may be employed including, for example, coders using the LDPC and Turbo families of codes. The coding scheme applied may be a coding scheme selected to achieve the computed supported bit rate for each of the channels.

[0090] The encoded substreams are modulated and weighted in act 514 to produce a set of frequency components for controlling signal transmission in the associated

subchannels. Any modulation schemes may be used, including, for example the digital and analog modulation schemes discussed with respect to FIG. 3. The modulation step may require a long FFT with unused spectral values zeroed out. The weighting step may employ weights that depend on the maximum allowed output power level in each frequency, and may, for example, be performed by a weighter such as the weighters 316 and 317 illustrated in FIG. 3.

[0091] In act 515, the weighted frequency components are then transmitted by a suitable radio, which may, as non-limiting example, be the radio 220 shown in FIG. 2. Aspects of the process 500 may be implemented in many different ways. For instance, the data substreams may be separately encoded, modulated, weighted and transmitted as in example system 300. Another example is that EC coding could be combined with modulation. Yet another example is that all subchannels may use the same EC coder, but different modulation schemes. Still other variations are possible.

[0092] After the data stream is transmitted in act 515, the process checks, in act 516, whether more data is to be transmitted. If there is no more data to be transmitted, the process 500 completes. Otherwise in act 517, the process checks whether the previously selected subchannels (both in licensed and unlicensed-domains) may still be used. This may be accomplished by checking whether a time interval allotted for transmission in each subchannel expired in act 517.

[0093] For each subchannel, if the time interval has not expired, the process 500 may continue in act 509 and further portions of the data stream may be transmitted on the subchannel, until the allotted time has expired and the process 500 needs to confirm whether or not the subchannel may still be used. This check occurs in act 518. If the subchannel remains available, then the process 500 may continue in act 509 and receive more data in the data stream for transmission.

[0094] If one of the subchannels on which a substream is being transmitted does not remain available, the substream allocation to subchannels may need to be altered. Any of numerous schemes for altering substream allocations to subchannels/frequencies is possible. For instance, given a first subchannel in a first regulatory (e.g. unlicensed) domain transmitting a first data substream and a second subchannel in a second regulatory (e.g., licensed) domain transmitting a second data substream, it may be that, at a certain time, the first subchannel remains available, while the second subchannel does not. In this case, one possibility is that transmission may continue only on the first subchannel (because e.g., no other licensed-domain subchannels are available). Another possibility is

that the second data substream, which may no longer be transmitted on the second subchannel, may be split into a third and fourth substream. The third substream may be transmitted on a third unlicensed-domain subchannel, and the fourth substream may be transmitted on a fourth licensed-domain subchannel. Yet another example may be that both the first subchannel and the second subchannel become unavailable, for instance, due to a combination of user interference in the unlicensed domain and regulatory limitations in the licensed domain. In this case, new subchannels in both regulatory domains may need to be selected. Still another example is that more data may need be transmitted than before and even though the first and second subchannels may remain available, other subchannels may need to be allocated for transmitting the additional data. Alternatively or additionally, the proportion of a data stream transmitted in each of multiple subchannels may be altered for reasons other than availability of the subchannels. As one example, determinations may be made of the conditions, such as noise, in each of the subchannels, which may lead to a new computed value for the supported bit rate in the subchannel as performed at act 505. It should be recognized, that a subchannel may comprise at least one frequency and, therefore, even the change of availability of the at least single frequency may necessitate a data substream to subchannel reallocation. It should be appreciated that many other allocations of data substream to subchannels are possible.

[0095] All the above examples, and many others, may be realized by the process

500. If in step 518, at least one subchannel comprising at least one frequency remains unavailable, the process may continue in act 502 and attempt to obtain any available subchannels. For instance, newly obtained channels may be white space channels obtained in act 502 or unlicensed subchannels obtained in act 507. The existing data streams may be repartitioned in accordance with the new subchannel assignments in act 509. Once all the data has been transmitted, the process 500 completes.

[0096] FIG. 6 shows an illustrative process 600 for the reception of information over at least two regulatory domains, in accordance with some embodiments of the present disclosure. The process 600 begins in act 601 by receiving subchannel assignments, which may comprise a list of subchannels in licensed and unlicensed domains that a transmitter may be transmitting data on. The subchannel assignments may have been sent to the receiver, for instance, by a computer executing act 508 of the process 500. Alternatively, the subchannel assignments may have been broadcast over a network, pre-programmed or otherwise communicated.

[0097] The process 600 may configure a first and a second demodulator in acts 602 and 603, respectively. Further, the process 600 may configure a first EC decoder and a second EC decoder in acts 604 and 605, respectively. In some instances, the first and second demodulators may be the same, while in other instances they may be different.

5 Similarly, the first and second EC coders may be the same in some instances and different in others. The configuration of demodulators and EC decoders may depend on the subchannel assignments. It may further depend on the EC coding and modulation schemes employed for the transmission of the data to be received on the assigned subchannels, so that the transmitted data may be appropriately demodulated and decoded.

10 **[0098]** In act 606, data is received on each of the assigned subchannels; whether or not each subchannel was in a licensed domain is determined in step 607. In the example process 600, all substreams transmitted over unlicensed-domain subchannels are demodulated using the first demodulator and decoded using the first EC decoder in acts 608 and 610, respectively. All substreams transmitted over licensed-domain subchannels
15 demodulated using the second demodulator and decoded using the second EC decoder in acts 609 and 611, respectively. However, the invention is not limited in this respect as each individual substream may have its own EC decoder and demodulator and/or various groups of subchannels (i.e., not just the licensed/unlicensed grouping) may share an EC decoder and a demodulator. For instance, a single modulation scheme may be used for a
20 contiguous block of multiple frequencies, of which at least one lies in a licensed domain and at least one lies may lie in an unlicensed domain.

[0099] The substreams are next combined, in act 612 and further processed in act 613. Further processing may involve sending the resultant data stream to an application (e.g., the application 401 shown in FIG. 4). It may also involve modifying the data in any
25 suitable way. Yet another example is rendering a display device, such as a television, to display the data stream, which is useful if the data stream is a media stream. The process 600 then completes.

[00100] FIG. 7 illustrates an example of a suitable computing system environment 700 on which the invention may be implemented. The computing system environment
30 700 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should be computing environment 700 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 700.

[00101] The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server
5 computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[00102] The invention may be described in the general context of a computer-executable instructions, such as program modules, being executed by a computer.
10 Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communication
15 network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

[00103] With reference to FIG. 7, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 710. Components of computer 710 may include, but are not limited to, a processing unit 720, a
20 system memory 730, and a system bus 721 that couples various system components including the system memory to the processing unit 720. The system 721, may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro
25 Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

[00104] Computer 710 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer
30 710 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program

modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 710. Combinations of the any of the above should also be included within the scope of computer readable storage media.

[00105] The system memory 730 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 731 and random access memory (RAM) 732. A basic input/output system 733 (BIOS), containing the basic routines that help to transfer information between elements within computer 710, such as during start-up, is typically stored in ROM 731. RAM 732 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 720. By way of example, and not limitation, FIG. 7 illustrates operating system 734, application programs 735, other program modules 736, and program data 737.

[00106] The computer 710 may also include other removable/non-removable volatile/nonvolatile computer storage media. By way of example only, FIG. 7 illustrates a hard disk drive 740 that reads from or write to non-removable, nonvolatile magnetic media, a magnetic disk drive 751 that reads from or writes to a removable, nonvolatile magnetic disk 752, and an optical disk drive 755 that reads from or writes to a removable, nonvolatile optical disk 756 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 741 is typically connected to the system bus 721 through a non-removable memory interface such as interface 740, and magnetic disk drive 751 and optical disk drive 755 are typically connected to the system bus 721 by a removable memory interface, such as interface 750.

[00107] The drives and their associated computer storage media, discussed above and illustrated in FIG. 7, provide storage of computer readable instructions, data structures, program modules and other data for the computer 710. In FIG. 7, for example, hard disk drive 741 is illustrated as storing operating system 744, application programs 745, other program modules 746, and program data 747. Note that these components can

either be the same as or different from operating system 734, application programs 735, other program modules 736, and program data 737. Operating system 744, application programs 745, other program modules 746, and program data 747 are given different numbers here to illustrate that, at a minimum, they are different copies.

5 **[00108]** A user may enter commands and information into the computer 710 through input devices such as a keyboard 762 and pointing device 761, commonly referred to as a mouse, trackball or touch pad. Other input devices may include a microphone 763, joystick, a tablet 764, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 720 through a user input interface 760 that is
10 coupled to the system bus, but may not be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 791 or other type of display device is also connected to the system 721 via an interface, such as a video interface 790. In addition to the monitor, computers may also include other peripheral output devices such as speakers 797 and printer 796, which may be connected through a
15 output peripheral interface 795.

[00109] The computer 710 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 780. The remote computer 780 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements
20 described above relative to the computer 710, although only a memory storage device 781 has been illustrated in FIG. 7. The logical connections depicted in FIG. 7 include a local area network (LAN) 771 and a wide area network (WAN) 773 and a wireless link, for example via a wireless interface 797 complete with an antenna, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide
25 computer networks, intranets and the Internet. While wireless interface 797 is shown directly connected to system bus 721, it is recognized that the wireless interface 797 may be connected to system bus 721 via network interface 770.

[00110] When used in a LAN networking environment, the computer 710 is connected to the LAN 771 through a network interface or adapter 770. When used in a
30 WAN networking environment, the computer 710 typically includes a modem 772 or other means for establishing communications over the WAN 773, such as the Internet. The modem 772, which may be internal or external, may be connected to the system bus 721 via the user input interface 760, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 710, or portions thereof,

may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 7 illustrates remote application programs 785 as residing on memory device 781. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

5 [00111] Having thus described several aspects of at least one embodiment of this invention, it is to be appreciated that various alterations, modifications, and improvements will readily occur to those skilled in the art.

[00112] For example, an embodiment was described in which power level is set in the regulated domain to be the maximum permitted power level. In some embodiments, 10 the power level in each subchannel in the licensed domain may be set in conjunction with the power levels used for transmission in the unlicensed domain. Such an approach may allocate more power to better subchannels, regardless of the domain, up to any limits imposed by regulation or limits on total transmission power that may be imposed by policy on the computing device. Such an allocation may be performed according to a 15 technique known as the Shannon water-filling technique. Allocating power in this manner may improve the overall rate at which information is communicated.

[00113] As another example, an embodiment was described in which a weight is applied to each of a plurality of substreams to be transmitted in a respective domain after a modulator generates frequency coefficients in the respective domains. Alternatively or 20 additionally, a weight may be applied to signals at the input of a modulator. Such an approach may be useful, for example, if the same weights are to be used in different subchannels within the same regulatory domain.

[00114] Also, an embodiment is described in which a separate modulator and other components may be used for separate substreams. In other embodiments, a single 25 modulator may be used. Such alterations, modifications, and improvements are intended to be part of this disclosure, and are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description and drawings are by way of example only.

[00115] The above-described embodiments of the present invention can be 30 implemented in any of numerous ways. For example, the embodiments may be implemented using hardware, software or a combination thereof. When implemented in software, the software code can be executed on any suitable processor or collection of processors, whether provided in a single computer or distributed among multiple computers. Such processors may be implemented as integrated circuits, with one or more

processors in an integrated circuit component. Though, a processor may be implemented using circuitry in any suitable format.

[00116] Further, it should be appreciated that a computer may be embodied in any of a number of forms, such as a rack-mounted computer, a desktop computer, a laptop
5 computer, or a tablet computer. Additionally, a computer may be embedded in a device not generally regarded as a computer but with suitable processing capabilities, including a Personal Digital Assistant (PDA), a smart phone or any other suitable portable or fixed electronic device.

[00117] Also, a computer may have one or more input and output devices. These
10 devices can be used, among other things, to present a user interface. Examples of output devices that can be used to provide a user interface include printers or display screens for visual presentation of output and speakers or other sound generating devices for audible presentation of output. Examples of input devices that can be used for a user interface include keyboards, and pointing devices, such as mice, touch pads, and digitizing tablets.
15 As another example, a computer may receive input information through speech recognition or in other audible format.

[00118] Such computers may be interconnected by one or more networks in any suitable form, including as a local area network or a wide area network, such as an enterprise network or the Internet. Such networks may be based on any suitable
20 technology and may operate according to any suitable protocol and may include wireless networks, wired networks or fiber optic networks.

[00119] Also, the various methods or processes outlined herein may be coded as software that is executable on one or more processors that employ any one of a variety of operating systems or platforms. Additionally, such software may be written using any of a
25 number of suitable programming languages and/or programming or scripting tools, and also may be compiled as executable machine language code or intermediate code that is executed on a framework or virtual machine.

[00120] In this respect, the invention may be embodied as a computer readable storage medium (or multiple computer readable media) (e.g., a computer memory, one or
30 more floppy discs, compact discs (CD), optical discs, digital video disks (DVD), magnetic tapes, flash memories, circuit configurations in Field Programmable Gate Arrays or other semiconductor devices, or other non-transitory, tangible computer storage medium) encoded with one or more programs that, when executed on one or more computers or other processors, perform methods that implement the various embodiments of the

invention discussed above. The computer readable storage medium or media can be transportable, such that the program or programs stored thereon can be loaded onto one or more different computers or other processors to implement various aspects of the present invention as discussed above. As used herein, the term "non-transitory computer-readable storage medium" encompasses only a computer-readable medium that can be considered to be a manufacture (i.e., article of manufacture) or a machine. Alternatively or additionally, the invention may be embodied as a computer readable medium other than a computer-readable storage medium, such as a propagating signal.

[00121] The terms "program" or "software" are used herein in a generic sense to refer to any type of computer code or set of computer-executable instructions that can be employed to program a computer or other processor to implement various aspects of the present invention as discussed above. Additionally, it should be appreciated that according to one aspect of this embodiment, one or more computer programs that when executed perform methods of the present invention need not reside on a single computer or processor, but may be distributed in a modular fashion amongst a number of different computers or processors to implement various aspects of the present invention.

[00122] Computer-executable instructions may be in many forms, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically the functionality of the program modules may be combined or distributed as desired in various embodiments.

[00123] Also, data structures may be stored in computer-readable media in any suitable form. For simplicity of illustration, data structures may be shown to have fields that are related through location in the data structure. Such relationships may likewise be achieved by assigning storage for the fields with locations in a computer-readable medium that conveys relationship between the fields. However, any suitable mechanism may be used to establish a relationship between information in fields of a data structure, including through the use of pointers, tags or other mechanisms that establish relationship between data elements.

[00124] Various aspects of the present invention may be used alone, in combination, or in a variety of arrangements not specifically discussed in the embodiments described in the foregoing and is therefore not limited in its application to the details and arrangement of components set forth in the foregoing description or illustrated in the

drawings. For example, aspects described in one embodiment may be combined in any manner with aspects described in other embodiments.

[00125] Also, the invention may be embodied as a method, of which an example has been provided. The acts performed as part of the method may be ordered in any
5 suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[00126] Use of ordinal terms such as “first,” “second,” “third,” etc., in the claims to modify a claim element does not by itself connote any priority, precedence, or order of one
10 claim element over another or the temporal order in which acts of a method are performed, but are used merely as labels to distinguish one claim element having a certain name from another element having a same name (but for use of the ordinal term) to distinguish the claim elements.

[00127] Also, the phraseology and terminology used herein is for the purpose of
15 description and should not be regarded as limiting. The use of "including," "comprising," or "having," “containing,” “involving,” and variations thereof herein, is meant to encompass the items listed thereafter and equivalents thereof as well as additional items.

CLAIMS

What is claimed is:

1. A method of operating a wireless computing device having at least one radio, the method comprising:

5 dynamically determining at least one frequency in a first regulatory domain and at least one frequency in a second regulatory domain;

decomposing data for transmission into at least a first portion and a second portion; and

10 with the at least one radio, concurrently transmitting at least the first portion using the at least one frequency in the first regulatory domain and the second portion using the at least one frequency in the second regulatory domain.

2. The method of claim 1, wherein:

the second domain is a licensed domain; and

15 dynamically determining at least one frequency in the licensed domain comprises obtaining information from a source external to the computing device indicating unused channels within the licensed domain.

3. The method of claim 2, wherein:

20 dynamically determining at least one frequency in the licensed domain comprises repeatedly accessing the source external to the computing device to determine whether the at least one frequency in the licensed domain remains available; and

when the at least one frequency does not remain available, subsequently transmitting data using the at least one frequency in the first regulatory domain without concurrently using the at least one frequency in the second regulatory domain.

4. The method of claim 2, wherein:

25 dynamically determining at least one frequency in the licensed domain comprises repeatedly accessing the source external to the computing device to determine whether the at least one frequency in the licensed domain remains available; and

30 when the at least one frequency does not remain available, subsequently concurrently transmitting a third portion of the data using the at least one frequency in the first regulatory domain and a fourth portion of the data using a frequency in the second regulatory domain different than the at least one frequency in the second regulatory domain.

5. The method of claim 1, wherein:

transmitting at least the first portion using the at least one frequency in the first regulatory domain and the second portion using the at least one frequency in the second regulatory domain comprises:

5 determining a transmit power level for each of the at least one frequencies in the second regulatory domain, wherein the transmit power level is determined at least in part based on regulatory constraints within the second regulatory domain.

6. A method of operating a wireless device having at least one radio, the method comprising:

10 with the at least one radio, receiving data, the receiving comprising:

receiving a first portion of the data over at least one frequency in a first regulatory domain and at least a second portion of the data over at least one frequency in a second regulatory domain; and

generating the data by combining at least the first portion and the
15 second portion.

7. The method of claim 6, wherein:

the first domain is an unlicensed domain; and

the second domain is a licensed domain.

8. The method of claim 6, further comprising:

20 adapting to a change in the at least one frequency in the second regulatory domain.

9. The method of claim 8, wherein:

receiving the second portion of the data over the at least one frequency in the second regulatory domain comprises receiving the second portion over a first set of
25 subchannels in the second regulatory domain; and

adapting to the change in the at least one frequency comprises subsequently receiving a third portion of the data over a second set of subchannels different than the first set of subchannels.

10. The method of claim 8, wherein:

30 adapting to the change in the at least one frequency comprises subsequently receiving the data over subchannels in the first regulatory domain without concurrently receiving data over subchannels in the second regulatory domain.





FIG. 1

2/8

**FIG. 2**

3/8





FIG. 3

4/8

400
**FIG. 4**

5/8





FIG. 5a

6/8

**FIG. 5b**

7/8




FIG. 6




FIG. 7